

Special Report

Managing storage for virtual environments: A complete guide

Explore the top tips and tools for virtual server storage management and get hypervisor-specific management advice.

- **Contents**
- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

To help you determine the best way to proceed with data storage management in your virtual environment, we've compiled some best practices, including an overview of the top third-party tools and data storage management advice based on an environment's hypervisor type.

Table of Contents

Managing storage for virtual environments: Best practices

Virtual server data storage management: Best practices
.....Page 4

Data storage management in a virtual environment
.....Page 6

Top management difficulties for storage for virtual environments.....Page 8

Examining storage for virtual workload management
.....Page 13

Best methods for improving storage for virtual environments
.....Page 16

Top tools for virtual server storage management

Top tools and methods for enterprise storage management
.....Page 20

Third-party tools aid storage management in VMware environments
.....Page 24

How storage capacity planning tools work in virtual environments
.....Page 27

Improve storage efficiency of virtual environments with capacity planning
.....Page 30

Using specialized tools for managing storage for virtual environments
.....Page 35

Using performance metrics to enhance storage for virtual environments

Performance monitoring for virtual server storage takes application approach
.....Page 40

• **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Using virtualization-aware storage for per-VM statistics
.....Page 45

Storage for virtual environments needs better performance monitoring
.....Page 46

Performance monitoring in VMware virtual environments
.....Page 49

Tools used to monitor the performance of storage for virtual environments
.....Page 52

Hypervisor-specific data storage management

Hyper-V sees more options for storage management
.....Page 55

Managing storage in Hyper-V and vSphere
.....Page 59

Best practices for managing storage in a virtual Citrix environment
.....Page 63

Understanding storage management with vSphere 5 implementations
.....Page 65

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

MANAGING STORAGE FOR VIRTUAL ENVIRONMENTS

Among the difficulties storage administrators can encounter when managing storage for virtual environments are capacity and bottleneck issues. The following tips and best practices can help work around these challenges and run a more efficient virtual infrastructure.

Virtual server data storage management: Best practices

It's not uncommon for storage admins to have difficulty with certain aspects of managing storage for virtual environments. Tasks such as dealing with a [growing infrastructure](#), managing storage I/O capacity and effectively using new management features can be overwhelming. In this podcast, virtualization expert [David Davis](#) gives pointers on how to get around those obstacles, and discusses some [best practices for managing storage](#) for virtual environments with SearchVirtualStorage Associate Site Editor [Sarah Wilson](#).

What would you say is the chief challenge IT managers encounter in managing storage for virtual environments?

David Davis: To me, the chief challenge in storage for virtualization at most companies where you've got an ever-increasing [workload from the virtual infrastructure](#) is just simply dealing with that. As you replace physical servers in the data center and the new server is much more powerful, you have so much more [CPU capacity](#) and so much more memory capacity. And then you upgrade the hypervisor and it becomes much more scalable and is able to run many more virtual machines more efficiently, and the density of the VMs that you want to put on your storage continues to increase much faster than your storage. The bandwidth and [I/O capacity](#) is increasing.

In many cases, at least in my experience, we would buy a large SAN and we would put it on a five-year lease because it was something we thought would be in the data center a long time, [and] it was difficult to make that storage migration. So, you planned on it being there for a long time, but you would get a new server that had 10 times the capacity as the old server, whereas the SAN is just staying the same. So, you have to closely [monitor the I/O utilization](#) and each LUN on the storage array to ensure that that virtual infrastructure doesn't saturate out the storage I/O capacity of the SAN. So, I'd say that dealing with the cost, the planning and the complexity related to staying up with the virtual infrastructure's pace -- assuming the virtual infrastructure is expanding, as it is with most companies -- is really the chief concern for IT managers.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Additionally, I think IT managers really struggle to stay up-to-date on the latest [virtualized storage advancements](#), the integrations we've been talking about and cloud storage options. When I go to speak at VMware user groups, that's where most of the questions come up -- around all these new features and capabilities related to cloud and further [integration between storage and virtualization](#).

Are there any common mistakes admins make when managing storage for virtual environments?

Davis: To me, the most common mistake for storage admins is simply not embracing the idea of virtualization. In some cases, you have SAN managers or admins that have been doing their job for many years. They might be close-minded to this idea of virtualization, and in some cases they might be opponents to changing the way that things are done when it comes to [managing a SAN](#) or selecting a storage solution. So, in some cases they can slow the progress of the IT department, and even the company, to achieving greater [efficiency through virtualization](#). I don't want to fault those storage admins and managers for trying to be pragmatic and conservative in protecting the storage infrastructure and ensuring that it's as high-performance and efficient as possible. I certainly respect their opinions and their experience, but I also encourage them to be open-minded and try to embrace this new idea of virtualization and storage, virtualization especially.

I hope storage admins will look for all the positives in what virtualization offers, which is all new ways to do backups, [new ways to do disaster recovery](#) and new ways to see the entire data center from one dashboard, one tool. No longer does it have to be multiple silos of storage, server and network. These can be broken down and can just be an efficient, fluid data center seen through one management interface. So, in the end, many of the [storage issues related to virtualization](#) are really, in my opinion, about education and communication between what have typically been the multiple silos of management in the traditional data center. So, the virtualization admins and storage admins need to work better together. They need to do whiteboarding sessions, attend the same classes, read the same books or blog posts, and really communicate about how to best make the data center more efficient by leveraging this new technology in virtualization, storage virtualization and network virtualization that's going to benefit everyone in the data center, especially the company as a whole.

What are some best practices to keep in mind when managing storage for a virtual server environment?

Davis: I would recommend to storage and virtualization admins alike to really work together to develop an [application-based approach](#) for a form of management. What I mean by that is look at the applications you support

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

and focus on providing the best performance possible for those applications. That's really why the IT department is there, to support the company's critical applications. At many times IT people, being technologists, get focused on the technical terms and the technical features and software applications. We focus on the [LUNs](#) [logical unit numbers] or the [IOPS](#) or the [VMs], and instead we need to find ways to measure whether the end users are seeing faster screen refreshers, faster data updates, anything so that they're able to do their jobs faster and more efficiently.

On a side note, why aren't IT departments called 'business technology enablement' departments? I mean, that's what they really do. IT departments aren't just about information technology, they're there to enable the business to be more efficient, to sell more products and to make the employees able to do their jobs faster.

Something else to always look for is [shared resource contention](#). This is something that [VMware Storage I/O controller](#), or SIOC, works to solve -- shared resource contention, specifically storage contention, in a VMware virtual infrastructure. So, don't get me wrong, it's not a magic bullet, but simply understanding the problem that SIOC works to solve I think would be enlightening to both storage and virtualization admins because it works to solve the most common form of resource contention in a [VMware virtual infrastructure](#).

One more thing is to make sure you map your storage and virtualization domains and create a baseline. It's always amazing to me how many large enterprises just don't have any documentation or diagrams as to how the virtual infrastructure and the storage work together. When there's a problem in the data center -- let's say an application is running slow -- they don't even know what storage that VM disk file is on, the type of storage that it is, what the expected capacity is out of it, and what the expected capacity in terms of I/Os is and how it's changed. And when that happens chaos ensues because there's no real documentation or plan. And this type of thing can even happen when you have to add a new SAN to the data center. If there's no existing documentation or diagrams or baseline, the process of [adding a new SAN](#) becomes complete chaos. So, make sure you monitor your capacity in the data center, not just in gigabytes but in I/Os because that's something that both storage and virtualization admins need to do their job. And that's really why this greater integration is happening between the storage and the virtualization -- to give those IT managers what they need, which is [virtualization-aware storage](#) and virtualization that's storage aware.

Data storage management in a virtual environment

What you'll learn in this tip: *Managing data in a [virtual server environment](#) can be challenging for data storage managers. The ratio of virtual servers*

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

and desktops to storage devices often creates performance issues and bottlenecks. Follow these top five tips to better manage your [virtual servers](#).

Virtual server environments and [desktop virtualization](#) have made it easier for users to consolidate their data, conserve energy and reduce their physical systems. However, enterprise data storage managers have run into problems with managing data in a virtual server environment. [Server virtualization](#) consolidates several servers on just a few storage devices, but that often creates bottlenecks as [virtual machines](#) (VMs) compete for storage resources. Desktop virtualization can also create bottlenecks because so many desktops run on a single host. Therefore, successfully [managing virtual servers](#) often takes a joint effort from storage administrators, as well as desktop and server administrators. To help data storage professionals make the most of these challenging environments, we've collected our top five tips on [managing storage in a virtual server environment](#).

Tip 1. Know your storage workloads. The storage workloads for virtual desktops and virtual servers can be very different. If you place a virtual machine on a host without considering its disk I/O usage, you might create instant resource bottlenecks.

To avoid this, you should have a basic understanding of how much disk I/O a virtual machine will generate based on the applications and workloads it will host. Also, high disk I/O virtual machines should be equal among both physical hosts and data resources. Too many VMs with high disk I/O on a single host might overwork the host's storage controller. At the same time, too many high disk I/O VMs accessing one storage system or LUN can also create performance bottlenecks. So even if you know a lot about your virtual machine's disk I/O workloads, you should still use [performance monitoring](#) tools to get a more in-depth look at your environment such as average and peak usage statistics.

Also, keep in mind that virtual machines may not always be on the same host. VMs are mobile and might be moved to another physical host. It's risky having a group of Exchange servers end up on the same host; it could bring the whole disk subsystem down. Many storage managers use [VMware Distributed Resource Scheduler \(VMware DRS\)](#) to balance workload among hosts, but be careful when using it; it only balances the workload based on CPU and memory usage, not virtual machine I/O disk usage. In cases where you do use VMware DRS, make sure you use DRS rules that keep specific virtual machines on different hosts.

Tip 2. Avoid intense disk I/O. Certain scenarios with your VMs may create periods of very intense disk I/O, which can create such high resource contention that all of your virtual machines will slow to a crawl. For virtual desktops this can be caused by time-specific events, like all of your users

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

turning on their desktops at approximately the same time each morning -- often referred to as a boot storm. While that kind of situation may be unavoidable, there are ways to deal with it such as by adding large cache controllers like NetApp's Performance Acceleration Module (PAM) to your storage device, or by using [automated storage tiering](#) technologies that can leverage faster storage devices like [solid-state drives \(SSDs\)](#) during periods of high disk I/O.

Other scenarios -- like [virtual machine backup](#) windows and scheduled VM activities such as antivirus scans or patching -- are controllable. Having concurrent backups running on multiple virtual machines on a host or data store can cause high disk I/O that will impact the performance of other VMs running on the host or data store. Try to schedule your backups evenly so you don't have too many backups occurring simultaneously on the same host or storage resource. You should also consider backup applications that avoid using host resources by accessing the VM data stores directly to back up VM disks. And some specific virtualization disk-to-disk backup products can shorten backup windows and allow tape backups of the disk repositories to occur afterwards without impacting hosts and virtual machines. For scheduled operations like patching and antivirus scanning, enable randomization or create staggered schedules to spread the operations over a period of time so they don't run simultaneously. You should also be careful when running disk defragmentation operations; defragmentation generates high disk I/O and can cause thin disks to rapidly increase in size.

Tip 3. Use space efficiently. It's easy to use up disk space with virtual machines, but there are ways to control and limit the amount of space they take up on your storage devices. For virtual desktops or lab-type server environments, using [linked clones](#) can save a great deal of disk space. Linked clones are similar to VM snapshots where a virtual machine's virtual disk file is made read-only and a smaller delta disk is created for any disk writes that may occur. Linked clones work by creating a master virtual disk image that's read by many VMs but all writes occur on each virtual machine's own delta disk. For example, if you create 100 virtual machines with 40 GB virtual disks, they would consume 4 TB of disk space without linked clones. If you used linked clones, however, you would have a single 40 GB virtual disk for all VMs to read from and smaller 1 GB to 2 GB virtual disks for writes -- a huge space savings.

Thin provisioning can also help manage [storage space in virtual server environments](#). It can be implemented at the virtualization layer or the storage layer. Almost all VMs are given more disk space than they usually need; thin provisioning allows you to overprovision storage by allowing virtual disk files to only take up the space they're actually using and not the full disk space they were allocated. The use of [thin provisioning](#) can greatly reduce the amount of disk space your virtual machines consume and will give you more control over costly storage capacity upgrades.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Tip 4. Avoid unnecessary I/O operations. Why generate excessive disk I/O if you don't have to? You should always try to limit the amount of disk I/O that virtual servers and virtual desktops create. This includes disabling any Windows services that aren't needed, uninstalling unnecessary applications, disabling file indexing, and limiting the amount of logging that both the operating system and applications generate. There are many other smaller things that can be tweaked and they can add up to greatly reduced disk I/O across your virtual machines. You can use end-point management tools or Active Directory group policy to help manage and control the configurations. You'll not only reduce virtual machine disk I/O, you'll reduce consumption of other host resources. Reducing the amount of unnecessary disk I/O that VMs generate is always a smart move as it allows your storage subsystem to operate at maximum efficiency.

Tip 5. Use the right storage for your workloads. Most hosts have local storage available in addition to being connected to shared storage for virtual machines. The types of storage available to your hosts will often have different performance characteristics, such as an [8 Gb Fibre Channel \(FC\)](#) SAN and a 1 Gb iSCSI or NFS storage device. Besides different storage protocols, you may have hard drives with different speeds (e.g., 10K rpm, 15K rpm) and interfaces (e.g., SAS, SATA, [solid state](#)). With so many different storage options to choose from, it makes sense to fit the VM to the right type of storage. Place less-critical virtual machines on the slower storage tiers and your more critical VMs with higher I/O requirements on the faster tiers. You can also use an automated storage tiering system like Compellent Technologies Inc.'s Fluid Data architecture or EMC Corp.'s Fast technology that moves data between [storage tiers](#) based on demand.

You can go a step further by splitting a VM into multiple disk partitions whose virtual disk files reside on multiple storage tiers according to their performance needs. One common way to do this is to create separate disk partitions for the operating system, Windows pagefile, applications and data. The faster storage tiers can be used for the data's higher I/O requirements, while slower tiers can be used for everything else. Even if you don't do that, you can still specify slower or local storage for the large virtual machine virtual swap file created for each VM and used when a host exhausts its physical memory. This also helps ensure that your virtual machine uses less disk space on the more expensive storage tiers.

Overall, simply understanding your storage environment is a crucial step in better management of your virtual server environment. And following these five tips will make managing virtual servers easier for the future.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Storage challenges in virtualized server environments

Managing storage in a [virtual server environment](#) comes with challenges in the areas of efficiency, deployment time and complexity. You have to navigate your way through virtual server sprawl, end user expectations and storage networking choices.

In this podcast interview, Andrew Reichman, principal analyst for Forrester Research, talks about the processes and tools to use to overcome these [storage challenges](#).

SearchVirtualStorage.com: What are some of the challenges with managing storage in a [virtualized server](#) environment?

Reichman: I'll point to three different challenges, and one is efficiency, the second is speed, and the third is lots of choices -- confusion and complexity. [Regarding] the first one, efficiency, we're seeing a lot of [virtual server sprawl](#); you make it really easy for people to spin up virtual servers, and guess what they do? They ask for lots of them. And if you do your storage environment in the traditional way that you did pre-virtual, you end up burning up lots and lots of capacity. So, keeping the storage environment efficient is a real big challenge, to make sure that you get the benefits you were hoping for, from building the virtual server environment in the first place.

The second one is speed. People's expectations are super high. We're in the [cloud](#) era. People want to get a virtual server right away, and they really don't want to wait for you to go buy storage and take six weeks, eight weeks, months, whatever it takes to build out storage the way you used to for big application development projects in the virtual server era. So, making sure that you can turn up those virtual servers very quickly and meet your internal customers' expectations is really, really important.

