virtual

DATA CENTER

Volume 33

TechTarget
Where Serious Technology Buyers Decide

One Hypervisor or Two?

DATA CENTER managers know that the choice of a hypervisor is central to their virtualization strategies. But as organizations gain experience in virtual environments, many are re-evaluating their needs and determining that two hypervisors may be better than one.

For example, an organization just beginning to deploy virtualization technology might choose a hypervisor that is simpler to use and offers fewer features. But as that organization gains experience, it may look for a more sophisticated hypervisor to meet its growing needs.

Knowing when to make the transition can be tricky. The right time to deploy a new hypervisor will depend on the unique requirements and circumstances of each organization. Stephen J. Bigelow describes the scenarios in which another hypervisor might fit the bill in "Time to Consider a Second Hypervisor?"

Before deploying virtual machines, it's critically important to understand how the hardware will be used. Underallocating resources to the VM can be detrimental to its performance, while overallocating resources may drain all others running on that host.

That's where proper resource allocation comes in. Read "Rev Up Virtualization Performance" for Bill Kleyman's tips on optimizing your virtual environment.

Packing more VMs onto fewer physical servers does wonders for server consolidation, but it can lead to oversubscription of CPU cycles, memory and other computing resources. Although higher consolidation ratios can increase the potential for saturation of core resources, it's not inevitable. There is plenty a data center manager can do to tweak performance and avoid VM contention before it gets out of hand.

Striking a balance is the goal. Mike Laverick addresses "right-sizing" as a science and an art form in "Managing Virtual Resource Oversubscription."

How do you right-size your virtual environment? Share your knowledge. Send an email to ccasatelli@techtarget.com. ■

CHRISTINE CASATELLI
Editor, *Virtual Data Center*

we're not just making servers. we're making server history.

While innovation comes rapidly in the IT industry, basic server architectures haven't changed for decades. That's why Cisco answered the need for innovation by introducing the Cisco Unified Computing System – which integrates compute, high-speed networking, storage access and virtualization in one system. Since its introduction, IT departments have dramatically reduced data center complexity while:

- Lowering operating costs by up to 30%.
- Reducing deployment times from weeks to minutes.
- Harnessing the power of over 30 world-record performance benchmarks.

The Cisco Unified Computing System, powered by intelligent Intel® Xeon® processors, signals the next evolution of the data center – where everything, and everyone, works together like never before.

Find out more at www.cisco.com/go/servers

together we are the human network. cisco

Time to Consider a Second Hypervisor?

ADDING A SECOND HYPERVISOR MAY BE KEY TO EXPANDING VIRTUALIZATION IN THE ENTERPRISE, BUT IT CAN ALSO OFFER SOME NEW CHALLENGES. **BY STEPHEN J. BIGELOW**

EVERY virtualization deployment starts with a hypervisor—which is software that forms the critical layer of abstraction between hardware and the operating system. Hypervisors allocate computing resources to each virtual machine (VM) and maintain the integrity of every VM within its own processing session.

It seems like virtualization technology has sprung up almost overnight, maturing from a noble experiment into a reliable and effective staple of enterprise computing. And virtualization vendors have made the most of this maturity, adding a suite of powerful features and tools that simplify and automate many of the associated tasks.

But as organizations expand their virtual environments, they're discovering that one hypervisor product isn't always enough for their data centers. Deploying a second hypervisor carries some potential benefits along with a few noteworthy risks. Before adding a hypervisor, it's important to carefully consider the pros and cons of a multi-hypervisor environment.

REASONS DRIVING THE SECOND HYPERVISOR

The rapid maturity of hypervisor technology begs an important question: If hypervisors are so well developed, why bother adding a second one to the environment? It turns out that there are several important factors.

In late 2010, TechTarget surveyed IT professionals about their virtualization decisions, and more than 50% of respondents said that cost was the principal factor in considering an alternative to VMware virtualization. More than 38% of 315 IT professionals responding to this question sought

a VMware alternative for more features and functionality, 33% wanted improved interoperability, and almost 30% wanted to avoid vendor lock-in.

The issues highlighted here are not really technical ones but rather a reflection of business concerns—a realization that one size does not fit all. For example, cost ranks high as a major consideration for enterprises. But in practice, it's the management costs that attract the most attention.

"The greatest concern that I see with users is not the actual hypervisor but the virtualization management strategy that they're going to deploy," said Tony Iams, vice president and senior analyst for system software research at Ideas International Inc. Iams said that x86 hypervisors share competitive costs and interoperability, so lock-in may not appear to be an immediate risk.

Still, virtualization management products can sometimes limit hypervisor choices. "That's really where the long-term lock-in is," he said.

