virtual

DATA CENTER

Volume 28

TechTarget
Where Serious Technology Buyers Decide

When Less Is More

MORE IS BETTER, right? Not in the case of server consolidation. We all know that consolidation increases the use of available computing resources, but too much consolidation can be bad for your data center. Even for today's virtualization-friendly servers, limits are necessary.

The problem with over-taxed servers is the effect it has on the rest of your data center. When that server fails, you can't start up the affected virtual machines (VMs) on other servers because there is no computing capacity available. As a result, all of the affected VMs remain offline until the ailing server is repaired and the VMs are restarted.

As with most things, moderation is key. Stephen J. Bigelow makes the case for restraint in "Achieving the Right Level of Server Consolidation." Adding more VMs to servers can potentially affect network connectivity as well. The secret to optimizing the network layer is making sure your configuration doesn't create a situation where critical resources become oversubscribed. There is plenty that can be done to improve and monitor network performance for VMs as your consolidation ratios grow. Learn Mike Laverick's tips and tricks in "How to Optimize Network Performance to Improve Server Consolidation" for better results in your virtual data center.

Maybe you're testing the cloud-computing waters. How do you know if your organization is ready to take the plunge? To become cloud-ready, you'll probably have to add features to your virtual environment to provide greater availability and automation. Don't forget additional features to measure resource usage for charge-back and to offer self-service capabilities. Read "Get Ready for a Private Cloud" by Eric Siebert to figure out if your virtual environment can meet the demands of a cloud architecture—and what changes you'll have to make if it can't.

Are you pushing the limits of consolidation in your data center? When is enough enough? Tell us what you think—send an email to ccasatelli@techtarget.com. ■

CHRISTINE CASATELLI
Editor, Virtual Data Center

Achieving the Right Level of Server Consolidation

TOO MUCH CONSOLIDATION CAN POSE SERIOUS PROBLEMS FOR ENTERPRISES. STRIKE A BALANCE TO IMPROVE COMPUTING EFFICIENCY. **BY STEPHEN J. BIGELOW**

IF YOU'VE ALREADY worked with server virtualization, chances are that you understand the importance of consolidation. It's probably the single most important consideration in a virtual data center.

Simply stated, consolidation increases the use of available computing resources and allows more virtual machines (VMs) to operate simultaneously on a physical host system. But there are practical limitations to consolidation—even on today's most powerful and virtualization-friendly servers. Too much consolidation is not a good thing, and administrators involved in any virtual environment need to consider the serious implications of excess consolidation on their data centers.

The role of consolidation has become so ubiquitous that it's easy to forget why we do it in the first place: Consolidation saves money.

Consider traditional nonvirtualized environments where one application operates on one server. The server rarely uses more than 10% of its total computing resources for the application, and each new service or application brought into the environment incurs the expense of a physical server, network support, power, cooling, maintenance and so on.

Virtualization encapsulates workloads and lets multiple workloads reside together on the same physical server, allowing administrators to use considerably more of the server's CPU, memory and I/O resources. This translates into fewer physical servers with fewer power and cooling requirements.

In addition, workloads can be moved between physical servers using live migration, allowing real-time workload balancing and minimizing application downtime for hardware maintenance and repairs. Even the licensing schemes for Windows Server 2008/R2 Data Center Edition make hosting VMs on the same server more cost-effective than ever.

Taken all together, consolidation initiatives can vastly improve computing efficiency and present a significant cost savings for modern enterprises.

TOO MUCH SERVER CONSOLIDATION

As with so many other things, too much server consolidation is not beneficial for data centers or their users. But over-consolidation is becoming a

WHEN TO AVOID CONSOLIDATION

VIRTUALIZATION PLATFORMS HAVE evolved a great deal over the last few years, and the major offerings from VMware, Microsoft and Citrix Systems can support almost any kind of workload. As a general rule, all modern applications can readily exist as a virtual machine (VM). Still, IT administrators would be wise to exercise caution when planning to move physical workloads into the virtual domain.

Old applications can be problematic—especially applications that are custom-built or otherwise rely on direct interaction with specific hardware. Because virtualization imposes a layer of abstraction between the application and the underlying hardware, any applications that "need" access to specific hardware may malfunction or experience unacceptable performance issues.