The third one is complexity. There are lots of different ways you can build storage for virtual server environments. You can use iSCSI, Fibre Channel, [NFS](#); you can do lots of different network designs and server designs and storage designs. It's just important to pick a best practice design and stick with it -- make sure you have consistency, you test it out, you understand the [business continuity](#) and disaster recovery aspects so that you can keep it safe and reliable. ... I think chargeback is a really important thing to make sure that your environment is really consistent and efficient. So, really just having the process to keep things from getting out of control is [really] important.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

SearchVirtualStorage.com: What are some storage management best practices to overcome these challenges?

Reichman: I'm going to talk about it again in the same three aspects that I talked about before. From an efficiency perspective, it's really critical to use both processes that can clean up and make sure you don't give out too much storage as well as technologies to be efficient. So, from a process perspective, really keeping an eye on what you're giving out; making sure you have a couple of different gradients -- like a gold, a silver and a bronze -- in terms of performance, reliability and redundancy; making sure you go back after the fact and audit the virtual servers you've given out and make sure they're still in use and have a process to disconnect them and put the storage capacity back in the free pool; keeping it absolutely consistent -- [a] process like that is really going to help with efficiency.

From a technology perspective, tools like [thin provisioning](#); deduplication; [and] [wide striping](#), [which] can allow you to use cheap disk rather than really high-performance storage capacity to satisfy those virtual server requests, [are] going to help you spend less. You know that you're going to start shifting capacity from disk on board the server to shared storage capacity, which is more expensive and more complicated, so you don't want to see your costs skyrocket. Thin provisioning is really critical because you're going to give out lots of gold images -- you don't want to have a custom storage allocation for every different virtual server. But the reality is that most of those virtual servers aren't going to consume all the storage capacity you give them. Thin provisioning allows those servers to pull from a common pool and to only use the actual amount of capacity that they're writing data to, rather than locking up physical storage for each of those.

Another key tool for efficiency is [snapshots](#) and clones, where you can quickly copy those and give them out. The snapshots and the clones actually [help] out with the second thing -- being quick, being fast. You can spawn off lots of gold images of a virtual server using a writable snapshot, and that's a great way to keep down the sprawl and make sure that you can turn up virtual servers very quickly and satisfy your customers' demands.

Finally, it's that process and complexity. I really encourage people to think about the whole company and build one consistent virtual server infrastructure on one storage architecture, one network architecture, one server architecture and one version of the virtual server technology that you're using. Keep it consistent. It can get out of control so quickly if you have lots of different versions and lots of different hardware floating around the environment. Really figure out what the requirements are and keep it simple, keep it consistent.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

SearchVirtualStorage.com: How much of a factor are the different storage architectures in creating storage challenges in virtualized server environments?

Reichman: Virtual servers create a lot of new challenges that we didn't have back in the physical server era. In the physical server era, things were fairly static. You had one server with one HBA that goes on, maybe, two redundant paths to storage -- and that's it. With virtual servers, things are much more dynamic. You're turning servers up and tearing them down much more quickly; the laborious process of zoning in a traditional Fibre Channel environment can really take a lot of time and it's not that well-suited to the dynamic nature of virtual server traffic. Thinking about NFS or iSCSI can make things much easier. Some of the testing results from VMware [show] that you really get the same performance out of Fibre Channel, or iSCSI, or NFS. So that's one architectural thing I would really encourage people to think about: What network protocol you want to use? Do some testing. Do some evaluation. Talk to some reference customers that are using it in a number of different ways. The reason why people needed Fiber Channel performance in the past might not be that relevant for virtual server environments, and you might be adding complexity, reducing flexibility and adding cost that's really not justified. So that's one thing [to] think about: Can you get this done with Ethernet for cheaper and faster?

SearchVirtualStorage.com: How are hypervisor vendors addressing storage management challenges?

I kind of touched on it before with that integration [discussion]. There's a lot of [APIs](#) that allow the hypervisor software to reach into the storage and call features that are built in the storage. This allows you to just simply give space out to the hypervisor and then let the virtual server administrators do self-service, build their own environment. You can build some templates and allow them to just grab the storage that they need, and that speeds things up a lot.

And that can be risky too. You're giving a lot more freedom and control to non-storage experts, and so you really want to build those templates and those gold images properly, and the tools that they're calling -- like thin provisioning, the different performance choices -- that they're well-architected from the start and that they work effectively [and] that you bought the right storage that allows for rapid provisioning and efficient usage. That's going to matter a lot.

That's really the approach that a lot of the hypervisor companies are taking -- rather than build it themselves, they're giving APIs that can call the native storage functionalities. So it's really incumbent on the storage team and the storage purchase decision makers to make sure that they're building the right

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

environment that's going to work well when it interacts with those APIs. But if you do it right, it can really make a great orchestrated environment where everything works together and you're not having lots of handoffs that slow the process down.

Data storage management for virtual workloads

IT engineers understand that storage is a valuable asset in any environment. One of the biggest problems facing virtual deployments is storage space. The issue is not the lack of storage, but data storage management. This results in wasted time, wasted space and wasted capital. This tip helps administrators to understand storage needs, as well as what they can do to monitor and manage their data center's storage capacity.

By knowing and understanding what the virtual machines (VMs) and workloads in an environment require, an IT engineer can see their storage infrastructure start to last longer and more efficiently. However, virtualization is constantly growing, and allocating storage is a daunting task. Fortunately, there are new technology advancements that can help with today's data storage management tasks, including better integration of storage and server virtualization, easier workload movement, on-the-fly formatting, better data security, superior thin provisioning and better dynamic provisioning tools.

But the underlying challenge remains. How do you manage and allocate storage with these new storage capabilities, and how can you dynamically manage workload storage requirements within a virtual environment?

Understanding VM storage requirements

Planning a deployment will help save time and money, and avoid headaches in the future. Before deploying a physical storage environment, it is essential to understand what is being done to the environment. Let's examine how to work with space requirements within a virtual infrastructure. Each environment is unique and therefore, some simple questions can help outline a data storage management plan:

1. The business needs to understand the scope of virtualization in the environment. Will it have a majority of its system virtualized, or will it run just a few VMs?
2. More users, more services and more applications require computing resources that the environment will need to accommodate. Where do you want to go with the environment, now and in the future?

Once a plan is established, the engineering team needs to understand what type of storage solution they will be rolling out. Some VMs require a set parameter for their storage requirements, while others can operate more

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

dynamically. Two schemes are seen as predominant across most virtual machine monitor (VMM) implementations: pre-allocated storage – [fat provisioning](#) – and flexible, on-demand storage – [thin provisioning](#). Fat provisioning sets aside a large amount of storage that rarely gets used fully while thin-provisioning doesn't allocate space until it's needed on a per-use basis.

Once a workload is established and all appropriate research has been done, it's time to determine how storage will be assigned to it. Even more importantly, now will be the time to determine *how much* storage the workload will need, because this is where you'll start the allocation process.

Dynamic storage distribution

Administrators can now monitor, allocate and manage their space requirements for all VMs through a hypervisor's interface. VSphere, XenServer and Hyper-V are now deployed with very sophisticated graphical user interfaces (GUIs), each providing a great deal of information. For example, an administrator can see the connected storage repository, how it is being utilized and the space requirements for each VM. Each new update to the hypervisor expands this storage-link capability to include more vendors, more features and more control over storage *directly* at the GUI level.

Keeping track of your resource pool is key to maintaining systems. Dynamic space allocation is nothing new. This feature has been available in most leading hypervisors for a few versions now. However, there are certain best practices for this data storage management tactic:

1. **Set an alarm for your space requirements.** Adding additional space is not difficult. In reality, it can be accomplished with about three mouse clicks. The challenge is knowing how much space there is to allocate, and if the data store is running low. To resolve this problem an engineer should set alarms within the hypervisor to properly manage thin provisioning. For some hypervisors, the alarm feature is new, but extremely important. These alarms can be customized to trigger alerts at certain thresholds so that an IT administrator can take the actions required to prevent an out-of-space issue. Alarms can be set to trigger when a data store has reached a percentage, or is a percent overcommitted.
2. **Document and monitor the environment.** Every major hypervisor's GUI is advanced enough that any IT engineer should be able to look at the storage repository and have a solid idea of where they stand on space. However, when working with space requirements, data storage management is a never-ending process that requires attention at all times. Running out of space is not a pleasant issue to

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

deal with and, for the most part, can be avoided by auditing and maintaining the storage environment.

3. **Keep the storage and hypervisor infrastructure updated.**

Watching over the workload is an important ongoing task, and keeping an eye on the storage hardware and hypervisor software is just as vital. New hardware and software releases promise better support and feature sets that help IT engineers manage their environments. Small changes, such as alerts and alarms, can go a long way in managing space needs.

Best practices and cautions

Every environment is unique; therefore, space requirements can span a tremendous range. However, there are some best practices and notes of caution for data storage management that every IT engineer should keep in mind:

- **Nothing is ever set in stone.** Modifying the size of a VM is very common. Some VMs cannot be changed and their space requirements are preset either by the IT manager or by the vendor. However, for the most part, a VM running in a storage pool has the ability to have its storage space modified. Administrators have the ability to add disk space as needed.
- **Always monitor your VMs.** As mentioned earlier it's important to know which resources VMs are using at any given moment. Workload management that involves watching VMs perform over time and seeing when storage demands fluctuate allows an engineer to properly distribute resources when needed.
- **Know your workloads.** Never assume that an application or workload will always run the same. With service packs, additional users and changes in the overall environment, certain workloads can require more storage at different times.

Keep your eyes on the storage prize

Allocating storage dynamically, or utilizing thin provisioning, is a great way to conserve storage resources by giving a workload only what it needs at the time, and then modifying the requirement later. There is no automatic solution for dynamically pushing out storage needs to a VM or workload. Using alerts, alarms and notifications within a hypervisor GUI, administrators can constantly monitor their storage repositories and know what they can utilize and what they cannot. While dynamically assigning storage on a per-need basis is a wonderful technology, it can potentially lead to errors, causing out-of-space conditions that could affect VM availability. But, when storage requirements are understood and managed correctly, IT engineers will become more effective at allocating dynamic space to their workloads.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

10 ways to improve virtual server storage

Virtual server storage operations can have a profound effect on the storage network. Here are 10 ways to help ensure that your storage yields its best I/O for the VMs it supports.

Enterprise virtual infrastructure will always use shared storage. It's a fact of life -- if you want to use the advanced features of VMware vSphere and Microsoft Hyper-V, all your hosts will need access to the files that make up your virtual machines ([VMs](#)). While the latest editions of VMware VMotion and Microsoft Live Migration don't require shared storage, most features like VMware vSphere High Availability and Microsoft Failover Clustering still do (and likely always will).

For the virtual infrastructure (and your critical applications) to perform well, they must be fed the virtual CPU, [virtual memory](#) and [virtual disk I/O](#) they hunger for. The monitoring of virtual CPU and memory has gotten easier as hypervisors have made more statistics available and monitoring tools have gotten smarter. However, performance monitoring for [virtual infrastructure storage](#) is still, in many cases, more of an art than a science. Certainly there are scientific statistics that are analyzed in storage performance, but the difference is that the storage is external to the hypervisor and physical server and, thus, the hypervisor (and your virtualization performance monitoring tools) doesn't typically have as much insight into [what's happening with storage](#) (but that's changing fast).

Nonetheless, there are a number of steps you can take to improve the performance of your virtual server storage. Here are 10 tips that can help you improve your storage systems' performance for your VMs.

1. Know your IOPS

The science of [tuning virtual server storage](#) is, in its simplest form, "IOPS in versus IOPS out." The drives in a LUN can provide a certain number of I/Os per second, the path in between the server and storage has a specific throughput, and the applications demand so much. While the simplest question is "Do you have enough or do you need more?," the answer isn't always easy. Many of the tools, technologies and solutions described in the following tips can help you [gain more I/O capacity](#) or find ways to reduce the I/O demand (usually by knowing the apps and working with app owners or developers).

In a virtual infrastructure, the best way to understand the typical I/O for VMs and datastores is to use virtualization-specific performance analysis tools. Almost every one of these tools will help you to "know your IOPS." Those

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

tools could be the included [VMware esxtop CLI tool](#), performance graphs inside vCenter (where IOPS are represented as *disk.numberRead.summation* and *disk.numberWrite.summation* counters) or commercial tools from third parties. Because vCenter stats must be divided by the sample time to produce IOPS, it can be more work to calculate. Third-party tools (or even VMware vCenter Operations Manager) are more costly but are much easier to use and much more intelligent.

Note that, among other things listed below, LUN configurations and RAID selection can also have a massive impact on I/O performance.

2. Know your applications

While storage and virtualization admins have many tools to help churn out more I/O, they're usually limited in the ways they can reduce the I/O demand. Also, too many IT pros are more than happy to get knee deep in the weeds of data center infrastructure while totally ignoring the applications running inside.

It's tough to be successful at truly understanding storage performance if you don't know what your most [I/O-intensive applications](#) are, what VMs they run inside, what server clusters they run on and what storage they use. Are the applications read- or write-intensive? Knowing the answers to these questions will guide you in your storage tuning efforts. By understanding the applications, and working with app owners and app developers, you may even be able to reduce the I/O workload and save the company a small fortune by not having to buy a new storage array or add high-performance drives.

If the application is already virtualized, you can use the tools mentioned above to get to know the typical I/O load over time and whether it's read- or write-intensive. If the app isn't virtualized, you may want to create a proof of concept for virtualizing the app, test the virtualized app with a realistic load and then model the production workload based on those tests. Alternatively, if the production app already uses a SAN, then you could use your SAN vendor's performance tool to measure I/O workload for the physical server before virtualizing.

3. Maximize the storage you have

One way to improve storage performance is to ensure that you're making the most of the storage you already have. You may have features that aren't even turned on, or perhaps you have features that could be enabled for a minimal cost. You could even consider [reconfiguring your storage LUNs](#) or RAID groups to gain better performance for the virtual infrastructure. I/O workloads change over time and perhaps the LUN layout and RAID groups

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

that were laid out a couple of years ago no longer apply. By rebalancing the workload across different LUNs or changing the RAID type for a LUN, you may improve performance for the virtual infrastructure as well as physical servers. In many cases, admins tend to look to new "bells and whistles" to improve performance when the source of virtual infrastructure storage performance issues are the same things that typically affect traditional physical servers.

4. VAAI or ODX

VMware's vStorage APIs for Array Integration ([VAAI](#)) allows the vSphere hypervisor to offload some storage-related tasks to the storage. For example, instead of cloning a VM through the hypervisor, that can just be done in the storage. VAAI also significantly reduces the time the disk is locked by using hardware-accelerated locking. Make sure your storage is VAAI-capable to ensure that your virtual infrastructure benefits from storage offload.

In the Microsoft Hyper-V world, Offloaded Data Transfer ([ODX](#)) is comparable to VAAI. If you're using Hyper-V, you want your storage to be ODX-capable.

To find out if your storage offers VMware VAAI, look up your SAN or NAS in the [VMware Compatibility Guide](#). To find out if your storage is ODX compatible, ask the vendor if their product is Microsoft ODX certified. Over time, more and more vendors will add VAAI and ODX compatibility to their enterprise storage arrays.

5. Limit snapshots

One of the most useful features of virtualization is the ability to [snapshot a virtual machine](#). That point-in-time picture of the VM's memory and virtual disk are useful should an application or OS upgrade go wrong or you need to test configuration changes. The snapshot is also used by virtualization backup and replication applications to capture virtual disk changes without causing downtime. However, too many times snapshots are overused, not only taking up disk space but causing poor performance for backup applications, VMotion and other storage-related functions. Snapshots should only be used temporarily and then deleted.