Also consider the implications of interoperability, feature sets and performance. An organization with uniform x86-based servers may do quite well with a single hypervisor, but that single hypervisor may not run as well or offer all of the required features on mainframe, RISC or Sparc-based servers. This is a major driver to adopt another more suitable hypervisor for those systems.

Virtualization has emerged quickly in the enterprise. The speed of adoption may lead some organizations to adopt several hypervisors.

DIFFERENT HYPERVISORS FOR DIFFERENT FUNCTIONS

A more common example may be the use of one hypervisor for server virtualization and another hypervisor for desktop virtualization. Also, as private clouds become more commonplace, organizations may adopt a different virtualization platform to deploy their private clouds.

A second hypervisor may also be appropriate when an organization matures technologically. For example, an IT staff just starting a virtualization deployment might select a hypervisor that costs the least, offers fewer features or is simpler to use. As the staffers gain experience with virtualization and learn how to use various capabilities, they may eventually identify new business opportunities that can be better met with a more sophisticated hypervisor (see **Figure 1**, page 6).

Virtualization has emerged quickly in the enterprise. The speed of adoption may lead some organizations to adopt several hypervisors. "Chances are good that you could see independ-

ent technical paths in different groups," he said, adding that one group might use a certain hypervisor for test and development, while another group might use a different hypervisor for production. Without a mandate to standardize on a single hypervisor, this can lead to interoperability issues or other problems if an organization needs to bring together the resources of those groups in the future.

Others aren't so sure that the speed of adoption is a driving factor. "Just having more choice in the market is obviously going to drive people to look at second platforms," said Gary Chen, research manager for enterprise virtualization software at IDC. Chen pointed to the search for better prices, particularly compared with long-time leaders in the market such as VMware. "These alternatives to VMware have certainly gotten more looks from people."

TESTING AND DEPLOYMENT

Introducing a second hypervisor into an environment requires careful plan-

FIGURE 1

WHY CHANGE HYPERVISORS

Respondents to a TechTarget survey offer reasons why they would consider a second hypervisor.

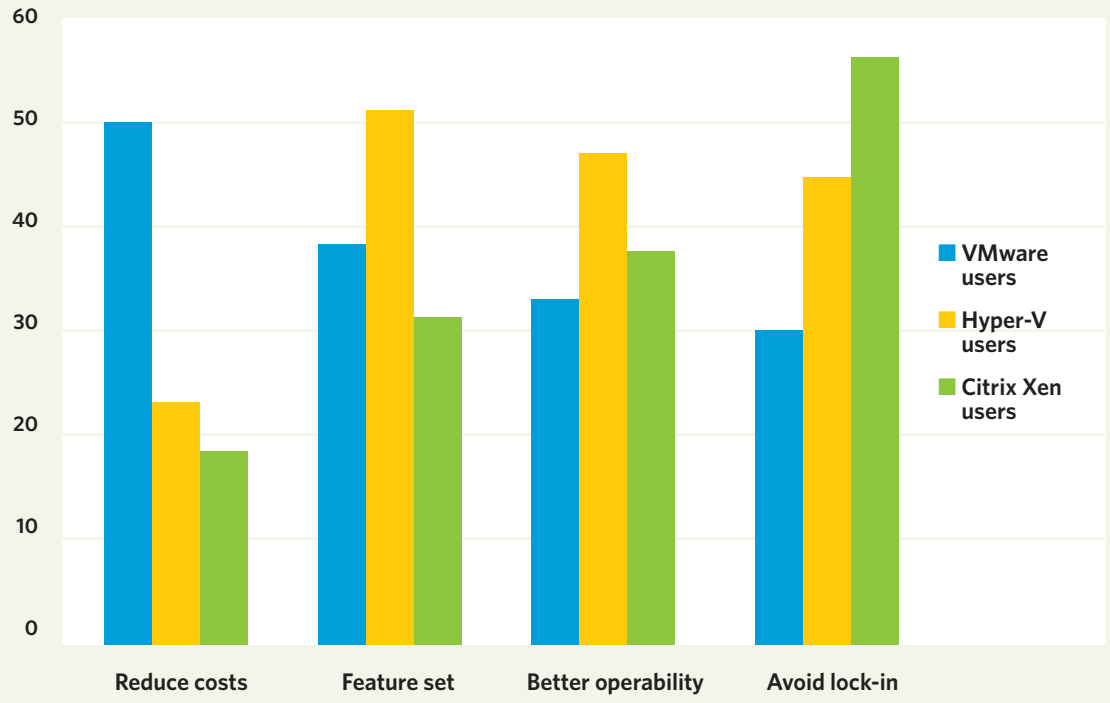

ning and testing that should start long before an actual rollout. As with the current hypervisor, a second hypervisor must fully support the intended server hardware platforms, device drivers, the operating system versions—and, ultimately, the applications.