One solution is to update the application using newer programming languages and techniques to create a more hardware-agnostic version. However, such upgrades can be costly and time-consuming. Similarly, it might be possible to replace the custom application with a commercial product that you can customize in-house to accomplish the same tasks. But the time and effort needed to customize the commercial application may sometimes be more than the resources needed to update the existing application.

In actual practice, it's often easiest to simply leave the custom application running on a non-virtualized physical server. And just because an application can exist as a VM does not mean it should—or even could—be consolidated.

Consider a demanding application like SQL Server or Exchange Server. It is possible to run SQL or Exchange as a VM, but it's doubtful that you would get acceptable performance out of the application if it's forced to share computing resources with 10 or more other VMs on the same physical machine. Demanding or resource-intensive applications are best approached with minimal levels of server consolidation.

Testing is an important part of the consolidation process and should be performed in a lab environment that doesn't impinge on the production data center. Testing ensures that the application is suitable for a virtual environment, verifies the computing resources required, provides insights into performance and checks interoperability in the virtual environment. ■

more familiar phenomenon as organizations seek to stretch their resources to the limit. The problem is that virtualization is too easy.

In years past, the addition of a new application or service meant a capital expenditure for the server and the labor to install it. The group or department requesting the expenditure faced financial scrutiny and had to budget for their decisions, and it might have taken weeks—or even months—to implement the deployment.

Some organizations overburden their servers as a matter of standard practice, with the goal of using 100% of the server's resources.

Virtualization completely changes the paradigm. Today's companies can provision a new VM on an existing server in a matter of minutes. There's no new server hardware to buy or install, and the only immediate costs involve the operating system and application licenses.

The desire to make IT faster and more responsive has fostered an "on-demand" climate where computing resources are perceived as a free commodity that is easily depleted with little—if any—regard for the actual implications on the business.

Some organizations overburden their servers as a matter of standard practice, with the goal of using 100% of the server's resources.

"If I'm paying for a four-socket box, and I've already paid for my [Windows Server] Data Center licenses, how many servers can I get on that box?" said Todd Erickson, president of Technology Navigator, which provides business intelligence to the financial industry.

Other organizations bloat their servers by accident, fitting new VMs onto any server with the available resources to make it work—but without any regard to the business implications involved. In reality, both approaches are doomed by the serious consequences of server over-consolidation.

Application performance and stability are the first casualties of an over-consolidated server as VMs compete for scarce computing resources. Every application on the server can be affected to one extent or another, including backup, disaster recovery (DR) and other data protection tools.

Less extreme cases may cause only hesitation in the application, while more extreme situations may crash one or more VMs—or even crash the entire server. It's a risk that most administrators choose to avoid because of the corresponding impact on business revenue, customer access and satisfaction, and the potential for data loss with its compliance implications.

A host failure with a high consolidation ratio on a single host can affect a lot of VMs when it goes down. And once a failure occurs, all of those affected VMs have to restart—either on the original server or across one or more other servers in the data center. The recovery process can place significant stress on the entire virtual environment.

Over-consolidation also impairs live migration capabilities within the environment. Even though most administrators don't allow automatic migration, the ability to move workloads on-demand is an essential benefit of virtualization. When servers are taxed to their limit, however, it's almost impossible to move workloads. Just consider what happens when a server fails. You can't start up the affected VMs on other servers because there is no computing capacity available, so all of the affected VMs remain offline until the server is repaired and the VMs are restarted.

Ultimately, many experts suggest remaining within a moderate level of server consolidation where only about 60% to 70% of a server's resources are normally used. The actual percentage will depend on your particular business situation, but the underlying goal is to leave some amount of computing resources unused.

Businesses still benefit from greatly improved computing efficiency—up from 5% to 10% utilization in a nonvirtualized environment—while maintaining enough reserve resources to restart VMs without straining the server. In addition, the remaining computing resources allow VM migration between servers to balance workloads or support maintenance efforts.