6. Thick provisioning versus thin provisioning

The ability to create VMs with [thin provisioned virtual disks](#) is very useful. I use thin provisioning all the time in my lab environment. However, thin provisioning in the virtualization layer and physical storage array must be carefully planned and managed. Most of the problems with thin provisioning

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

relate to simply running out of storage. However, if you don't have VAAI-enabled arrays, there can be some performance degradation when the thinly provisioned virtual disk needs to be increased in size. In that case, you may want to use [thick provisioning](#) (*lazyzerethick*, to be specific) for VM disks and then use thin provisioning in your storage array.

7. Virtualization-aware storage

Some new storage arrays are billed as "virtualization-aware." That means the storage talks to VMware vCenter, and so the storage knows what VMs are running and the path to their virtual machine disk files (VMDKs) on the SAN/NAS. Both the virtualization admin and the storage admin benefit significantly from this communication. For example, you could find out what VMs are creating the most IOPS or which VMs are experiencing the most storage latency. Some [virtualization-aware storage systems](#) can even provide per-VM snapshots and clones. You don't necessarily have to replace your storage array with one that's virtualization-aware, but it's something to consider for future projects and something to ask your current storage vendor about to see what similar features they're planning for future release.

8. Partition alignment

When older operating systems (OSes) are run in VMs without taking [partition alignment](#) into account, you can cause performance issues. However, with Windows Vista, Windows 7 and Windows 2008 (or later, such as Windows 8 and Windows 2012), this is no longer an issue. Some virtualization-aware storage can automatically align VMDKs, no matter the OS. But if your company is using older OSes, you should verify that your VMDKs are aligned with your virtual machine file system (VMFS) and SAN LUN.

There are a number of tools available to help you verify partition alignment of pre-Windows Vista and pre-Windows 2008 OSes, including the free [UberAlign](#) or the commercial [vOptimizer Pro](#).

9. Storage I/O control

With multiple virtualization hosts sharing the same storage, something needs to ensure the storage isn't monopolized by an I/O-hungry application or that more critical apps aren't starved for I/O. [VMware's vSphere Storage I/O Control](#) does this and is enabled with a single checkbox on a VMFS datastore.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

10. Storage Distributed Resource Scheduler (SDRS)

What VMware's [Distributed Resource Scheduler](#) does for CPU and RAM, [SDRS](#) does for storage. SDRS will move virtual machine disks (with no downtime) from one datastore to another (or even from one array to another) should that datastore be experiencing high latency or high disk space utilization. Make sure your storage can speak to vCenter using the VMware vStorage APIs for Storage Awareness (VASA) so vCenter can identify the capabilities of the storage for SDRS to be as successful as possible in preventing storage hiccups.

And if all else fails...

While budgets are tight for everyone, you may need to spend some money and purchase a new storage system to get higher performance for the virtual infrastructure. The virtual infrastructure's I/O workload of 50-plus VMs all hammering on disk LUNs will be very different from the previous 1:1 mapping for a server and a disk LUN. In addition to the better performance your new storage will provide, you may gain additional features such as virtualization-aware storage, hybrid storage, VAAI, VASA and more.

Virtualization is a new mindset in the data center; when VMs are consolidated, storage must be rethought and re-equipped.

TOP TOOLS FOR VIRTUAL SERVER STORAGE MANAGEMENT

When it comes to storage capacity and performance in virtual environments, storage administrators have a number of choices available to them. Some tools, such as those for automated tiering or deduplication, might be built into the storage itself and can help maintain top-notch performance. And there are third-party vendor tools that can provide reports on a per-VM basis. The links below provide examples of some of the most helpful tools available from a variety of vendors.

Top tools and methods for enterprise storage management

There are a slew of tools available designed to ease enterprise storage management in a [virtual server environment](#), whether they're products or services provided by third-party vendors or features of the virtualization hypervisor. More recently, [storage virtualization](#) and cloud technologies have been recognized for their impact on storage management. But do all of these approaches really make enterprise storage management easier? In this podcast, expert [David Davis](#) explains to associate site editor [Sarah Wilson](#)

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

whether virtualizing storage and cloud services [help enterprise storage management](#), and provides examples of some of the best tools available to storage administrators for simplified management.

What are the [benefits of virtualizing storage](#) within a virtual server environment?

David Davis: Just as server virtualization was able to reduce the number of virtual servers in the data center, these new [virtualized storage solutions](#) may be able to eliminate storage arrays completely. If they can't completely eliminate a storage array -- because that's a huge task in a large enterprise data center -- they'll be able to make those storage arrays so much more efficient when it comes to providing [storage for the virtual environment](#). What this means is that you'll be able to save money through the efficiencies that are gained. Even if you can't completely [eliminate an enterprise storage array](#), you gain benefits like [per-virtual machine \[VM\] replication](#) through your storage array because it can now understand the virtual infrastructure; you'll gain per-virtual machine statistics through the storage perspective, which is very important, and then features like super-fast cloning, snapshotting and much more.

And then there's another type of virtualized storage feature that I haven't yet talked about, and that's an exciting new group of products that are [server-side caching](#) for virtualization. They come from companies like Infinio Systems and [PernixData](#), and what these do is they take SSDs [solid-state drives] in many cases -- you don't always have to have SSD in your physical vSphere host -- but in many cases they utilize SSD to provide massive [performance increases for your VMs](#), and then of course whatever applications are running on those VMs. In effect, they give you much greater VM density on each of your vSphere hosts, and also on your storage arrays. So the idea, which I really like, is taking what you already have and just increasing the performance and scalability of it -- and that's what they call server-side caching solutions.

What about cloud services? How can those fit into a virtual server storage management strategy?

Davis: Of course, every company under the sun is talking about the cloud, and I don't want to completely throw that out just as market-speak because there are some very valid solutions. Some of them are like [second-tier storage](#) for primary data, like the TwinStrata cloud array. What it does is have a virtual appliance that you install, and you start storing your virtual machines there -- and whatever data you want can be stored on this virtual appliance. Then it begins replicating that data to the cloud and it acts as a high-speed cache for the most common data -- that data is stored on the virtual appliance. And so it's kind of like an infinite cloud array. So you have

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

unlimited capacity as it's [stored in the cloud](#). So I think that's a really innovative solution, but it's not a solution that I see most companies adopting or most companies interested in.

The more likely solutions that I think you would see companies be interested in would be [backup as a service](#) and disaster recovery as a service (DRaaS) solutions. And that's really the best place for enterprises to get started with cloud computing, at least in my opinion. If you could send your backup data off-site to the cloud instead of using some arcane tape drive that you have to wait for to write to and put in boxes and have someone come pick up to take offsite, and then worry about their reliability if the tapes should fail or get lost -- if you could just replace that with [backup to the cloud](#), it's just a much more simple and beautiful solution. And the same for disaster recovery: If you could have a [DRaaS solution](#) in the cloud so that at any time your VMs are already replicated there, you could just flip a switch and start using those VMs and that data instead of your primary data center when a disaster occurs. That's just a much simpler solution than maintaining your own off-site secondary data center, replication solutions and so forth.

I think if a company can afford those solutions, and they can test them and they're a good solution for that company, I think it's something that should definitely be considered. Just to name a few examples, backup solutions that have cloud integration come from [Veeam backup and replication](#), PHV Virtual and there are many others. And then for DRaaS, for example, you can use [Zerto](#) and VMware [vCenter] Site Recovery Manager, with a bunch of [different cloud providers](#). There are a lot of great solutions out there, so I just encourage you to do your homework and try to think past all the cloud-related marketing campaigns.

Let's talk about what products or services are available to [make enterprise storage management easier](#). What do you think are some of the most helpful tools for virtual server environments?

Davis: The tools that you use for managing [storage and virtual storage](#), in my opinion, should be the same. It should just be a storage tool whether that storage is virtualized or that storage is on a traditional enterprise storage-area network. Now, you also [have virtualization management tools](#), so you kind of have two different categories of tools there: virtualization tools and storage tools. But both of these categories of tools are getting smarter, and they are able to understand the storage if it's a virtualization tool, or the virtualization if it's a storage tool. So they're becoming more aware of one another.

One example is in these new [virtualization-aware storage arrays](#) where you can see per-VM I/O statistics. To me, that's a huge advancement because if you can make storage decisions based on [I/Os per VM](#), you'll know where to

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

place that VM, you'll know where your hot virtual machines are -- if you want to move those virtual machines into an SSD-based LUN or something like that, a [hybrid LUN](#) with greater performance or faster caching -- that's something you can do. Then in your virtualization tools, if you can quickly sort all your [VMs by IOPS](#) and see how many virtual machines are on a LUN, how many I/Os that LUN is serving, then I think you're doing well and you're using some smart tools there. Traditionally, these types of statistics that are related to I/Os per VM have been difficult to get and the enterprise has typically [turned to third-party tools](#). For example, [vCenter Operations Manager](#) has some storage views to help you keep up with [storage issues in virtual environments](#). You can also look for tools that are traditionally storage tools that support SMI-S [Storage Management Initiative Specification], which is the storage monitoring protocol, and those can talk directly to the storage and dig even deeper. They can get really deep information related to performance, utilization and capacity out of the physical storage arrays, but you also want to make sure those tools can understand your virtual infrastructure.

I think more and more of these types of management tools are becoming one -- they're bridging the gap. You have virtualization tools that use SMI-S, you have storage tools that talk to the virtualization layer, and it's just becoming a matter of selecting the right tool for your company.

What about some of the most useful [storage management features](#) supplied by prominent hypervisor vendors like VMware or Microsoft. How can those help with enterprise storage management?

Davis: Both VMware and Hyper-V support the movement of VMs, virtual disk files, from one virtual machine to another with no downtime. In VMware vSphere that's called [Storage vMotion](#), and in Microsoft Hyper-V that's called [Storage Live Migration](#). VMware has the unique features of VAAI, vStorage APIs for Array Integration, and they also have VASA, which is the vStorage APIs for Storage Awareness. And [Microsoft has ODX](#), which they call Offloaded Data Transfer. All of these are great ways for the hypervisor and the storage to work better together.

I think they're great steps in the right direction, but most of these features are relatively 1.0 features that immediately provide some benefits. But there's lots of room for improvement, and I think many more advancements will come.

[Storage DRS](#), VMware's Distributed Resource Scheduler, is also very innovative. It will automatically move running VMs from one data store to another. But I'm not really sure how many enterprises have deployed it in production because it's going to take some time for them to get a comfort level around their virtual disk moving from one array to another array during

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

the day with end users using it -- applications are actually hitting that storage while this virtual machine disk file is transitioning across the network. It's going to take some time for them to feel good about that and just click the checkbox to enable it on full automatic mode, and I'm sure it's the storage admin who's going to feel the most discomfort about doing that. Still, it's a very innovative feature and I think it's a peek at the types of [innovations that virtual storage can provide](#) us in the future. There are also advanced graphical enterprise storage management tools out there, but so many VMware admins still monitor storage, believe it or not, with tried and true command line tools like [ESXtop](#).

Third-party tools help with VMware monitoring for storage

Third-party monitoring tools provide a single interface from which an IT shop can monitor the performance of the storage and servers associated with a virtual environment based on VMware Inc.'s hypervisor.

Some [VMware monitoring](#) tools can provide insight into network traffic and drill down into the guest operating systems and the applications as well as supply detailed statistics on the storage to help companies address problems with their virtual infrastructure, according to [David Davis](#), author of [TrainSignal](#) Inc.'s vSphere video-training library, a VMware vExpert and VMware Certified Advanced Professional.

In the following podcast interview with Senior Writer Carol Sliwa, Davis discusses the advantages and disadvantages of third-party [storage performance monitoring](#) tools for VMware environments and offers up advice

What types of tools are available to help IT shops monitor the performance of the storage they use with their VMware server environments?

David Davis: In today's data centers, when you implement a virtual infrastructure using VMware, you're going to have a storage array, a SAN or a NAS, physical servers, of course VMware [vSphere](#), and then the VMware portion is monitored by [vCenter](#). When it comes to the storage, your storage vendor is going to provide you with some sort of storage [performance monitoring] tool. The limitations of those tools that you have by default, once you implement the virtual infrastructure, of course, is that they're tailored towards monitoring just that one thing. So, your storage tool is just going to monitor the storage, but it really doesn't have the insight into the virtual infrastructure. And then VMware vCenter has basic storage monitoring information and insight into the virtual infrastructure, but the reporting and the knowledge that it offers is very limited.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

How do the storage performance monitoring tools from third parties distinguish themselves from the traditional monitoring tools from an IT shop's storage vendor or VMware?

Davis: VMware released their own performance monitoring tool called VMware vCenter Operations, and this has further justified why you need third-party or extra performance and monitoring tools in the first place. You've got these storage performance monitoring tools from your SAN vendor, and of course, they only manage and monitor the performance of the SAN. Then you've got VMware vCenter, which can monitor and manage the performance of the vSphere virtual infrastructure, and it also offers some storage statistics. Then there are third-party tools, or additional tools not made by VMware, and these tools really distinguish themselves by giving you the single pane of glass, or the complete package, where you can go to a single interface and not only see performance information about the virtual infrastructure but also information about the storage area network and then tie it all together so that you can see which virtual machines are on what VMware data store and then what SAN LUN that data store is stored on and what SAN it's running on and what the performance of the SAN is. In other words, they're able to tie all these things together along with best-practices information about the virtual infrastructure, things like thresholds, that you might not know just on your own. These third-party performance tools have come up with, over time, recommendations to help you make the most out of your virtual infrastructure.

What are the main advantages of going with a third-party tool to monitor the storage performance with a VMware server environment?

Davis: Again, that single pane of glass is very beneficial, but also [some of these tools] can talk directly to the storage, using something like SMI-S [[Storage Management Initiative Specification](#)] to get additional storage statistics. Other tools have great insight into the virtual network, where they can actually analyze what type of traffic is traversing the virtual network and what virtual machine is talking to other virtual machines. Other tools can talk down or drill down into the guest operating systems and even the applications to understand what applications are running and how those applications are performing and even understand what virtual machines are using a three-tiered, client/server application -- who's talking to who and how are the virtual machines related; what hosts are they running on -- and then eventually, what storage they're running on, and how is the performance of all those things related to help you to solve the problems in your virtual infrastructure and analyze the performance of all the pieces of the virtual infrastructure over time.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

What are the downsides of using a third-party [VMware monitoring tool for storage performance]?

Davis: Like I said earlier, when you purchase your storage, you've got a tool there. When you purchase VMware vSphere and vCenter, you've got a tool there. So, the first downside would be the additional cost of purchasing a third-party tool. But, to me, even VMware with their recent release of their own performance tool [justifies] that you do need your own performance tool to monitor a virtual infrastructure. So I think that ... everyone has agreed over time that [the additional cost is] well worth it. You want to make sure that the tool that you find is best suited for your needs. Of course, I mentioned how certain tools have insight into different things -- the network or the applications. And your tool may not go as deep as the tool from your storage vendor or as the VMware-specific tools do in certain areas. For example, if you wanted to go and modify or analyze the multi-pathing that a host takes to access a storage array through the Fibre Channel fabric, these third-party tools probably aren't going to go into that. So you'd still need to use your vendor's tools to do something like that. Or they might not analyze or allow you to modify the queue depth on a host bus adapter. And so, there are some very deep and very detailed areas that a third-party tool probably isn't going to go into.

What are the top three pieces of advice you would offer to VMware users on storage performance monitoring?

Davis: The first one would be: Make sure you take into account that virtualization is all about increasing the utilization of your physical servers and increasing the utilization of your storage area network. So while you're trying to increase those utilizations to get the most out of your return on investment, you also want to make sure that you don't overutilize those resources and end up in slow application performance and even potentially downtime for end users. So, make sure that you have some kind of performance tool that can show you everything that's going on in the virtual infrastructure, including the servers and the storage, to help make sure that you get the most out of what you're doing without overutilizing it. Also, make sure that your tool can take into account the advanced features of virtualization -- things like thin provisioning, Storage I/O Control, Storage VMotion and so on. Third, make sure that you have a smart tool. You don't just want a tool that has some pretty charts and graphs. You want to make sure you have a tool that really knows the best practices of a VMware virtual infrastructure and can help make recommendations to ensure that you get the most out of your virtual infrastructure.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Storage capacity planning tools: Why virtual machines change everything

Put simply, storage capacity planning is the process of understanding how much storage is available, where it is, how to divide it between the various applications and users that require it, and how to maximise its utilisation.