If a third-party hypervisor management tool is in place, it might be necessary to investigate its compatibility with the new hypervisor as well. If the same management tool can handle both hypervisors, this can significantly simplify the management process (see **Figure 2**).

Figure 2

THE CHALLENGES OF HYPERVISORS

A range of issues can make system management difficult for hypervisors, according to a TechTarget survey.

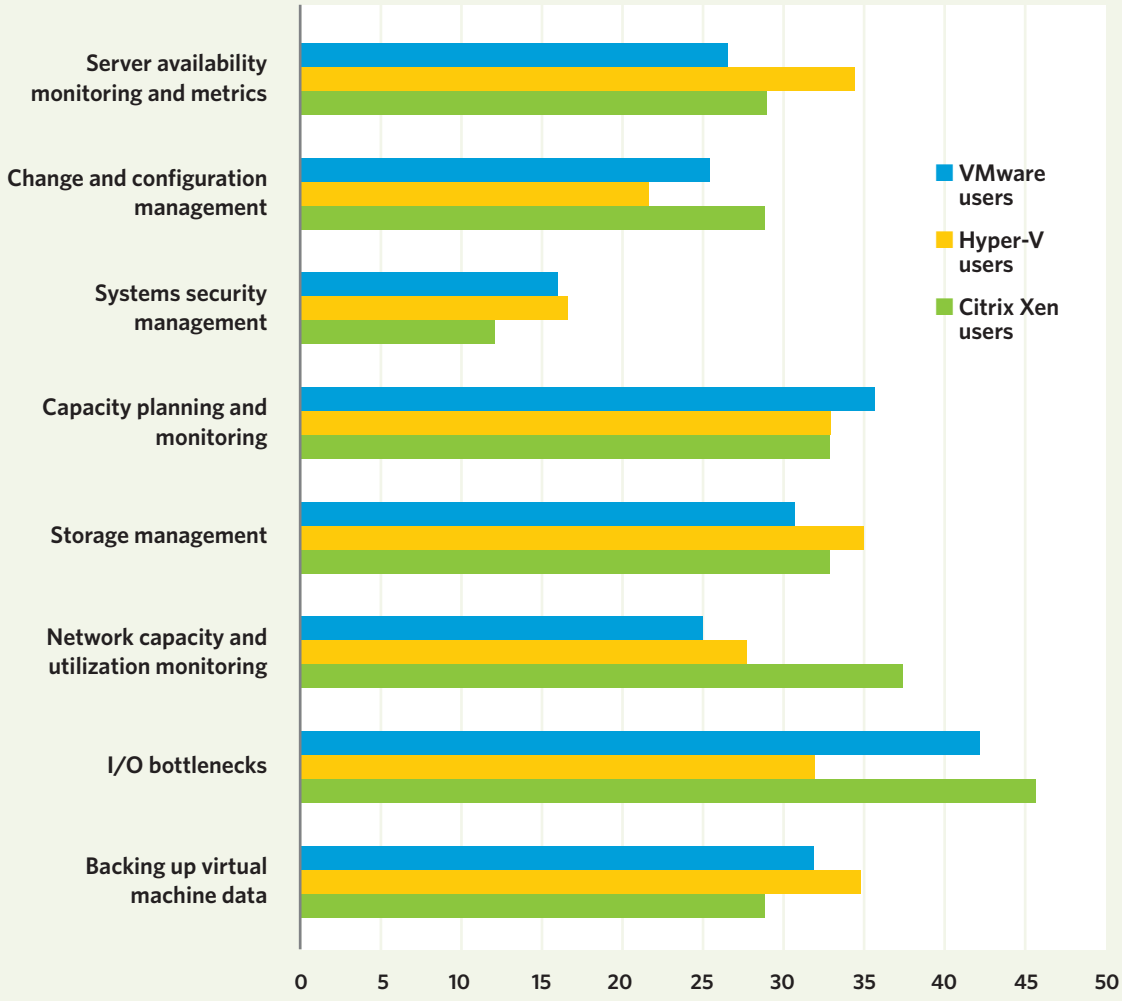

The new hypervisor's compatibility list is the place to start researching that information, but there is no substitute for hands-on testing in a lab environment. Once initial testing at the server level has been completed, expand the testing scope to verify network and storage compatibility under the application's workload.

NAVIGATING THE LEARNING CURVE

Testing a second hypervisor can provide an opportunity for IT staff to become proficient at installing and managing the new hypervisor. There are bound to be similarities, but understanding the differences and nuances among hypervisors can eliminate confusion and streamline the

THE CASE FOR COST

CHOOSING A hypervisor always involves a consideration of cost. In reality, the cost of a hypervisor is extremely competitive. They're virtually free to try, and even Hyper-V is included with the Windows Server operating system. But there's a lot more to a hypervisor's total cost of ownership (TCO) than just acquiring a software license.