With the growth in importance of storage and its growing technological diversity, along with the increasing demands being placed upon it, especially by virtual machines, the need for [storage capacity planning](#) has never been greater. The alternative is to wait until problems arise, but that can mean spending far more time firefighting and implementing poorly thought-out quick fixes.

Virtualisation complicates the picture. A physical server has natural physical limits; to increase its capacity you need to buy more memory, storage and [CPU](#), or a whole new server. A virtual server has none of these limitations, and its needs can quickly scale up and down as demands change. This rapid scaling means that VMs add complexity from a capacity planning point of view, whether storage, CPU or memory, and so make the process of planning and growth prediction harder.

This is especially true in an orchestrated data centre, where VMs move from host to host as resource needs and availability change. In such cases, the VM management system will need to know how much and which types of storage are available, as well as whether the subsystem meets certain criteria, such as high availability and data protection.

As an example, virtualisation admins often need to know how many more VMs storage systems can support. The storage manager will then need to model different scenarios and scale them to match likely real-world workloads. This will involve categorising or profiling VMs, rather than just averaging capacity usage across all VMs. Profiles might be email, database, application or Web server, for example, reflecting the fact that a VM hosting a database will manifest different growth and usage patterns from one providing email services.

To start the planning process, gather data from VMs over a period of days or even weeks to ensure that you have captured underlying trends. You -- or your storage planning tool -- should collate data on CPU, memory and storage usage, and network and storage throughput and I/O. Quantifying and analysing how different types of VM use storage helps to answer questions about capacity and performance. For example, trending the data allows you to predict how much longer existing storage capacity can support existing

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

and expected levels of demand and when you will need to acquire more or reallocate capacity.

Storage capacity planning can help you decide not just how much but what sort of storage you will need and where best to deploy it. Profiling capacity, throughput and I/O loads helps with understanding and determining what protocols, redundancy protection and array features you need to implement. Identification of issues could point towards the need, for example, to change drive types or RAID configuration used, to reroute some data, or simply to add disk capacity at a faster rate.

Further steps you should take include developing warning and critical thresholds, which define the upper levels of I/O and capacity utilisation at which performance starts to become constrained. These will vary, of course, by workload and application.

Positive steps

Modern storage technologies can help reduce the cost of the future storage requirements that a capacity planning process predicts.

For example, [automated storage tiering](#) can help to reduce the demands of high-speed storage not just by managing application data but also by ensuring that only appropriate elements of a VM reside on the first tier. Page files change constantly and are very frequently accessed so their virtual containers can reside on Tier 1 SAS drives or even Tier 0 SSDs, for example, while OS images -- whose I/O requirements do not greatly affect overall performance -- might live on Tier 2 SATA disks.

Meanwhile, [data deduplication](#) can reduce overall volumes of data, especially OS images, where there will be a high degree of commonality across maybe dozens or even hundreds of VMs.

[Thin provisioning](#) means that storage is only allocated when required rather than being allocated upfront. In addition, a good planning tool should be able to reclaim unused disk space, even space that is thinly provisioned but unused. An example of that would be space provisioned using a standard VM template with a large default disk size but which one or more VMs have not actually claimed. Such volumes could then be resized and the capacity reallocated.

Storage [capacity planning tools](#) should be able to aid all these processes by modelling multiple what-if scenarios, taking into account not just today's usage patterns but also future deployments of new technologies during the next budget cycle.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

The key benefit of storage capacity planning tools is that they should help reduce or contain costs by avoiding overprovisioning, and providing actionable intelligence that is easier to use than building spreadsheets, which, in a virtual environment containing hundreds of servers, can quickly grow to become large and complex.

Capacity planning tools

Aptare StorageConsole Capacity Manager offers views across storage systems including utilisation per array and [LUN](#). It can highlight LUNs that have exceeded an allocated storage threshold and are in danger of running out of storage, and it can forecast storage array capacities based on current and historical storage allocation.

Dell EqualLogic SAN Headquarters (HQ) aims to identify performance bottlenecks and help with planning by identifying storage growth requirements. It also provides consolidated performance and event monitoring. This product only works with Dell EqualLogic SANs. All other capacity planning tools from the storage system vendors mentioned below claim to work with other makers' arrays.

EMC Ionix ControlCenter StorageScope offers capacity utilisation reporting and trend analysis across heterogeneous storage including VMware and virtually provisioned storage. It allows you to plan for future capacity requirements using customised reports, and it will identify and reclaim stranded and underutilised storage assets.

IBM Tivoli Storage Productivity Center's features include automated system discovery, provisioning, configuration management, performance monitoring and replication. It provides storage metrics and analytics, including data on storage capacity, utilisation and performance, with trend analysis to help identify application workload contentions.

Microsoft System Center Operations Manager 2012 offers wide capabilities for Microsoft shops, including resource utilisation optimisation from a number of management packs from Dell, HP, NetApp, Quest and vKernel, and reporting of resource utilisation, which then allows you to trend the data.

NetApp OnCommand Balance is designed to optimise infrastructure performance and capacity. It uses modelling and analytics to provide an understanding of how application workloads, utilisation levels and resources interact, and it provides reports to aid future planning, such as predicting when particular servers will run out of capacity.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

SolarWinds' Virtualization Manager (acquired from Hyper9) performs capacity planning via a dashboard, supports Hyper-V and VMware ESXi/vCenter, and allows you to predict future bottlenecks. It also helps to determine how much it would cost to run all or part of your virtual deployment in the public cloud and provides reports on usage by line of business and departments.

SolarWinds also sells Storage Manager, Powered by Profiler. Using a dashboard and a drill-down approach, it offers a range of features. Its storage capacity planning module automates data collection with the aim of making it easy to view storage growth rates, project when storage capacity will be reached and forecast costs. It can track performance over time to identify hot spots, peak hours and potential outages.

Virtual Instruments' VirtualWisdom hardware-based system provides real-time monitoring and analysis, helps to identify and reduce overprovisioned SAN ports, and provides data on latency and throughput per port.

VMware View Planner is aimed at customers [planning VDI rollouts](#). It consists of a virtual appliance that simulates workloads for different user types. It does this by reproducing typical actions of a Windows end user running productivity applications. As you would expect, it integrates with ESXi, vCenter and VMware View.

New storage capacity management tools can make efficiency a reality

Poor provisioning and a lack of storage capacity management tools leads to underused storage systems. New tools and improved processes can make storage efficiency a reality.

Storage managers rarely admit they have a [capacity management problem](#). Instead, they're more likely to talk about how big a slice of their IT budget storage eats up or the unpleasantness of unplanned purchase requests. In some cases, the conversation focuses on the high cost per gigabyte of storage.

Other managers may be preoccupied with seeking a solution to seemingly unattainable backup windows or impossible disaster recovery scenarios.

Some are looking for capacity management tools or processes that can identify and prune obsolete data, while others are [buying storage in large chunks](#) annually to get "quantity discounts."

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

What do all of these scenarios have in common? In each case, storage managers are trying to address a symptom without taking a holistic view of a fundamental problem: the lack of an [effective storage capacity management regime](#).

Don't look to the cloud for answers

Let's state up front that cloud storage is not the solution to a capacity management problem. Increasingly, cloud is portrayed as the cure-all for what storage ailments are afflicting companies. Cloud may mask the pain with a somewhat lower cost per GB, but it does nothing to fundamentally [address uncontrolled capacity expansion](#). Cloud has a role in storage service delivery, but [solving capacity problems](#) isn't one of them.

It would be charitable to say that some organizations' [storage utilization](#) is less than stellar. Many companies have as little as 20% to 30% average utilization as measured by storage actually consumed. Those organizations whose consumed utilization is more than 50% are the exception. This metric is one of the fundamental obstacles to better utilization.

There are three basic ways to measure storage capacity:

- **Formatted** (sometimes referred to as raw, though there is a technical difference)
- **Allocated** (sometime expressed as provisioned)
- **Consumed** (or written)

When asked what their utilization rate is, most storage administrators will quote the allocated figure. From their perspective, if it's allocated to an application, it's as good as consumed because it's unavailable for [new provisioning](#). It's a legitimate perspective, but it can cover an insidious incentive to overprovision because it allows that portion of storage to be ignored for a long period of time. Some administrators will tout an 85% utilization rate, even though perhaps only 20% of the array has actually been consumed. Such poor utilization, however, ultimately drives up the average cost per GB consumed by 2x or more with management none the wiser. Moreover, most capacity purchases are triggered when allocated capacity hits 85% regardless of how much is really being consumed. Responsible teams husband an organization's resources more diligently.

Why is data getting so big?

The biggest driver of storage growth is "secondary" data, copies of original data or primary storage. Secondary data includes snapshots, mirrors, replication and even data warehouses. The secondary data multiplier can be

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

as high as 15:1. It would seem the obvious solution is to reduce the number of data copies, which may indeed be the case. However, the secondary copies were likely created for a reason, such as for data protection or to reduce contention for specific sets of data. The unintended consequence of [optimizing storage capacity management](#) may be reduced data recovery capabilities or worse performance. Thus, storage managers must be aware that there's an inverse relationship between data recovery, performance and capacity management; if you improve one, you're likely to impede the other. Consequently, it's important to start with service-level requirements for recovery and performance. Capacity management can be optimized only to the point that other service levels aren't jeopardized.

Tools to take control of capacity management

Thin provisioning

- Eliminates overallocation and increases utilized capacity from 30% to 60%
- Cuts the cost per gigabyte (GB) stored by 50%

Compression

- A 2:1 compression allows twice as much data in the same array, for another 50% reduction in cost per GB stored

Deduplication

- A 2:1 deduplication rate further halves the cost per GB of storage and the deduplication rate could be higher for some data types

Storage resource management applications

- Manages storage as an enterprise, not as individual arrays
- Measures storage metrics to drive best practices
- Spots trends that could become serious problems without proper attention

Capacity management toolkits

Fortunately, storage managers have numerous [tools to assist them in tackling capacity management](#). These include two general categories: utilities and reporting tools. Array vendors have a number of useful utilities that are now available with most systems.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Perhaps the most common of these is [thin provisioning capability](#), which is supported by nearly every vendor. Thin provisioning allows administrators to logically allocate storage, but automatically keeps the physical allocation only slightly above the actual capacity used. Storage is automatically allocated from a common pool as a volume demands more space. Because the array itself may be logically overallocated, it's possible to have an out-of-space train wreck if administrators don't ensure that enough physical capacity is available as data grows. This is uncommon, however, as automated alerts should keep administrators on top of the situation. Thin provisioning alone can largely alleviate the problem of high allocation/low utilization. In most cases it's complemented by a space reclamation feature that returns unused space to the common pool. While array vendors may offer this feature, reclamation can also be performed by Symantec Corp.'s Veritas Foundation Suite for those who use that product.

Another useful and near-universal utility is compression. Most vendors are willing to guarantee a 2:1 compression on primary storage, or a 50% space savings. Compression is normally applied at the LUN or volume level, depending upon the vendor's specific implementation. Compression does incur some performance penalty, though it can be as little as 5%. Of course, your mileage may vary, so a proof of concept is worth the effort. From a management standpoint, the benefit of compression is cutting the cost per GB stored by 50%.

[Compression is complemented by data deduplication](#), though deduplication is not yet supported on primary storage by every vendor; EMC Corp. and NetApp Inc. are examples of vendors that do. Here again, deduplication differs in its implementation on primary storage versus backup appliances. On primary storage, data deduplication is an idle-time process and isn't nearly as aggressive in eliminating duplicate blocks as deduping backup appliances. Because it's a background process, the compression itself doesn't impact operations. Decompression, known as "rehydration," may have minimal or significant effect on performance, so a proof of concept is advised. Rehydration is more like reassembly of parts. Unlike compression where vendors make efficiency guarantees, there are no such guarantees with deduplication because it's highly dependent upon data type. Media files generally dedupe poorly, whereas text files may dedupe quite well.

Capacity management reporting tools

The other category of tools is reporting tools, or more accurately, [storage resource management \(SRM\) products](#). Both array vendors and independent vendors offer SRM products, examples of which include EMC ControlCenter, Hewlett-Packard (HP) Co.'s HP Storage Essentials, NetApp OnCommand Insight (formerly SANscreen) and Symantec's Veritas CommandCentral Storage. All of them offer the ability to comprehensively manage and monitor

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

an enterprise storage environment. Yet few organizations leverage them, largely because SRM has gained a reputation as being unwieldy and resource-intensive. These limitations can be overcome by focusing on only those aspects of an SRM application that are truly beneficial, otherwise known as the 80/20 rule. In the [context of storage capacity management](#), you should focus on the following:

- **Thresholds.** Individual arrays provide threshold alerts, but SRM applications can consolidate them and give an enterprise-wide picture to administrators. This allows far more comprehensive planning and provisioning to prevent one array from being oversubscribed while another is undersubscribed, for example.
- **Utilization.** Again, SRM consolidates information that otherwise must be manually aggregated (and who has the time to do that?). Utilization figures to monitor include:
 - **Consumed as a percent of raw.** Know how much the array is truly utilized. Target 55% or higher as a best practice, though this will vary with the age of the array and growth rates.
 - **Consumed as a percent of allocated.** Know whether or not the array is overallocated. Target greater than 70% (85% if thin provisioning is used) as a best practice. Allocations lower than 70% may be acceptable for newly provisioned LUNs or those with high, unpredictable growth.
 - **Secondary data.** Know how much data is consumed by snapshots, mirrors and the like. Target no more than 3x the primary storage. More than 3x may be justifiable for various reasons, but this ensures that space isn't consumed unnecessarily. This feeds into data/information lifecycle management.
- **Trends.** Thresholds and utilization are points-in-time. Identifying trends is the key to optimizing capacity.
 - **Growth rates.** Knowing growth rates fosters accurate forecasting, thereby avoiding unnecessary "safety factor" purchases. Storage prices decline approximately 10% per quarter on a per-GB basis, so delaying an organization's purchases can yield substantial savings over time.
 - **"Days storage in inventory."** Using growth rates, calculate how many days of storage growth capacity is on the floor. Target 90 to 180 days. Less than 90 days doesn't give purchasing enough time to do their job most effectively. More than 180 days and you could have purchased the storage later at a cheaper price.

Organizations can dramatically cut the cost per gigabyte stored by using the array utilities that in many cases are already paid for. Implementing thin provisioning, compression and deduplication (where applicable) can reduce

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

this cost by 50% to 75%, which isn't bad by any measure. However, best-organizations will implement SRM products to take their [storage management](#) to the next level. With it, storage managers can balance and optimize performance, data protection and capacity utilization simultaneously.

Server virtualization and storage management: Better integration, specialized tools

In the past decade or so, [storage management](#) has developed into a discipline in its own right, driven by big increases in the amount of data being stored, as well as the rise of storage networking protocols that facilitate shared storage.

Virtualization, meanwhile, has become the foremost technology for server and PC optimization. In this environment, shared storage provides functionality that would otherwise not be possible, such as nondisruptive virtual machine migration.

But virtualization adds a layer of complexity in the association between a server and the storage that supports it. That layer of abstraction between [virtualization and storage](#) means it's a challenge to translate storage-focused concepts such as RAID groups and LUNs to virtual objects such as VMDKs and [virtual hard disks](#). So, to successfully provide storage for virtual environments, storage admins should take a new tack.