Hypervisors include a rich suite of features. But using a complete set of enterprise-class features along with the management tools needed to keep tabs on a virtual data center can lead to a significant price tag. "Cost still matters," said Tony Iams, vice president and senior analyst with system software research at Ideas International Inc. "There's cost of acquisition, cost of operation and cost of ownership," he said. Running multiple hypervisors also imposes a cost on the IT resources that support and maintain a number of concurrent platforms, Iams added.

Organizations that opt for a second hypervisor will need to make critical cost choices. For example, the value of a second hypervisor's features and capabilities—such as more efficiency or better performance—should justify the costs associated with the new platform. A justification that makes sense for one organization may not make financial sense for another. So each organization needs to understand the TCO for a second hypervisor and make its own informed determination of whether a second platform is worth supporting.

Something else that executives and IT administrators often overlook is the replacement cost when a second hypervisor replaces an existing platform. This can happen when an organization has outgrown the features of its existing system and must move to a more sophisticated virtualization platform. Each organization must consider the costs of potential business disruptions as well as other risks—such as data loss or impaired performance for some virtual machines—involved with a platform transition. ■

deployment. Some organizations may choose to certify administrators in the new hypervisor, but personnel who are already fluent with an existing hypervisor can usually achieve proficiency with some hands-on experience.

In the early days of virtualization, hypervisor rollouts were usually ad hoc, with IT shops installing the platforms on various servers on an as-needed basis. Today, the rollout process is far more formal in most cases. This is mostly because there is a hypervisor already in use, so knowing where to install a second is particularly important.

Another factor is that organizations seriously pursuing multiple hypervisors are usually large enough to require a level of formality and documentation in new platform deployments. It's definitely coming from higher level initiatives around dynamic data center, private and public clouds, Chen said. "We're going to go through a vetting process step by step."

WHEN IS THE RIGHT TIME?

Knowing when to switch from one hypervisor to another can be nebulous. There is rarely a single trigger that precipitates a switch to the new platform. The right time to deploy a new hypervisor will depend on the unique needs and situation of each organization. For example, if cost is a primary driver, the time to switch may simply be the point at which administrators feel comfortable deploying and managing the new platform.

In other cases, the trigger may be the need for new features or services that the old hypervisor just doesn't provide. Or maybe there's a service-level agreement requirement that a new hypervisor can help resolve. What's important is to understand the needs that are driving a second hypervisor and consider the timing or other situations that will push the new hypervisor into production.

Today, hypervisors are typically well developed, stable and reliable software platforms, but it's important to remember that there are risks in moving to a new hypervisor. The biggest risk is organizational resistance. For example, suppose management forces IT to adopt a hypervisor that is radically different from the previous platform. If there isn't enough time for administrators to test and master the new hypervisor or if IT gives up features and capabilities that the old hypervisor provided, there may be a backlash that undermines the expected benefits.

The development of isolated silos is another risk. For example, a Windows part of the organization adopts one hypervisor, and a Linux part of the organization adopts another hypervisor. Organizations that manage the hypervisors separately often fail to optimize the use of technology and IT expertise. "Today people are really looking at moving toward a private cloud where things are abstracted and IT delivers a service," Chen said. "Having isolated pools is counter to that. ■

An intuitive, integrated IT lifecycle management solution for the control freak side of you.

Managing your IT environment means challenges coming at you from all sides. You've got to track assets, distribute software, manage costs and stay compliant with licensing — all while maintaining airtight network and data security. The last thing you need is a complex set of tools to manage it all. • That's why we developed Hitachi IT Operations Director. • As one of the world's largest software providers, we know the best solution needs to be both powerful and intuitive. IT Operations Director is an integrated IT lifecycle management solution that gives you control over all of your IT assets and costs, automates tasks and helps you maintain the security of your systems and data. • So let the control freak side of you (not to mention the efficiency expert and risk-averse sides) experience the power and ease of use of IT Operations Director with a 30-day free trial. Download it today at **www.itoperations.com.**

HITACHI IT OPERATIONS DIRECTOR

Covering all sides of your IT: asset management, security management and software distribution.

Take control of your IT today with a 30-day trial offer!
www.itoperations.com

Hitachi Data Systems

HITACHI
Inspire the Next

Rev Up Virtualization Performance

PROPER RESOURCE ALLOCATION IS VITAL TO THE PERFORMANCE OF PHYSICAL HOSTS AND VIRTUAL MACHINES. **BY BILL KLEYMAN**

OVER THE past 10 years, virtualization technologies have allowed IT engineers to consolidate, save money and help their businesses grow. But just how much of an effect has virtualization had? Any engineer working with modern server technology would be foolish not to use the resources on these machines for business growth and expansion.

By placing numerous virtual machines (VMs) on top of a physical host, data centers are running more streamlined and have a smaller hardware footprint. A new Hewlett-Packard Co. ProLiant Server with an Intel Corp. eight-core processor is capable of running five to six VMs without much trouble. So why hold on to dying hardware when it can be virtualized, mothballed and recycled?

As virtualization gets easier, network engineers are becoming increasingly comfortable with the technology. But there is a downside to this ease of use. Lately, many companies have begun to follow this scenario: buy a server, load the workloads on a storage area network (SAN), deploy the hypervisor and load the VMs. Although this approach may work, IT administrators are beginning to skip important steps that can hurt virtual server performance.

UNDERSTANDING VIRTUALIZATION TECHNOLOGY

Generally speaking, there are two types of virtualization technologies—hosted and bare metal. A hosted environment uses a server with a preloaded operating system on it, such as Windows Server 2008. Once that OS is installed, an administrator would load virtualization software on top of the original OS.

Bare-metal architecture removes the OS from the equation and uses a Linux/Unix-based kernel as the hyper-

visor. The virtualization technology is installed right on top of the hardware. VMware offers both types of virtualization technologies, while XenServer from Citrix Systems works around a bare-metal hypervisor.

The biggest advantage of using a bare-metal design is that there is no extra software that VMs have to go through to access base hardware resources. But a bare-metal deployment requires newer hardware because the onboard processors must be either Intel-VT- or AMD-V-ready. This means that administrators looking to upgrade older servers will not be able to use this technology. Still, there is always the hosted hypervisor at their disposal.

The biggest advantage of using a bare-metal design is that there is no extra software that VMs have to go through to access base hardware resources.

THREE CORE UPGRADE PATHS

Proper resource allocation is vital to the performance of both the physical host and the VMs. Prior to deployment, it's imperative to understand what the hardware will be used for. Is the environment running high-end SQL queries nonstop? Or is the company planning on hosting simple applications for a few users to access on an infrequent basis? By understand the goal, an administrator can roll out an environment capable of handling the load and capable of growing with the infrastructure.

There are three core upgrade paths when it comes to physical hardware:

1. **Hard drive.** There is little debate that by upgrading a machine's hard disk with a faster drive, an environment will see an increase in VM performance. If a given environment does not use a SAN as a central point for workloads, consider upgrading with several faster disks. Smaller IT shops that do not require the use of a centralized storage array work with onboard space requirements built into their physical hardware. That is, by upgrading the RAID array with better and higher performance drives, the environment will see a boost in performance and redundancy.

For larger deployments where a SAN is present, consider the technology on hand. Is the SAN older? Do the drives on it spin fast enough to allow quick and seamless access to a workload? Often, IT engineers will won't replace a SAN and then wonder why their virtual infrastructure is running slowly, despite having a new server and new virtualization software. Although this can be a bit of an investment, using older SAN technology can be detrimental to a new virtual environment.

2 CPU. As VMs are loaded onto a physical box, the onboard processor is used more. From an IT engineer's perspective, a faster CPU will always provide faster processing. Physical hosts have the capacity to have their

BEST PRACTICES FOR VM PERFORMANCE

TO ENHANCE the performance of a virtual machine (VM) or physical host, remember the following points:

- **USE NATIVE SOFTWARE VIRTUALIZATION TOOLS.** Both XenServer and VMware come with paravirtualization technology. That means they will both have Xen tools or VMware tools to allow better VM usage of OS resources. For example, VMware will have its own network interface card driver to better communicate with the underlying hypervisor. These tools come with most virtualization platforms and should be installed as soon as the VM is deployed.
- **MAKE SURE THAT DISKS ARE ALIGNED.** For administrators running a Windows OS older than Windows Server 2008 or Vista, disk alignment may be a problem. Server 2003 and XP VMs will format their virtual disks and will be misaligned with the underlying storage by default. This will cause performance degradation in VMs. To avoid this, it's a good idea to brush up on how to fix disk alignment.
- **MAKE SURE THE ANTIVIRUS SOFTWARE IS CONFIGURED CORRECTLY.** By allowing real-time virus scanning to run on a VM, an engineer runs the risk of seeing a host OS experience a severe speed decrease. Optimizing your antivirus configuration will significantly improve VM performance.
- **MAKE SURE THAT ALL FIRMWARE AND BIOS UPDATES ARE APPLIED** when deploying a new physical host. Many times, little tweaks are deployed by manufacturers that have big effects on VMs running on that hardware.
- **NEVER FORGET TO RUN REGULAR MAINTENANCE ON THE VM ENVIRONMENT.** Make sure to defragment the virtual disks as necessary, and apply any needed updates from both the host OS and the virtualization software.
- **REMOVE VISUAL EFFECTS.** Just like a regular OS, VMs function better without visual effects. By removing these little settings, a VM will operate faster.