The challenges

Virtualization creates new operational headaches. Because many [virtual machines](#) (VMs) can exist on a single storage LUN, the I/O profile of virtual servers and desktops tends to be more random and unpredictable in nature. The functionality of today's hypervisors enables large amounts of I/O to be generated when moving virtual machines around the storage infrastructure through the use of features such as VMware Inc.'s [Storage vMotion](#) and Microsoft Corp.'s [Hyper-V Live Migration](#). Virtualization may also impact heavily on storage utilization as virtual machines are copied, cloned or otherwise replicated across the environment.

In considering virtualization and storage, we must examine the operational structures that have been built up in many large organizations. As IT infrastructure has grown, the component technologies have tended to split into silos covering the disciplines of storage, networking, servers and databases. Once, it was possible for storage administrators to go about their business with little regard for the operation of other parts of the infrastructure.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

But virtualization has changed that world and made it necessary for those isolated silos to integrate like never before.

Choosing a strategy

Efficient storage management in virtual environments entails meeting two basic metrics: capacity and performance. While this could also be said of nonvirtualized environments, performance is the primary consideration in virtual storage designs as it has more of an impact on the operation of a virtual infrastructure. Slow response times from a single LUN are likely to affect only a single host in nonvirtualized environments; however, poor responses from a large LUN supporting many virtual machines can have a much wider impact. This is especially so with a virtual desktop infrastructure (VDI). There are a number of strategies a storage administrator should consider.

Best practices: VMs and storage

HERE ARE SOME suggested best practices for managing storage for virtual server environments:

- Configure storage for performance, then capacity.
- Use performance-enhancing features such as wide striping.
- Deploy storage that uses the hardware acceleration via APIs.
- Don't use traditional backup methods for virtual machines; look for snapshots and third-party software that use hypervisor backup APIs.
- Use vendor plug-ins to management frameworks such as VMware vCenter and Microsoft System Center Virtual Machine Manager (SCVMM).
- Consider custom solutions.
- Move toward automation.

Use hardware acceleration and APIs

Many vendors (including the top six storage vendors: Dell Inc., EMC Corp., Hewlett-Packard Co., Hitachi Data Systems, IBM and NetApp Inc.) today support hardware acceleration of [virtualization I/O](#). This is implemented through application program interfaces (APIs) in the hypervisor, such as [vStorage APIs for Array Integration](#) (VAAI). VAAI offloads some of the "heavy lifting" from the hypervisor by letting the storage array choose the best way to

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

perform key operations, such as sub-LUN locking, bulk copying and zeroing out ranges of data. Most recently, in vSphere 5, VMware added the thin reclaim feature, which lets the hypervisor release deleted storage from thin-provisioned LUNs without directly writing data to deleted blocks.

Offloading storage management tasks to the array provides numerous benefits. First, it reduces the workload on the hypervisor, lessening the CPU load and traffic across the storage network. Second, it lets the storage array optimize and prioritize I/O-intensive operations that may be best achieved internally within the array. As the leading hypervisor vendor, VMware has developed a number of APIs, including vStorage APIs for Data Protection (VADP) and [vStorage APIs for Storage Awareness](#) (VASA). VASA is of increasing importance in the delivery of scalable storage environments, providing configuration information to the hypervisor about storage LUNs, including replication and [performance metrics](#).

Configure for performance

When delivering I/O to virtual environments, performance is everything. Typically, virtual environments create more random workloads, making the work of optimizing I/O workloads much harder for the array. There are techniques that can be employed to ensure performance is delivered optimally, including:

- **Wide striping.** This involves spreading I/O across as many physical disk spindles as possible. Wide striping can be achieved by using large RAID groups (being mindful of rebuild times for disk failures) or by concatenating RAID groups into storage pools. This technique is applicable to both file- and block-based storage platforms.
- **Dynamic tiering.** Like any storage environment, virtual servers will have I/O "hotspots," data that generates a large proportion of the I/O workload. Hotspot areas can be difficult to predict, so platforms that offer dynamic tiering provide an automated way to ensure the hottest data stays on the fastest disk. This technique is particularly useful where virtual machines have been cloned from a single master image.

Use thin provisioning

It's very easy for storage in virtual environments to grow out of control, as virtual machines are relatively easy to create. This is especially true in on-demand environments. Thin provisioning ensures that disk space is consumed only by data that's written to the disk by the host, rather than reserving a fixed image for each VM. The feature can be implemented in the hypervisor and is a common option with most storage platforms.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Use vendor plug-ins

Almost all enterprise and midrange storage platforms offer plug-ins for centralized management tools such as [VMware vCenter](#). This provides a "single-pane-of-glass" view of both virtualization and storage systems, in many cases allowing the storage to be configured directly from the vCenter console. In organizations without dedicated storage teams, this can significantly reduce the work of an IT administrator.

Sampler: Vendor tools for managing storage for VMs

IN ADDITION TO plug-ins for virtualization management tools, there are a number of storage and third-party vendor tools to consider. Some examples include:

- **EMC Corp. Virtual Storage Integrator (VSI).** This VMware vCenter plug-in from EMC provides a rich range of storage information directly within the vCenter console and can integrate with Citrix XenDesktop in virtual desktop infrastructure (VDI) environments.
- **iWave Software LLC Storage Automator.** iWave's Storage Automator allows policy-based deployment of storage for virtual environments, managing the workflow processes of delivering virtual servers in public cloud-type environments.
- **NetApp Inc. SANscreen VM Insight.** NetApp's SANscreen platform provides storage visualization, with integration into VMware VirtualCenter, operating on NetApp or heterogeneous storage configurations.
- **SolarWinds Storage Manager.** SolarWinds has a variety of tools that both visualize and help in the optimization of storage for virtual environments, including virtual machine to storage mapping.

Storage built for virtual servers

A number of startup storage vendors have rolled out hardware and software storage solutions specifically designed for virtual server environments. These include Atlantis Computing Inc., SolidFire Inc., Tintri Inc. and Virsto Software. In essence, these products are architected to address the issues described here, including random I/O challenges.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Automate it

Managing dynamically changing virtual environments to optimize capacity and performance can be a time-consuming process. As virtual environments scale and mature, there's a need to move toward more automation of manual optimization processes. Hypervisor vendors are starting to include capabilities in their products that allow some of these features to be semi-automated, reducing the onus on the administrator to continually tune the storage environment. In vSphere 5, VMware introduced [Storage Distributed Resource Scheduler](#) (SDRS), which provides some degree of automation of storage allocations. SDRS provides automated initial placement of VMDKs, automated migrations of virtual machines to meet capacity and performance goals, as well as affinity rules -- ensuring, for example, that high I/O virtual machines are placed on separate hardware.

The move to more automated storage management will be an absolute requirement as virtual infrastructures scale and become more service-orientated in their delivery. Already, storage vendors are coming to the market with new products that provide provisioning APIs to hook directly into virtual server automation.

Don't forget backup

Backup always seems to be treated as a poor relation in storage management; however, it's of vital importance for delivering high-availability storage environments. In virtual infrastructures, traditional backup solutions aren't the most efficient way to back up and restore data, and other techniques can be used to optimize the backup and restore process.

In [block storage deployments](#), traditional backups use the host itself to back up data. This is because the storage array has no awareness of the format of data on a LUN. The host places the file system onto the LUN, so the backup software relies on the host to provide a stream of files for backup.

On all virtual platforms, a VM is stored as a file or series of files, even when using block-based storage arrays. This makes the backup process easier, as backups can be taken simply by taking a copy of the files that make up the virtual machine.

Some hypervisor vendors, such as VMware, offer APIs that allow third-party software to view changed block data within the virtual machine itself, providing a highly efficient way of backing up only those files that have changed since the last backup was taken. All hypervisor vendors also provide the ability to [snapshot virtual machines](#). Although this results in a "crash-consistent" copy in some instances, with agent software the

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

snapshots can be coordinated by quiescing the host file system to allow consistent snapshots to be taken.

Storage tools will evolve

Storage continues to be a key feature in deploying scalable virtual infrastructures. As these environments scale and mature, storage administrators will need to employ tools and techniques such as automation and visualization software that will allow them to meet the challenges of an ever-integrated IT world.

USING PERFORMANCE METRICS TO ENHANCE STORAGE FOR VIRTUAL ENVIRONMENTS

Performance monitoring is an essential aspect of managing storage for virtual environments. Recently, more storage performance monitoring tools designed specifically for virtual server environments have been hitting the market. These tools can report statistics on a per-VM or per-application basis to help administrators avoid performance problems or easily locate them. The content that follows outlines some of the different performance monitoring approaches for virtual environments.

Storage performance monitoring for virtual servers takes an app approach

Like other areas of storage-area network (SAN) management, virtual server proliferation has made the job tougher for [storage performance monitoring](#). Traditional tools that help you track I/O resources and alert you when performance thresholds are challenged may not be able to distinguish between a [virtual machine \(VM\)](#) and a physical host, leading to inaccurate data.

However, the latest storage performance monitoring products have been adapted to virtual environments, and some are even specifically built for hypervisor technologies.

[Performance monitoring](#) in virtual environments is moving away from siloed tools that individually gauge storage, networking and physical host performance. Instead, VMware Inc. and third-party vendors such as BlueStripe Software Inc. are developing an application-focused approach that aggregates performance metrics and presents an overall picture of system health to virtualization, storage and networking administrators. When issues arise, the tools can drill down for detailed information and find the source of the problem.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Storage performance monitoring tools in virtual environments

In a physical server environment, [storage performance monitoring tools](#) watch the operating system (OS) to measure the server's performance statistics. Typically, there are only two ports for data flowing in and out of the hardware. It's easier to identify the physical links to the storage, including direct LUN connections to applications.

A [virtual environment](#) introduces OS emulators and multiple virtual ports within each physical host, which will render performance monitoring tools built for physical environments unreliable.

"Virtualization has been a wonderful boon to the cost of platforms and the flexibility of where you host applications," said Vic Nyman, BlueStripe's founder and chief operating officer. "But it has also become an increasing challenge in how you assess what a given business application is using in terms of storage, how you measure its availability and how you gauge its performance."

Typical virtual environment storage issues include [storage mapping](#) and faulty configurations. Storage mapping problems arise when there's a storage-related performance issue, but you don't know which physical host the application's storage resides on. Misconfiguration issues arise when administrators make errors provisioning VM storage, or when the assumptions made at the time of the original provisioning no longer apply.

"Things are happening so fast," said Bob Laliberte, a senior analyst at Milford, Mass.-based Enterprise Strategy Group (ESG). "You used to have these heavily siloed domains, and now all of a sudden everything is collapsing."

SCSI reservation troubles also plague data storage administrators. vSphere and Microsoft Cluster Service use SCSI reservations to ensure that VMs writing metadata changes have exclusive access to shared storage LUNs. If multiple hosts concurrently access and update shared resource metadata files, data could easily become corrupted. But if too many reservations are allowed, hosts attempting to access locked storage LUNs will get I/O failures and, after multiple attempts, the operation may ultimately fail.

According to Paul Turner, general manager at NetApp Inc.'s SANscreen business unit, these physical environment tools need to become virtualization-aware. "In a virtual environment, a lot of the existing tools [can] actually work really well," he said. "They just need to become virtualization-aware. They need to know what the VM mapping is through to their storage."

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Virtualization-aware performance monitoring tools

While the need for [virtualization-aware performance monitoring tools](#) is becoming well known, there still aren't that many of them on the market. According to Jeff Boles, senior analyst and director of validation services at Hopkinton, Mass.-based Taneja Group Inc., "performance monitoring in virtual environments is still massively under-served."

A discussion of storage performance monitoring tools in virtual environments begins with those built into [VMware vSphere](#). The vSphere Client monitors storage performance for whole data centers, clusters, physical hosts or individual virtual machines. The Performance tab within the vSphere Client dashboard shows both an Overview and Advanced view. The Overview option shows key statistics, while the Advanced option provides more detailed information. By itself, the vSphere Client has only limited historical statistics for trending and planning purposes.

VMware vCenter Server, formerly called VMware VirtualCenter, provides much more detailed information, as well as better alerting. vCenter AppSpeed, which VMware obtained when it acquired B-hive Networks Inc. in May 2008, tracks transaction performance and measures latency and throughput in the virtual environment as a tab in vCenter. It also can give you application service-level agreement (SLA) status and troubleshoot application performance directly inside vCenter.

Vendor approaches to storage performance monitoring in the virtual world

Several vendors are tackling storage performance monitoring in the virtual world, including BlueStripe Software, NetApp, VMware and Virtual Instruments.

BlueStripe's FactFinder v5: BlueStripe Software has a different take on storage performance monitoring. The company's flagship product, [FactFinder v5](#), was released in March as an application-focused management software package. "We have a whole new approach to managing application systems," BlueStripe's Nyman said. "On the fly, we will automatically discover the applications and the transaction paths. We'll also go through the stack and tell you where that transaction is stuck. It's like a bridge between the customer transaction and the technology and systems that support it."

The BlueStripe agents and passive observer peers into every aspect of an application's performance, including hypervisors, networking and heterogeneous storage systems. "We see the application, its dependency on the storage, and its performance and interaction," Nyman said. But it doesn't necessarily drill down to find the root cause of a problem. "We see when the

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

application has issues with storage,” he explained. “We see when it has [performance bottlenecks](#). We don't necessarily know why it may not be performing well on that storage.”

BlueStripe is not setting out to replace traditional storage performance monitoring tools. Just like VMware vCenter Operations, it aggregates data and presents an overall system-health view. When performance issues arise, administrators can drill down into dependent systems and get the right system manager to look at the performance monitoring and diagnostic tools for more specific analysis.

BlueStripe's FactFinder and vCenter Operations are two products that show the market's increasing interest in overall system monitoring tools, and decreasing focus on separate system monitoring silos. But a proven and robust storage performance monitoring tool like Virtual Instruments' VirtualWisdom is essential in today's increasingly complex virtual environments. Either way, keeping tabs on your storage performance is essential today.

NetApp's Akorri BalancePoint: NetApp acquired Akorri Networks Inc. in February, and moved [Akorri BalancePoint](#) storage performance monitoring technology into the OnCommand management software suite with its existing SANscreen Service Insight storage management solution. SANscreen performs capacity planning and trending in VM environments. It can see all the capacity being used by all the virtual machines, how much each VM is using and what each cluster is using. It can also provide capacity planning reports and analytics.

“The reason we acquired Akorri was because we were seeing quite a significant need for better tools in virtual environments,” NetApp's Turner said. “In particular, what we saw was the need for really good performance modeling tools and performance prediction tools.”

Turner said the Akorri Performance Index watches the CPU and memory headroom and utilizes queuing theory, a mathematical model based on simulated transactions and queues, to determine when a server limit will pass the set threshold. Alerts go into the vCenter console or email.

VMware's vCenter Operations: VMware released its own advanced analytics engine in March called [vCenter Operations](#). VMware's new advanced analytics engine aggregates storage, network, CPU and memory performance data from the vSphere hypervisor and presents it in one view as an overall system health metric.

“We view this as a new approach to both infrastructure and operations management,” said Rob Smoot, VMware's director of product marketing

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

management. "It helps you get an aggregate view of the health, performance and capacity of the environment. Increasingly, we think the three disciplines of performance, [capacity management](#) and [configuration management](#) need to come together."

vCenter Operations uses sophisticated algorithms to determine when the aggregate systems as a whole are behaving abnormally. Then you can drill down into the specifics to see which system -- storage, networking or the physical host -- is causing the bottleneck. Smoot said the goal is to aggregate alerting systems into three cores: workload, capacity and system health.

"What operations teams deal with today are alerts for specific silos and the underlying aspects of the infrastructure," Smoot said. "So they get these alert storms where there's just tons of noise in the environment. What vCenter Operations does is cut through that noise and alert you when there's a building performance problem that spans all of those individual metrics."