Tasks such as constantly monitoring VM resources and keeping up with simple preventative maintenance can easily be overlooked. Always take the necessary time to ensure that both VM and physical host are healthy. ■

CPUs upgraded or even added to. Many machines come with open CPU slots that are capable of expanding.

3 **RAM.** Upgrading RAM is probably the most cost-effective way of tweaking a virtual host's performance. By upgrading memory on a host, you will be able to allocate more RAM to each VM. Any server-class machine is capable of handling more RAM than it is usually shipped with. When new RAM is added, an engineer can now re-examine how the VMs are using resources. From there, additional memory can be allocated to a given machine to improve performance.

VM SIZING IS CRUCIAL

A VM requires resources delivered from the underlying hardware to function. The more resources given to it, the better and faster it will run. This is where appropriate VM sizing is absolutely crucial. Underallocating or overallocating resources to the VM can be detrimental to the performance of not only that VM but also all others running on that host.

"Verifying the appropriate resource to give to a VM is an important ongoing process," said Tim O'Brien, systems consultant at MTM Technologies Inc. "When working with live machines, an engineer should always monitor how much utilization is happening on that VM. Take the time to dedicate enough resources to a workload, and remove resources when they are not being used," he said.

When an environment is supporting multiple physical hosts with numerous VMs in the infrastructure, resource management becomes even more important. "Virtual technology has come a long way in how it deals with resource usage," O'Brien said. "I always recommend using dynamic memory allocation to more efficiently use resources and improve overall VM density."

Underallocating or overallocating resources to the VM can be detrimental to the performance of not only that VM but also all others running on that host.

Many environments now have built-in failover capabilities where other hosts are able to take virtual workloads away from failed physical servers. This is where an overallocated VM can have a lot of trouble failing over to a device that doesn't have the resources to share. ■

Managing Virtual Resource Oversubscription

IDENTIFY THE SYMPTOMS OF RESOURCE OVERSUBSCRIPTION, AND LEARN WAYS TO LIMIT IT OR EVEN PREVENT IT. BY MIKE LAVERICK

IN TODAY'S data centers, organizations are packing more virtual machines (VMs) onto fewer physical servers. That's great for server consolidation, but it could lead to oversubscription of CPU cycles, memory and other computing resources. When considering consolidation, know what the symptoms of resource oversubscription are and how organizations can limit it or even prevent it.

The consolidation of servers and desktops raise concerns that concentrations of workloads bring with them inevitable risk of downtime and the potential of oversaturation of the core computing resources—memory, processor, network and disk.

In the ideal world, a hypervisor should be consuming these resources as much as possible, while leaving headroom to accommodate the growth in the VMs as well any unexpected surges in workload. Expressed simply, a hypervisor consuming only 1% of memory, CPU, network or disk is underutilized. Likewise, a hypervisor running at 99% of memory, processor, network or disk is likely to provide poor performance and be a major bottleneck in any clustered environment.

STRIKING A HAPPY BALANCE

It's the administrator's job to strike a happy balance between these extremes. It's as much an art as a science, and successful administrators know their environments and are able to work with app owners to resolve any performance problems.

VMs don't typically cause most performance problems. It's badly configured applications that are usually the culprits. The virtualization layer always takes the heat on this because application owners will say that the service worked brilliantly on a physical server. When they say this, they are subconsciously betraying the fact that they are skeptical of virtualization to some degree.

Unless the VM was created by a physical-to-virtual (P2V) process, it is likely that the way the application was configured does not mirror the physical system. That said, most customers paradoxically experience improvements in performance because virtualization projects often bring in new and improved hardware at the server and storage layers.

MEMORY IS KING

The biggest single constraining resource in virtualization is memory. This is the resource that environments run out of before they run out of CPU cycles or bandwidth to the network or storage array.

To avoid this, begin by "right-sizing" your VMs relative to the demands of your application. This means resisting the demands of application owners who request VMs with the same specifications as the physical server.

A dose of reality is needed here. It's totally unrealistic to think a Tier 1 application such as Microsoft Exchange, SQL or Oracle will sit happily with just the allocation needed to make a 64-bit operating system. In general, most environments are risk-averse, and administrators have a tendency to over-spec the VMs in hopes that they will not experience any blowback from disgruntled application owners.

This type of approach to resource allocation should be avoided at all costs. VMs that go beyond spec cost the environment in wasted resources that could have been allocated to more deserving VMs. It also systematically and unnecessarily degrades the performance of features and wastes resources elsewhere.

WASTED RESOURCES

Another area to review is any system that has been converted to a VM through the process of P2V. More often than not, IT folks choose not to downgrade the memory allocation, leaving the VMs with the same allocation as the original physical machine. This can contribute to a massive waste in resources.

In general, most environments are risk-averse, and administrators have a tendency to over-spec the VMs in hopes that they will not experience any blowback from disgruntled application owners.

Remember why you virtualized in the first place? You had many physical systems that were using only 10% to 20% of their resources, taking up space and power in the data center.

If you think you are experiencing memory problems, check the following areas:

- Does the VM have enough memory allocated?
- Has the physical server run out of memory?
- Is there swap activity taking place inside the guest operating system and at the hypervisor level?
- Are there unusually high statistics in the hypervisor's memory management systems?

CPU BOTTLENECKS

In some environments CPU bottlenecks can happen. In these situations, the memory payload of the guest OS and application is small, but the service is carrying out a high volume of transactions per second.

Despite all the fancy footwork that modern hypervisors bring to the table, it is still the case that one vCPU VM can send threads to only a single CPU core. In this respect, raw performance is the number of cycles per second that core can provide.

Bear in mind that with virtualization it is still unlikely that a VM with one vCPU would gain exclusive access to a core. It's more likely that the VM will have to share the CPU with other VMs.

If this sharing is allowed to continue unchecked, it is possible that the CPU could become saturated with requests. In this scenario, CPU contention takes place. Fundamentally, if a VM needs more CPU cycles than the core can provide, the only way to deliver CPU cycles greater than 100% is to configure the VM with two or more vCPUs to deliver true symmetric multiprocessing.

Before you do that, though, you need to decide if the CPU is the constraining resource and confirm that the application within the guest operating system is multi-threaded. There's little point in giving the VM two vCPUs if threads are executed only on CPU0 while CPU1 is still there twiddling its thumbs.

Despite all the fancy footwork that modern hypervisors bring to the table, it is still the case that one vCPU VM can send threads to only a single CPU core.

As time goes on, you might find yourself running on the more modern Intel Nehalem architecture. Studies have shown that a single Intel Nethalem core can actually outperform a SMP-enabled system using the older CPU types. To identify CPU bottlenecks, investigate the following areas:

- Using your hypervisor, identify if any VMs are using 100% of the CPU. Avoid using the guest oper-

ating system tools in the VM, and look for performance data delivered by the hypervisor—it will be more accurate.

- Look for high %Ready values in VMware ESX because this is an indication that your VM would like more CPU cycles but isn't receiving them.

- Look for high co-stop values because this can show excessive use of SMP in your VMs.

That last tip needs some explanation. Sometimes administrators go overboard in giving every VM more than one vCPU as a standard—even when it's not entirely necessary. If you use virtual SMP excessively, you can give the hypervisor more work than it's expecting. It has to work harder to schedule multiple vCPUs across mul-

SOFT NETWORK OPTIMIZATION

ONE WAY to stop network oversubscription is to buy new hardware and implement 10 Gbps networking for all the high-level, bandwidth-intensive traffic. Another way is through soft network optimization, which can be configured in most modern hypervisors and within the VM. Consider modifying and enabling the following features:

- **If your hypervisor and physical network support a larger maximum transmission unit, commonly referred to as Jumbo Frames, enable their use.**

- **Investigate if your onboard network interface card supports iSCSI offloading in the BIOS.**

- **Configure the requirements for iSCSI Port binding on your vSwitches to take advantage of true multipathing I/O.**

- **Ensure that software integration tools from virtualization vendors are installed and up to date because they often provide optimized network drivers to the guest operating system.**

- **Consider using network I/O controls within the hypervisor to allocate bandwidth more intelligently to environments using 10 Gbps.**

The use of a 10Gbps network can greatly simplify your virtual switch configuration. It creates one master virtual switch to handle all network traffic rather than having a complicated series of switches that handle discrete traffic types. ■

tiple cores inside the CPU socket.

What can happen is this situation can actually increase the very contention you were trying to avoid in the first place. So strictly control the use of virtual SMP, especially on the modern CPU architectures where its benefits may be limited.

NETWORK BOTTLENECKS

It's a common misconception that with many VMs sharing the same physical networks, network bandwidth would become scarce. As with CPU resources, most networks work in a nonlinear way so that many systems can coexist on the same network without treading on each other's toes.

In most environments you will struggle to see VMs saturate even a 1Gbps interface. It's common that these interfaces are teamed together for fault-tolerance and load balancing. So if it's unlikely in a 1Gbps environment, it's even less so when you have bundles of network interface cards.

The reality is that your network bottlenecks are more likely to be seen during the process commonly called live migration in which VMs with large memory allocations are moved from one physical server to another. Again, right-sizing your VM's memory allocation is critical in reducing the performance hit that live migration brings. So it's important to have dedicated gigabit-and-above network interfaces dedicated to this ancillary process.

STORAGE BOTTLENECKS

Storage oversubscription is usually caused by simple administrative errors. What's surprising are the I/O demands that virtual desktop projects can sometimes impose at the storage layer.