Virtual Instruments' VirtualWisdom: Virtual Instruments [VirtualWisdom](#) SAN optimization and troubleshooting software specifically focuses on storage performance monitoring within virtual environments. As the successor to NetWisdom, VirtualWisdom provides monitoring, optimization and troubleshooting for heterogeneous Fibre Channel (FC) storage networks.

Skip Bacon, chief technology officer at Virtual Instruments, said that without a detailed understanding of what's happening under the hood, server virtualization can go bad quickly. "The good news and bad news about server virtualization is that you can spin-up new VMs very, very quickly, you can move VMs around very, very quickly, and in some cases with vMotion automatically," Bacon said. "The bad news is that if you don't have a firm grasp of the underlying storage picture from both a capacity and performance perspective, then all that [dynamic computing] does is get you to the scene of the accident that much more quickly."

The VirtualWisdom platform has several components to deal with the problem. The VirtualWisdom Server runs on a Windows server platform. The ProbeVM software acquires [performance metrics](#) from physical servers and VMs. The ProbeV software gets data from the SAN FC switch fabric, while ProbeFCX sends the base server SCSI device transactions and link metrics. The VirtualWisdom alarms are policy based and can trigger email notification and SNMP traps, as well as execute scripts and vMotion transfers.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Want per-VM reports and statistics? Turn to hypervisor-aware storage

Servers, network devices and storage need to be more hypervisor-aware to keep up with today's expanding virtual infrastructures. Fortunately, some storage vendors have recognized this need and developed virtualization-aware storage that could benefit your infrastructure.

Bare-metal hypervisors virtualize everything on the server, and the underlying guests never know the difference. As these hypervisors have matured, however, they have asked a bit more from the server hardware. For example, the latest version of [VMware vSphere](#) now requires a 64-bit CPU and either an [Intel VT or AMD-V CPU virtualization extension](#). These CPU features allow the software -- the vSphere or Hyper-V hypervisor -- to work with the hardware in ways not previously possible. Now that the CPU cooperates with the hypervisor, storage should follow suit.

How hypervisor and storage cooperate with VAAI and VASA

Today, VMware's [vStorage APIs for Array Integration](#) (VAAI) and Microsoft's [Offloaded Data Transfer](#) allow for some cooperation between the storage and the hypervisor. Many enterprise storage vendors have also adopted VMware's [vStorage APIs for Storage Awareness](#) (VASA).

VAAI speeds common hypervisor storage-related operations, such as virtual machine (VM) cloning or disk space reclamation. When you perform a [snapshot of a VM](#) located on a VAAI-capable array, you'll prevent locking on the entire logical unit number (LUN) unique identifier.

VASA, on the other hand, lets storage report to the hypervisor on the storage capabilities. The hypervisor benefits from these storage reports, but storage could also [benefit from VM reports](#). With that knowledge, the storage could offer the most important VMs the best performance and give the virtualization and storage administrators per-VM statistics.

Let's look at some of the benefits admins can reap from virtualization-aware storage.

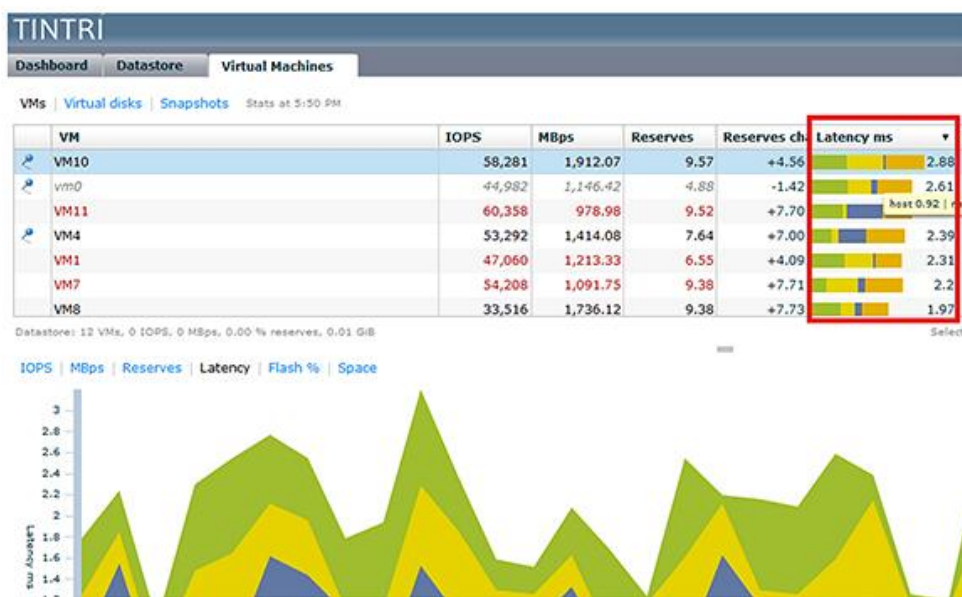
The benefits to hypervisor-aware storage

As an admin, you want to know how many [IOPS](#) each VM generates, which VMs are write-intensive, how many IOPS the whole virtual infrastructure demands and how many VMs are using that storage. Virtualization-aware storage provides you with that information.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Another benefit to virtualization-aware storage is end-to-end [bottleneck identification](#) from the perspective of the storage. Previously, admins blamed the storage for problems they couldn't identify. Now, storage admins and virtualization admins with access to the storage can get more accurate storage statistics. The storage can then combine its knowledge of the virtual infrastructure to quickly solve performance problems and identify bottlenecks. Figure 2 shows latency identification from [Tintri](#) Inc.'s storage management interface.



Finally, virtualization-aware storage goes beyond statistics to actually perform operations on virtual machines, thanks to its connection to vCenter. For example, traditional storage can only take a storage snapshot on an entire LUN. Tintri, on the other hand, can take a snapshot or [clone an individual VM](#) from the storage layer.

Tintri, [Tegile](#) Systems and [Nutanix](#) all provide hypervisor-aware storage, but all hardware vendors must create more virtualization-aware products to keep up with evolving and maturing virtual data centers.

Better server performance monitoring tools needed with virtualization

As workloads grow more complex and become virtualized or are moved to cloud environments, IT professionals can no longer measure infrastructure performance with gut instinct and basic tools. As a result, server performance

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

monitoring tools must evolve to holistically monitor and analyze data center operation.

Data center infrastructure performance monitoring lets IT professionals keep a watchful eye on the performance of servers, as well as other critical infrastructure components. Monitoring tools are increasingly necessary to manage virtualized resources, which can change at a moment's notice. IT professionals proactively shift workloads based on resource use and performance data, and can even anticipate changes in workload demand and plan capacity accordingly.

"Everything we have is virtualized or will be, from our SQL database and related time and billing applications to our Exchange server, which is slated for migration to a virtualized environment soon," said Michael Ferguson, the director of IT for Miami law firm Rennert Vogel Mandler & Rodriguez P.A.

Ferguson's production servers are in a colocated bunker and also in a separate location in Boca Raton, Fla. "My daily concern is keeping consistent bandwidth between Miami, where we have 60 people, and our server in Boca Raton. We are always tweaking and watching -- and fortunately we have high availability in both VMware and storage virtualization in the form of SANsymphony from DataCore," he said.

In that new virtualized world, [monitoring is a challenge](#). "We have spikes and valleys ... in the flow of what is happening at that particular moment," Ferguson said. "For instance, if the real estate department is handling a lot of closings and paralegals are executing a lot of documents within the software, we see spikes. But that is only occasionally." Future performance is also an area of concern. "I have to anticipate things like how [VMware] vCenter running ESX servers is performing. Last week, the reports indicated I needed to increase RAM, so I have ordered more," he said.

Monitoring capabilities built into VMware and SANsymphony are currently his mainstays, because his biggest concern is running short on storage. Ferguson said while his existing management tools are adequate, he would like a single tool that could provide a comprehensive overview of operations and performance.

The market is [saturated with server performance monitoring tools](#), O'Donnell said. The key players are the usual suspects in management: Hewlett-Packard, CA Technologies and BMC Software, and there are even a number of open source server monitoring options, particularly [Nagios](#), which is getting very popular but "suffers from complexity," said Glenn O'Donnell, Forrester analyst.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Monitoring on critical production systems should be pervasive, according to Jonah Kowall, Gartner research director for IT operations management, in the report "How to Augment APM with Conventional Monitoring." Beyond server performance monitoring, tools should focus on [application performance management](#) (APM) and monitoring, since users ultimately rely on application performance.

A combination of server, network, database and storage monitoring, along with more in-depth app performance monitoring, can deliver a complete view of the application and its performance, Kowall said.

Analytics in the infrastructure monitoring toolset

Advanced algorithms can "help analyze the mess on the back end, including understanding behaviors and why things are happening," added Glenn O'Donnell, Forrester analyst. This is how [application performance monitoring and other kinds](#) of monitoring are coming together, he said. "APM tells you that you have an issue and, in theory, the analytics will tell you why and will help you drill down and map the problem."

The key is encouraging IT professionals to acquire more sophisticated analytics capability as they can get it. O'Donnell said there are a lot of useful management features coming from big companies, such as IBM with its [Tivoli management software](#), but most of the innovation is still coming from little "peanut companies." Over the next few years, the further development of analytics will make monitoring much more exciting, he said.

Ferguson finds that kind of broad capability -- to supplement what's built into his SANsymphony storage manager -- very desirable. "The second thing on my wish list after 'Don't bust the budget' is a dashboard app that would reach into the infrastructure and provide at-a-glance reports I can drill down into," he said.

If a utility included application performance, that would be a plus, Ferguson said. The billing, accounting and financial management software application at his firm "gets used constantly, but not with a big hammer." Application performance monitoring would be more welcome for the firm's most demanding applications -- built on SQL Server and Linux-based MySQL -- especially during end-of-month closing activities.

"If you have a good dashboard, things go well. The instruments on your panel, if you trust them, can get you a long way," he said.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

VMware storage performance analysis

To run either a physical server environment or virtual server environment efficiently, [storage performance analysis](#) and storage performance monitoring tools are essential to IT administrators. They're used to track I/O resources, alert administrators of [boot storms](#) and potential performance thresholds, and possibly improve [storage capacity](#) and utilization.

In this podcast interview, VMware expert Mike Laverick explores [VMware storage performance](#) analysis. Find out about the pros and cons of implementing storage performance analysis in a VMware environment, whether storage performance monitoring tools for a physical server can be used in a virtual server environment, and what storage performance analysis tools are offered by VMware and third-party vendors.

Read the transcript or listen to the podcast on VMware storage performance analysis. Click on the corresponding podcast player to listen to Laverick's answers to each question.

SearchVirtualStorage.com: [How is VMware storage performance analysis handled?](#)

Laverick: Well, the first thing to remember is that ... the I/O is being created ... inside the virtual machine. Therefore ... the analysis and the data collection [have] to start ... with the ESX host—the hypervisor, if you like—because that's where the I/O is being generated. But it doesn't stay on the ESX host. It then reports its data to Virtual Center, assuming that the customer has that. And Virtual Center then starts to manipulate that data in SQL over a period of time. So it will take the data that is recorded almost on a second-by-second basis but roll it up into five minutes or an hour, daily, weekly, monthly, even yearly [view]. So you've got a view of your performance data over a longer period of time.

The important thing is if you want that information, [the] SQL agent, on the SQL server, which is the back-end database to Virtual Center, must be running; otherwise, it can't run the job that would then roll up the statistics. And generally, the information that you get is very detailed. The only thing I would say is that historically there's always been a kind of lack of data around [NFS](#) and the performance that we get from NFS data stores. And to a degree that persists, even to this day, even though NFS has become increasingly popular.

- **Contents**

- Managing storage for virtual environments: Best practices
- Top tools for virtual server storage management
- Using performance metrics to enhance storage for virtual environments
- Hypervisor-specific data storage management

SearchVirtualStorage.com: [What are the advantages of implementing storage performance analysis in a VMware environment?](#)

Laverick: I think you've got to remember that the resource that most people run out of first is memory. But I don't think we can underestimate the impact of [storage in a virtualized environment](#). [That's] especially [true] if you've got very read- or write-intensive virtual machines like SQL or Exchange, or if you experience random I/O issues or even peak activity ... or if you work in the retail sector, spikes in demand caused by Black Friday, Thanksgiving, Christmas and that kind of thing. On a more specific level, if you're doing any VDI work, you can get boot storms once the virtual desktops power up. And also if you're running virus analysis software inside a VM, that can be the cause of this activity.

The bottom line is if you're not monitoring, you don't know what resource, where, is actually creating a problem. But I think just generally [there's an advantage in] validating your best practices and making sure those are being followed, because in my experience most problems are caused by people not following the best practices.

SearchVirtualStorage.com: [Can storage performance monitoring tools for a physical server be used in a virtual server environment?](#)

Laverick: I would never really use tools inside a virtual machine if I was monitoring CPU or memory. But there is more temptation, I think, to actually look at disk activity in a virtual machine using the guest operating system tools. After all, a write to a disk is a write to a disk; whatever message it gets to the physical storage, you can't avoid the physical world. But what makes it different is [that] the VM [just] sees a SCSI adapter; it has no idea if you've got Fibre Channel HBAs or whether you've got 10 [Gigabit Ethernet] storage. It just sees a virtual SCSI adapter. And any read/write [activity] that takes place inside the virtual machine gets translated [down] through some sort of storage driver or controller on the host.

Now if that's NFS or iSCSI, it's going to be ... an Ethernet driver. Obviously if it's Fibre Channel, it'll be a driver for Emulex or QLogic. So although most in-guest tools will give you reliable information, you have to treat them with great caution, because the big thing is that one virtual machine is totally unaware of all the other virtual machines that are on the same physical host. And all those virtual machines together are sharing the same connection to the storage back end, whether that's Ethernet or whether that's Fibre Channel. So if you really want to know what's going on it has to be ... done outside the scope of the virtual machine, outside of the guest operating system, using either tools from the hypervisor vendor, whether that's VMware or Microsoft. And really what you're looking for when you're doing that kind of

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

analysis outside of the guest is you're looking for contention. You're looking for volumes or LUNs that are oversubscribed, or you have too many VMs on a single LUN or a volume and they're just generating far too much I/O. Or [in] most cases, you're looking at the array; has the LUN that's being presented to the host got insufficient spindles backing it to deal with the I/O, or has the wrong RAID level been used for that particular profile of storage?

I think the key thing to get across from doing any analysis of storage is that you're going to probably need a bit of information from both environments—from your virtualization environment and probably from your array. And, in most cases people find that the problem isn't with the virtual machine but with the storage array. It's a classic example of mistaking the symptom—poor performance in the virtual machine—for the cause—which is often something external, something to do with the way the array is being configured. Perhaps the most common mistake that people make is when they're going to create a new virtual machine they sort the data stores that they've got available by free space, and they just plunk their VMs on the data stores that have the most free space. What I always suggest to people is sort your data stores by their I/O capabilities—fast or slow storage, is it SATA or is it SAS or even [SSD] storage—and use that as your guide for where you put your virtual machines, not just which ones have got the most free space on.

SearchVirtualStorage.com: [What tools does VMware offer to properly implement storage performance analysis?](#)

Laverick: There [are] lots of different tools. At the physical ESX host there's something called ESXTop, which if you're from a Linux background, it's kind of VMware's version of Top. And it can show not just CPU activity, but also memory, disk and network activity. So from the disk view of ESXTop, you can see read/writes per second, and any latency you might have, that's significant for networking—Ethernet-based storage—and any excessive queue depth you've got. One of the things I've often found with ESXTop is that it does generate an awful lot of data very rapidly. So I've often used a free tool from VMware ... called ESXplot, which can take a .csv file and then import that and give you a better way of monitoring and viewing the data. If you're looking for tools to control the storage, since, I think, vSphere 4.1, [Storage I/O Control](#) has been possible. So it's possible to kind of control on a virtual machine basis, VM by VM, which ones have what access to the storage layer. And as I said before, there's Virtual Center . If you're looking to spend more money on top of what you might already have, VMware [has] the [vCenter Operations Manager](#), and that will produce some very interesting statistics and give a thumbnail sketch of what your environment's currently like, and it learns over a period of time what your expected behavior [is]. And you can have [a dashboard], so if something unusual happens, it's [flagged] straightaway.

- **Contents**

- Managing storage for virtual environments: Best practices
- Top tools for virtual server storage management
- Using performance metrics to enhance storage for virtual environments
- Hypervisor-specific data storage management

SearchVirtualStorage.com: [What third-party vendors are tackling VMware storage performance analysis in a VMware environment, and what tools do they offer?](#)

Laverick: That's a very good question. I remember when I was at [a VMUG meeting] in San Diego a couple of months ago. Somebody was explaining that they had a storage problem, and another guy around the table said immediately, "You need to look at these products." And I took a step back from that and said, "Before you start looking at third-party vendors and start parting with your hard-earned cash, use what tools you have available," because a lot of the time I find customers are a little bit too willing to go out and spend money on a third-party solution when they haven't really exhausted all of the "free" tools they have already. But putting that caveat to one side, when I was in Boston recently at the [Tech Field Day](#) organized by Stephen Foskett, I was very impressed by Akorri. [Akorri's tools] can actually track the full path of the virtual machines through the connections into the array, even into different controllers and to the volumes. You see a full window of where the IOPS are being generated, which I hadn't seen before. I think [vKernel](#) and [vMonitor](#) have quite good tools. And I guess there's always the possibility of taking the spindle out of the equation. In the last 18 months to two years, we've seen increasing use of cache or SSD at the front of an array to absorb the IOPS requests of things that are very disk-intensive. And if we can take the spindle out of the equation, which is the most constraining physical resource, there's no reason why a virtual machine would be affected by storage performance problems, but I guess that's very much dependent on what generation of array you're on and whether you're prepared to spend that money to get those kinds of performance levels.

Virtualization storage performance monitoring: Metrics and tools

[Virtualization storage performance](#) has a direct effect on application performance and user experiences, so admins need to implement comprehensive monitoring to track storage performance.

Virtualization storage performance monitoring provides three distinct advantages:

1. [Storage performance monitoring](#) provides real-time warnings of serious problems -- usually summoning an administrator or technician when thresholds are exceeded.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

2. It yields important historical data about the virtualization storage that could reveal usage trends and help IT departments predict future resource needs before [storage performance](#) is affected.

3. Storage performance monitoring helps reveal problems. Virtualization obfuscates the relationship between storage and applications, and that's fine until trouble occurs. Abstraction can make it difficult (or even impossible) to locate the source of storage performance problems, so storage performance monitoring helps cut through that fog. Proper monitoring can help identify the real location of deteriorating virtualization storage performance.

Performance monitoring is key to capacity planning

Storage performance monitoring can highlight problem areas long before storage is exhausted or performance is compromised. It can lay the groundwork for investments in additional capacity and new technology. A LUN with only 50 MB of space remaining, for example, may require additional disk capacity or migration to another platform.

"Without monitoring ... you wake up one morning to find that your file server is out of disk space, and you're trying to figure out how that happened because six months ago, you tripled the size of that volume," said Dwayne Sye, chief information officer at Cvent Inc., an event management software company in McLean, Va. "With a little more analysis, you can perhaps demonstrate that an increase in disk usage was the result of a hiring spree three months ago," he said.

In some cases, the focus is on applications and their performance -- not storage. For example, monitoring may reveal a decline in the application's performance, prompting a closer examination of the underlying storage performance.

In other cases, IT managers use monitoring as a last-minute sanity check following an incident or technology initiative. Experts say that's a mistake because either event would disrupt normal daily operations.

Experts point out that the actual baseline for virtualization storage performance monitoring points are almost identical to those used for traditional (nonvirtualized) storage. These include metrics that should already be familiar to many storage professionals, including LAN and/or SAN bandwidth, I/O operations per second (IOPS), storage queue length, throughput and response time.

Don't give all storage performance metrics the same weight. Palmer says that the importance of each metric in storage performance monitoring will

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

vary by application and server. For example, response time may be most important for a front-end, customer-facing application, while IOPS may be critical for small-block, batch-processing applications.

Storage performance monitoring tools

Virtualized environments are assisted by a new generation of storage management tools and application programming interfaces that are more "virtual machine-aware" than they were in the past. Many storage vendors are now capable of displaying the properties of virtual machines (VMs) from within the storage management system. Previously, storage performance monitoring tools displayed only logical unit numbers (LUNs) and volumes, and there was little or no visibility into the contents of those storage units and into the VMs within them. This often led to difficulty in calculating the true amount of disk space consumed by a VM and identifying which VM was the source of excessive I/O demands to the storage layer.

In addition, great strides have been made in optimizing the demands made by the hypervisor to the virtualization storage layer, and it's now possible to offload a great deal of the IOPS that were generated by common VM tasks such as disk-to-disk backup, taking a snapshot or cloning a VM to the array. The new generation of arrays whose firmware is compatible with VMware's [vStorage APIs for Array Integration](#) (VAAI) initiative is a good example.

Even though service levels will vary between different storage tiers or pools, storage performance monitoring should ideally be applied uniformly. This gives a more comparative view of virtualization storage resources.

But experts are quick to note that there are exceptions to this rule about virtualization storage performance.

"There's always a cost consideration," said Dwayne Sye, chief information officer at Cvent Inc., an event management software company in McLean, Va. "If it's cost-prohibitive to monitor a SATA array at the same level as a Fibre Channel array, economics may play a factor. But certainly, if it's possible ... I think that would be the optimal situation."

Tony Palmer, senior engineer and analyst at Enterprise Strategy Group, takes a slightly different approach, underscoring the role of applications rather than virtualization storage. For example, monitoring the storage used for Exchange Server may take a higher priority than monitoring the storage used for backups, resulting in different storage performance monitoring procedures. "It should all be monitored using the same tools, but you're going to change your tactics appropriate to the importance of the application rather than the importance of the storage," he said.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

The future of storage performance monitoring

Virtualization admins can look forward to a time when virtualization storage systems support something known as "auto-tiering." This is the ability for LUNs to be "migrated" from one storage array to another or from one class of storage, say SATA, to another, such as solid-state drive. This "Storage vMotion," for want of a better phrase, is already present in some newer storage arrays, and it can eliminate "hotspots" where IOPS are oversaturating a particular array or subsystem.

Solid-state storage is increasingly used tactically to remove the disk spindle from the performance equation. This has opened the door to applications that have historically been regarded as too dependent on disk I/O to be easily virtualized.

As with [thin provisioning](#), customers should expect this to happen seamlessly at the storage layer but controlled and governed at the virtualization layer by a policy-based system. The requirement for administrative approval should stop VMs from being moved unexpectedly. Sadly, this kind of storage vendor-optimized volume or LUN migration will likely be available only in the top tier of enterprise-class storage.

Someday, dynamic resource management systems will include not just optimizations for memory and CPU-bound workloads, but also the capacity to optimize the network and storage I/O paths.

HYPERVISOR-SPECIFIC DATA STORAGE MANAGEMENT

More management tools have added support for Microsoft Hyper-V, and the latest version of VMware vSphere made improvements to existing features. From the content below, learn what management features the various hypervisors provide and how to best work with them.

Hyper-V storage management options expanding

VMware is the most popular choice for virtualizing servers by a wide margin, but it isn't the only option. Microsoft's [Hyper-V hypervisor](#) is gaining users because it costs less and requires little extra training for Windows users. As a result, more [Hyper-V storage](#) management tools are becoming available.

In addition to Microsoft's own [System Center Virtual Machine Manager \(SCVMM\)](#), vendors have released third-party products to manage performance issues in Hyper-V virtual environments. Here's a rundown of some of Microsoft's inherent Hyper-V storage management resources as well as some third-party tools.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

System Center Virtual Machine Manager

SCVMM is Microsoft's virtual machine (VM) support center. Microsoft added Hyper-V support in SCVMM 2008, and subsequent releases included enhanced support for Hyper-V and VMware Inc.'s hypervisor.

Microsoft's SCVMM 2008 includes typical virtual environment management functions, such as VM provisioning and monitoring, physical-to-virtual conversion (P2V), and host resource management. It includes storage-related capabilities, such as the ability to dynamically add and remove storage resources. Its Live Storage Migration allows migration from one LUN to another with minimal downtime.

Other SCVMM 2008 storage features include [Data Protection Manager \(DPM\)](#) combined with [Volume Shadow Copy Service \(VSS\)](#) for VM snapshots, and [Clustered Shared Volume \(CSV\)](#) support for storing multiple VMs on a single LUN to allow Windows failover clusters. All of these features are accessible from the System Center GUI.

[SCVMM 2012](#), scheduled for general release before the end of the year, will add more application-driven monitoring and resource orchestration, as well as virtual VM network and storage hardware provisioning through automated wizards. To compete with VMware's [vCloud Director](#), SCVMM has self-service portals for [Infrastructure as a Service](#) (IaaS) support.

Microsoft also increased its support of the Storage Networking Industry Association (SNIA) [Storage Management Initiative-Specification \(SMI-S\)](#) standard in SCVMM 2012 for [private cloud storage](#) integration with storage arrays. SMI-S specifies how clients can communicate with storage arrays through SMI-S Provider modules.

SCVMM 2012's new feature set prompted Citrix to discontinue its [Citrix Essentials for Hyper-V](#) storage management software in May 2011 due to overlap and SCVMM's new support for XenServer VMs.

Storage system vendors are also integrating their storage management tools with SCVMM and extending its native data protection capabilities. Enterprise Strategy Group (ESG) analyst Mark Bowker said that may be the best starting point for Hyper-V [storage management](#) tools. "If I'm sitting in IT's shoes, I'm actually starting with the vendor who I'm dealing with now," Bowker explained.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

For example, you can manage NetApp Inc. storage and data protection tools through the SCVMM console. NetApp will automate Hyper-V snapshots and recovery, as well as set policies for automated VM-level backups. SnapManager for Hyper-V will also automatically detect new VMs and identify those without backup snapshot policies. You can even use custom scripts to automate tape offload before or after a backup.

Dell's OpenManage Integration Suite for Microsoft System Center and Advanced Infrastructure Manager (AIM) offer storage and virtualization administrators unified tools for monitoring, deploying and configuring a suite of integration tools for Hyper-V environments using Dell hardware.

Performance management

Bowker said performance is the No. 1 storage issue among all hypervisors. While storage system vendor integration with System Center can manage a number of necessary tasks, storage system performance management might be better done with third-party tools specifically designed for Hyper-V environments.

Virsto Software Corp. chose to support only Hyper-V when it launched its initial [Virsto One](#) hypervisor storage performance management software in February 2010. Gregg Holzrichter, Virsto's vice president of marketing, said Virsto wanted to be one of the first providers for Hyper-V environments instead of just one of dozens of competitors supporting VMware's vSphere.

"I think that was a smart move on [Virsto's] part because its technology can transfer to the other hypervisors when it needs to do it anyway," said Henry Baltazar, a senior analyst for storage and systems with The 451 Group.

Virsto One virtualizes underlying heterogeneous storage and manages the virtual environment's random and intensive writes, which Holzrichter calls "the biggest challenge of server virtualization writing over to storage."

Virsto One intercepts writes from the VMs to the storage system and sequentializes and optimizes them to improve performance. Virsto also supports storage tiering and [thin provisioning](#) for the storage allocated to the VMs.

Virsto does not support automated, policy-based or dynamic storage tiering now, but Holzrichter said those features are on the company's roadmap.

The Virsto software installs on each physical server, and the management tool is a snap-in module to the Microsoft Management Console (MMC). Administrators use the Virsto management console for initial storage

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

configuration and to add capacity. Otherwise, the management is done through the System Center console. When provisioning a new VM, an administrator selects a Virsto virtual disk instead of a fixed [virtual hard disk \(VHD\)](#).

Sanbolic Inc.'s Melio product line focuses on Hyper-V's limitations with concurrent file-system access. Melio FS is Sanbolic's symmetrical cluster file system, which enables shared, concurrent read and write access to iSCSI or Fibre Channel heterogeneous SAN volumes and VHDs. Melio allows the addition of new physical hosts and storage resources for scale-out infrastructures.

Melio FS also has a VSS-based distributed snapshot capability for data protection with third-party backup products, quality-of-service (QoS) management and I/O transactional performance reporting.

Windows familiarity sways L.A. medical group

Because Hyper-V runs on Microsoft's Windows Server 2008 R2 and is fully integrated into SCVMM, provisioning and handling a Hyper-V virtual environment may be easier than having to implement an entirely new platform such as with VMware vSphere and Citrix Systems' XenServer.

Femi Adegoke, the IT director for the Los Angeles-based West Gastroenterology Medical Group (WGI), cited price and his familiarity with the Windows-based Hyper-V as reasons for switching to it from the old Virtual Iron (now part of Oracle VM) server virtualization software in May 2009. He said adopting Hyper-V meant he didn't have to learn a new operating system.

He currently has Hyper-V running his PBX phone, security card access and payroll systems. Now that vendors are starting to develop products for Hyper-V environments, Adegoke plans to put his entire IT business infrastructure in a virtual environment.

"The ecosystem suddenly now seems to have a lot of activity, including backup providers and monitoring software," he said.

Adegoke manages his Hyper-V environment with SCVMM 2008 and uses Windows PowerShell custom scripts to automatically provision VMs with Virsto One virtual storage disks and desktop image clones. Adegoke said he was one of Virsto's first five customers and began using it in early 2010.

"For me, the ability to thin provision and still get the high performance is huge," Adegoke said about using Virsto within his Hyper-V environment.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

“Otherwise, I’d be out of storage space looking to go spend some more money to buy another shelf of spindles.”

More Hyper-V support coming

Because it is the market leader, VMware’s vSphere garners the most attention from vendors looking to support virtual environments. But as those supporting applications and devices mature, their technologies will almost certainly make their way into the Hyper-V and XenServer marketplaces. That will make more storage management tools available to expand the SCVMM tool set and bring more advanced management and data protection technologies to Hyper-V storage environments.

“They are going to trickle their technologies down to the hypervisors,” Baltazar said of third-party software vendors. “In some cases I don’t think it’s a major development effort to [switch] over to the other hypervisors. I think it’s more a matter of there’s not enough value in doing it now because they don’t have enough customers beating down the door asking for them.”

Bowker said all hypervisors have many of the same storage issues. “It really comes down to configuring it and setting it up specific to what application you may be running,” he said, adding that the differences between the hypervisors come down to familiarity, features and function.

LUN storage management for vSphere and Hyper-V

The LUN has long been a bedrock of storage configuration for physical servers with [LUN storage](#) partitions carved out from [RAID](#) groups to provide logical chunks of capacity for applications.

Now, virtual server environments can abstract the physical characteristics of a server into software and so provide increased scale and utilisation of hardware resources.

But, storage must still be provided to virtual server and virtual desktop machines, with the [hypervisor](#) taking on an important role as the virtualisation layer, abstracting physical storage resources to virtual devices.

So, [what has happened to the LUN?](#) That depends on the virtualisation environment you’re using.

Physical and virtual drives and LUNs

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Regardless of hypervisor type, the persistent retention of data needs some form of storage device, either a traditional hard drive or a solid-state disk (SSD). For block storage, VMware's vSphere suite, including ESXi and Microsoft's Hyper-V use fundamentally different approaches to presenting physical storage.

vSphere systems take [LUNs configured](#) on a storage array and format them with the [VMware File System \(VMFS\)](#). This is a proprietary file format used for storing virtual machine files that takes advantage of on-disk structures to support highly granular levels of object and block locking.