It's not uncommon to have storms of storage and CPU I/O caused by "boot storms" and antivirus scanning.

As with CPU resources, most networks work in a nonlinear way so that many systems can coexist on the same network without treading on each other's toes.

Storage vendors can help by allowing customers to purchase caching modules that add solid-state storage to elevate the I/O chokepoints in a virtual desktop infrastructure (VDI).

Adequate planning needs to occur at an early stage so the costs of scaling up a VDI solution are exposed early. The tipping point for the use of such caching technologies appears to be 500 to 600 VMs.

Below this point, simply distributing the virtual desktop around various arrays, LUNs and spindles appears to be enough. Beyond the 500 to 600 VM range, businesses should serious-

ly begin considering solid-state solutions as a way of taking the disk spindle out of the equation.

This caching approach can also elevate some of the IOPS generated by routine VDI tasks such as creating and destroying virtual desktops as users log out or join the system.

For server-based VMs that are disk I/O-bound, many simple tasks can be used to boost performance. Although it's perfectly fine for 10 or 20 VMs to occupy the same volume/LUN, it is reasonable for disk I/O-bound VMs to dedicate a volume/LUN to specific application or merely reduce the ratio of VMs to a data store to increase the available disk IOPS shared among the competing VMs and shorten the disk queues to storage. Other optimization techniques include the following:

- When possible, use the hypervisors paravirtualized SCSI controller inside the VM.
- Distribute the virtual disks of the VM across multiple volume/LUNs to ensure virtual disks do not compete against each other for I/O.
- Adopt permissions on data stores to prevent rogue administrators from creating VMs in the wrong location by sorting storage according to the amount of free space without considering the IOPS needs of the VM.

Although it's true that higher consolidation ratios increase the potential for saturation of core resources, it's not inevitable. Today's hardware and software are keeping up with these increased resource demands.

Although it's true that higher consolidation ratios increase the potential for saturation of core resources, it's not inevitable.

In the world of virtualization, memory continues to be key constraint. But there is plenty an administrator can do to tweak performance and avoid VM contention before it gets out of hand.

Eventually the constraint on consolidation ratios may be more about businesses feeling anxious about putting too many eggs in one basket. In this respect it could be that an availability gap is opening. Although hardware and hypervisors increase resource capabilities, the risks of very high consolidation ratios remain despite the widespread use of hypervisor clustering, fault tolerance and in-guest service protection tools. ■

ABOUT THE AUTHORS

Stephen J. Bigelow, a senior technology editor in the Data Center and Virtualization Media Group at TechTarget Inc., has more than 15 years of technical writing experience in the PC/technology industry. He holds a bachelor of science in electrical engineering, along with CompTIA A+, Network+, Security+ and Server+ certifications, and has written hundreds of articles and more than 15 feature books on computer troubleshooting, including *Bigelow's PC Hardware Desk Reference* and *Bigelow's PC Hardware Annoyances*. Contact him at sbigelow@techtarget.com.

Bill Kleyman, MBA, MISM, is an avid technologist with experience in network infrastructure management. His engineering work includes large virtualization deployments as well as business network design and implementation. Kleyman is the director of technology at World Wide Fittings Inc., a global manufacturing firm with locations in China, Europe and the United States.

Mike Laverick is a professional instructor with 17 years experience in technologies such as Novell, Windows and Citrix. Involved with the VMware community since 2003, Laverick is a VMware forum moderator and member of the London VMware User Group Steering Committee. He is also the owner and author of the virtualization website and blog RTFM Education, where he publishes free guides and utilities aimed at VMware ESX/Virtual Center users.

TechTarget
Where Serious Technology Buyers Decide

Cathleen Gagne
Editorial Director
cgagne@techtarget.com

Christine Casatelli
Editor
ccasatelli@techtarget.com

Linda Koury
Director of Online Design
lkoury@techtarget.com

Marc Laplante
Publisher
mlaplante@techtarget.com

Peter Larkin
Associate Publisher
plarkin@techtarget.com

TechTarget
275 Grove Street
Newton, MA 02466
www.techtarget.com

- VMware Acceleration Kits
- VMware vSphere Case Study
- VMware vSphere Demo

See ad page 3

- Gartner: Fabric Computing Poised as a Preferred Infrastructure for Virtualization and Cloud Computing
- Data Center Management: The Key Ingredient for Reducing Server Power While Increasing Data Center Capacity
- Exploring the Potential Benefits of End-to-End Convergence of Data Center Networks

HITACHI
Inspire the Next

See ad page 10

- The Value of Integrating Availability and Performance Management and Service Desk Tools
- Citizen Machinery Achieves Real Time Asset Management Accuracy with Hitachi IT Operations Director
- Manage the Entire IT Environment with One Easy-to-use Product