The reason this is necessary is that most vSphere deployments use a small number of very large LUNs, with each [LUN](#) holding many virtual machines. An efficient locking method is needed to ensure performance doesn't suffer as virtual environments scale up.

A single virtual machine is comprised of many separate files, including the [VMDK or Virtual machine Disk](#). A VMDK is analogous to a physical server hard drive, with a virtual guest on vSphere potentially having many VMDK files, depending on the number of logical drives supported, the number of snapshots in use and the type of VMDK.

For example, for thin provisioned VMDKs, where storage is allocated on demand, a guest hard drive will consist of a master VMDK file and many VMDK data files, representing the allocation units of each increment of space as the virtual machine writes more data to disk.

By contrast, Microsoft has chosen to incorporate all components of the virtual machine disk into a single file known as a [VHD \(virtual hard disk\)](#). VHD files are deployed onto existing Microsoft formatted file systems, either using [NTFS](#) or [CIFS/SMB](#).

There is no separate LUN format for Hyper-V. VHD files allocated as thin volumes (known as dynamic hard disks) expand by increasing the size of the file and consuming more space on disk. Inside the VHD, Microsoft stores metadata information in the footer of fixed size VHDs and in the header and footer of dynamic VHDs.

VHDs have advantages over VMDKs and VMFS in block-based environments in that the underlying storage of the data is NTFS, Microsoft's standard file system for storage on Windows servers. This means VHD files can easily be copied between volumes or systems by the administrator without any special tools (assuming the virtual machine isn't running of course).

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

It also makes it easy to clone a virtual machine, simply by taking a copy of the VHD and using it as the source of a new virtual machine. This is particularly beneficial using the new [deduplication](#) features of [Windows 2012](#) which can significantly reduce the amount of space consumed by virtual machines that have been cloned from a master VHD.

Designing for performance and capacity

The aggregation of servers and desktops into virtual environments means that the [I/O](#) profile of data is very different to that of a traditional physical server. I/O workload becomes unpredictable as the individual I/O demands from virtualised servers can appear in any order and so are effectively random in nature.

This is referred to as the [“I/O blender”](#) effect and the result is that storage provisioned for virtual environments must be capable of handling large volumes of I/O and for virtual desktops, to cope with [“boot storms”](#), which result from high I/O demand as users start their virtual PCs in the morning and close down at the end of the working day.

To guarantee performance, typical storage deployments will use a number of options:

- [All-flash arrays](#) are becoming increasingly popular for virtual environments. For virtual servers, they provide consistent, predictable performance; for virtual desktops they handle the boot-storm issue with lots of I/O bandwidth.
- [Hybrid flash arrays](#) use a mix of traditional spinning media and solid state, targeting active I/O at the solid-state storage using dynamic tiering technology. This provides a more attractive price point than all-flash arrays, as many deployments have large amounts of inactive VM data.
- **Advanced features** – For vSphere these include [VAAI \(vStorage APIs for Array Integration\)](#) and for [Hyper-V ODX \(Offloaded Data Transfer\)](#). Both of these features offload repetitive tasks from the hypervisor and reduce the amount of data transferred over the storage network, when performing common tasks such as replicating virtual machines or initialising file systems.

Ultimately, provisioning storage for virtual environments is all about getting the right [IOPS](#) density for the capacity of storage being deployed. This may seem difficult to estimate but can be taken from existing physical servers as part of a migration programme, or by pre-building some virtual servers and measuring IOPS demand. For virtual desktops, a good estimate is around 5-10 IOPS per desktop, scaled up across the whole VDI farm. This would require additional capacity to be built in for boot storm events.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

LUN performance and presentation

For block devices, LUNs can be presented using [Fibre Channel](#), [Fibre Channel over Ethernet \(FCoE\)](#) or [iSCSI](#). Fibre Channel and FCoE have the benefit of using dedicated [host bus adapters \(HBAs\)](#) or CNAs (Converged Network Adaptors) that make it easier to separate host IP traffic from storage network traffic. However, there are still some important design considerations even where a dedicated storage network is available.

Firstly, there's the option to present LUNs across multiple Fibre Channel interfaces for both resiliency and performance. We'll take resiliency as a given, as that would be standard practice for storage administrators, but for performance, multiple HBAs (or dual-port HBAs) allow physical segmentation of vSphere and Hyper-V LUNs by tier for performance purposes.

This may not seem like the most logical approach, but bear in mind LUNs presented to vSphere and Hyper-V are typically large, and so queue depth to individual LUNs can become an issue, especially with workloads of different priorities. This can be especially important where high performance all-flash devices have been deployed. For iSCSI connections, dedicated NICs should be used and multipathed for redundancy. Both Microsoft and VMware have deployment guides to show how to enable iSCSI multipathing.

LUN sizing

While on the subject, it's worth discussing LUN sizes. vSphere (and less so Hyper-V) are limited in the number of LUNs that can be presented to a single hypervisor. Typically, storage for these environments is presented using large LUNs (up to 2TB) to maximise the presentable capacity. As a result, the users of that LUN, which could represent many hosts, all receive the same level of performance.

Creating many LUNs of 2TB in size is quite expensive in storage terms. So, [thin provisioning](#) on the storage array presents a useful way to enable LUNs to potentially expand to their full 2TB capacity, while enabling multiple LUNs to be presented to a host to ensure I/O is distributed across as many LUNs as possible.

Limits of the LUN and the future

The grouping of storage for hypervisor guests at the LUN level represents a physical restriction on delivering quality of service to an individual virtual machine; all guests on a LUN receive the same level of performance.

- **Contents**
- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Microsoft recommends using a single LUN per VM, which may be restrictive in larger systems (and certainly represents a significant management overhead), but is still possible to achieve.

VMware has stated its intention is to implement [vVOLS – virtual volumes](#) – to abstract the physical characteristics of the virtual machine storage from the storage array to the hypervisor. This would enable better granularity in terms of prioritisation of virtual machines and their I/O workload, even when they exist on the same physical array.

But while some companies focus on removing the storage array completely, it's clear there are benefits in retaining an intelligent storage array, one that understands and can communicate with the hypervisor.

Managing storage in a Citrix environment

There are many similarities with managing storage in a [Citrix environment](#) compared with a Hyper-V or VMware environment. For instance, all three environments have similar functionality, such as [raw device mapping](#) (RDM) and the ability to write to a file on a dedicated file system. With the Citrix XenServer environment's management platform, a storage administrator can access any given storage device, set up new [NFS](#)-based storage through a wizard-like interface and more.

In this podcast interview, Sander van Vugt, Linux high availability and virtualization expert, discusses best practices for storage in a Citrix environment. Discover the differences between the Citrix XenServer, [vSphere](#) and Hyper-V platforms; what tools Citrix and third-party vendors provide to ease [XenServer storage](#) management; and how to enhance storage performance in a Citrix virtual server environment. Listen to the podcast or read his answers below.

SearchVirtualStorage.com: How is managing storage for a Citrix XenServer virtual server environment different from managing storage for a VMware or [Hyper-V environment](#)?

Van Vugt: Basically, it's not that different. Storage is based on the same principles on all three environments. So in VMware you have raw device mapping, and in Citrix and Hyper-V, there are similar solutions. Likewise, they can all write the storage to a file on the dedicated file system if you don't have to write directly to the storage. The main difference between ... them is in the management interface that XenCenter provides for storage management, so the interface gives easy access to the storage device of choice, and that seems to be different in all the solutions.

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

SearchVirtualStorage.com: Tell us about that XenCenter interface. How does it work?

Van Vugt: [With that management interface], different storage options are set up through an easy wizard-like interface and allow customers to set up NFS-based storage or connect to a storage repository on software, [iSCSI](#) or a hardware HBA.

It's also easy to connect to ISO libraries using NFS or CIFS protocols. And apart from that, Citrix offers advanced [StorageLink](#) technology as an add-on. So this solution is developed to integrate with storage that is provided by third-party vendors. One benefit of the solution is that it provides a generic interface that allows users to manage different SAN fabrics and arrays. So that means that you can manage interfaces from one Citrix environment but not necessarily from the SAN fabric environment.

Also, because it integrates so well with these storage solutions, storage [managers] can take advantage of features by device, such as [thin provisioning](#), fast clones, fast snapshots and all those kinds of pieces that make storage really nice to work with and which aren't necessarily known to the native Citrix environments.

A third benefit is that StorageLink allows administrators to set apart their storage profiles for new virtual machines (VMs). These storage profiles can then be linked to storage volumes in specific arrays.

SearchVirtualStorage.com: What third-party tools are available to help IT shops manage storage to support their Citrix environments?

Van Vugt: What you're looking at is mainly hardware-specific versions of XenServer. For instance, XenServer Dell Edition is optimized to run from flash storage in Dell servers, and in such environments the [XenServer EqualLogic](#) storage data can be used to allow thin volumes to be created from a XenCenter interface. So we're not really looking at generic tools, but more at specific tools that allow you to integrate specific kinds of storage devices in virtualization environments.

SearchVirtualStorage.com: Finally, what techniques can IT shops use to fine-tune performance of storage in a Citrix environment?

Van Vugt: First, they need to choose the right kind of storage for the solution they want to implement. If, for example, ultimate speed is needed, they are better off connecting directly to a LUN instead of creating a file on a file system. The next step, after you have chosen the right tool, is to choose the right hardware, because there are so many differences in hardware. Specialized solutions for storage optimization exist on the market, like

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Atlantis for Citrix XenServer. This is a software appliance that runs on top of the XenServer hypervisor to optimize access for mainly the XenDesktop Virtual Desktop. But these kinds of storage optimization solutions are used to optimize the storage process for a [VDI environment](#) and not necessarily for a virtualized server environment. So for server environments, it's a different story. An optimization takes place on the SAN fabric itself, mostly.

VMware storage management in vSphere 5: VSA, VMFS and VASA

VMware vSphere 5 takes storage management to a new level. Not only does the new version include [Storage DRS and storage intelligence in VMware HA](#), but it also improves the file system and storage APIs, and added a new storage appliance.

VSphere Storage Appliance

VSphere 5 ushered in a new [storage appliance](#), which is available as a separate purchase. Virtual storage appliances are virtual machines (VMs) that run on your hosts and add capabilities to your existing storage. For instance, it's possible to populate a physical server with cheap direct-attached storage and then have the new vSphere Storage Appliance (VSA) present this storage to the outside world.

You might think performance would be subpar compared to other methods of presenting storage on the network, but the surprising thing is how well these appliances work. Storage virtualization will be commonplace within a few years, with many of the storage vendors creating a layer of virtualization that separates the inbound connection to the storage from the disks, which will allow customers to move these virtual instances from one array to another.

In vSphere 5, multiple VSAs can be clustered natively together, offering up a single NFS export point and eliminating the VSA as a single point of failure. VSA requires ESX 5. Before it can be enabled on a cluster, the installer confirms the hosts that are valid for its use. The installer then aggregates the local storage of the host the appliance is running on, and the VSA Manager and the VSA Cluster Service configure and monitor the appliance's availability and the NFS export.

Like with [vSphere Replication](#), VSA includes a wizard that imports the appliance and a plug-in to manage the service directly from vCenter. The manager deploys the appliance and automates the mounting of the NFS export to the host – saving the administrator the hassle of using the Add Storage wizard to mount it to each host. Using the VSA alongside vSphere Replication with Site Recovery Manager, you could potentially store your

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

VMs while also offering a DR service -- all without the expense of an enterprise storage array.

Improvements to VMFS

Not content with creating a new storage and replication appliance, VMware has also been busy upgrading existing technologies. VSphere 5 introduces a new version of the Virtual Machine File System (VMFS) that does not require you to power down VMs or use Storage vMotion to upgrade from the older VMFS 3.

VMFS 5 allows for much larger partitions than its predecessor. VMFS 5 abandons the MBR or cylinders-heads-sectors method of returning the size of a VMFS volume. Previously, this limited VMFS to a maximum single size of 2 TB. The only way to create a file system beyond that limit was by sticking them together using a feature called VMFS Extents, prompting many customers to use NFS instead. VMFS 5 marks a shift to using GUID Partition Tables (GPT) within the file system, which allows for the maximum single VMFS partition to be a maximum of 64 TB.

That being said, VMFS 5 doesn't yet alter the maximum size of a virtual disk, though it does allow for pass-through Raw Device Mappings (RDM) to be 60 TB. RDMs allow a VM direct access to the logical unit number on a storage array.

VMware-Aware Storage APIs

[VMware-Aware Storage APIs](#) (VASA) in vSphere 5 supersede the previous integration introduced in vSphere 4.1.

This VMware storage management feature is implemented differently than the [vStorage APIs for Array Integration](#) (VAAI). VAAI uses T10 standard "primitives" to drive storage improvements to vSphere, while VASA requires a storage vendor plug-in. This model allows VMware and the storage vendors to cooperate. It also means the storage vendors can present features that are unique to their arrays to the ESX host. In vSphere 5, VASA finally solves an issue from last year, where some storage vendors were supporting a different subset of primitives than their competitors.

Now, all the vendors support the VAAI fourth primitive, "thin provisioning stun." Thin provisioning is increasingly popular for VMware storage management of data stores and virtual disks. But you have to be careful with storage overcommitment, because you risk writing more blocks of data than you have actual disk space. If you run out of physical disk space, the VMs crash and halt. With "thin provisioning stun," you get dialog boxes warning of

- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

an imminent problem and VMs are “stunned” (stopped) rather than merely crashing.

Another thin provisioning issue that VASA resolves is space reclamation. When a file is deleted in Windows or Linux, it isn’t actually physically removed from the disk. Instead, the file is marked for deletion, and it is eventually destroyed by the creation of new files. In most case this isn’t a problem, but with thin virtual disks on thin data stores, it can result in thin volumes growing uncontrollably. That’s because when files are deleted, the free space is not gracefully handed back to the storage array.

As you might think, this situation does not play well for audit trails and chargeback. Could you imagine being a customer in a private or public cloud, where you would be paying for all the files you had as well as all the files you had deleted?

VMware really improved storage management in vSphere 5, but don’t be deceived, there are [numerous changes](#) all over vSphere that make it an epitome of innovation and technical engineering. VMware storage management is just the beginning.



- **Contents**

- **Managing storage for virtual environments: Best practices**
- **Top tools for virtual server storage management**
- **Using performance metrics to enhance storage for virtual environments**
- **Hypervisor-specific data storage management**

Free resources for technology professionals

TechTarget publishes targeted technology media that address your need for information and resources for researching products, developing strategy and making cost-effective purchase decisions. Our network of technology-specific Web sites gives you access to industry experts, independent content and analysis and the Web's largest library of vendor-provided white papers, webcasts, podcasts, videos, virtual trade shows, research reports and more—drawing on the rich R&D resources of technology providers to address market trends, challenges and solutions. Our live events and virtual seminars give you access to vendor neutral, expert commentary and advice on the issues and challenges you face daily. Our social community IT Knowledge Exchange allows you to share real world information in real time with peers and experts.

What makes TechTarget unique?

TechTarget is squarely focused on the enterprise IT space. Our team of editors and network of industry experts provide the richest, most relevant content to IT professionals and management. We leverage the immediacy of the Web, the networking and face-to-face opportunities of events and virtual events, and the ability to interact with peers—all to create compelling and actionable information for enterprise IT professionals across all industries and markets.

Related TechTarget Websites

- > [SearchDataBackup](#)
- > [SearchCloudStorage](#)

- **Contents**

- Managing storage for virtual environments: Best practices

➤ [SearchDisasterRecovery](#)

➤ [SearchVirtualStorage](#)

- Top tools for virtual server storage management

➤ [SearchStorage](#) ➤ [SearchSMBStorage](#)

- Using performance metrics to enhance storage for virtual environments

- Hypervisor-specific data storage management