EASING EXCESSIVE SERVER CONSOLIDATION

One of the easiest ways to prevent excessive consolidation on a server is to implement suitable IT practices right from the outset. Erickson points out the dangers in over-committing the server's available computing resources—a practice he calls "thin provisioning" the server. For example, vSphere and XenServer both support a memory over-commit feature that allows an administrator to provision more memory than is physically available on the server.

A host failure with a high consolidation ratio on a single host can affect a lot of VMs when it goes down.

"Nobody ever does thin provisioning as a best practice," Erickson said. "If you're doing any kind of thin provisioning, you're probably already bumping up against a consolidation ceiling." The problem is that exhausting an over-committed resource such as memory can potentially impair

the VM's performance or stability, he said.

And even when a company embraces over-committing resources as a practice, the amount of over-commitment that is considered acceptable is always prone to increase as more VMs are crammed onto physical servers.

For example, a server with 48 GB of physical memory and 52 GB of allocated memory—4 GB of memory is over-committed—may seem acceptable at about 10% over-commitment, but that server is over-consolidated. "You're setting the stage for problems," said Erickson, adding that the organization inevitably accepts increasing levels of over-consolidation, further increasing the risk of failures over time.

Suitable management tools can help identify over-consolidated servers and allow administrators to forestall problems before they arise. Experts emphasize that proactive management allows proactive development of the environment—finding and fixing potential resource problems before they even arise. No IT department should have to run out and buy a new server in response to computing resource shortages.

"At the end of the day, you've got to be constantly looking into your management console and understand what you're running at for resources," said Scott Roberts, director of

TRACKING CONSOLIDATION AND PERFORMANCE

NO MATTER HOW much consolidation you support in your environment, it's always a good practice to run benchmarking tools or other utilities to report on performance or computing resource levels. This practice puts administrators into a position where they can correlate a numerical relationship to application performance or user experience.

By reviewing the benchmark reports in the wake of an alarm or user complaint, it might be possible to spot a discrepant benchmark counter and troubleshoot potential problems more quickly. When benchmark counters are studied over time, an administrator can observe changes in resource use that might signal the need for upgrades, workload balancing, a new server, or to make other capacity planning decisions.

The need for objective benchmarking or resource reporting increases along with consolidation—more computing resources are used on the server, and more virtual machines are impacted by any changes in the server's state. Fortunately, there are benchmarking and reporting tools available with all three main virtualization platforms. There are also numerous third-party tools available, including Novell Inc.'s PlateSpin Recon and VKernel Capacity Analyzer. ■

information technology for the Town of South Windsor, Conn. "You don't want to get to the point where people are calling with problems."

The information derived from modern management consoles can also help with other important tasks such as workload balancing and capacity planning. Workload balancing analyzes the distribution of VMs and the resources they require and then generates recommendations for organizing or moving workloads that result in a more efficient environment.

Those tools can sometimes "find" capacity that may have been obscured by careless or inefficient workload deployment. Sound capacity planning practices are needed to evaluate resource use over time and ensure that resources are available to meet the future demands.

REEVALUATING CHARGEBACK POLICIES

Another means of limiting VM growth and easing the spiraling demand for server resources is to consider or reevaluate chargeback policies. Chargeback has been a delicate matter for many organizations, and the consolidation of multiple VMs onto fewer servers has only complicated the concept. But an organization that can assign or allocate IT costs back to the business units that request them will be able to make informed financial justifications for computing services and resources. "As soon as [departments] start paying for it, they understand that there's value," Erickson said.

Implementing some form of VM lifecycle management can also help undo and prevent the waste of VM sprawl. Tools like VMware Lifecycle Manager can identify VMs that may not be needed anymore. Removing an obsolete VM will free those computing resources for other VMs.

It's not just a savings of memory and CPU. Removing unnecessary VMs also saves storage space and eliminates those superfluous backup needs that sometimes forestall expenses like upgrades and new server acquisitions—saving considerable capital for the organization.

Finally, it's important to consider the role of new hardware in virtual server consolidation, and regularly funded technology refresh cycles are ideal opportunities to ease over-consolidation by integrating newer and more powerful server hardware into the environment. For example, simply moving from dual-core to quad-core processors can multiply a server's available CPU cycles, yet the server itself will probably cost little more than the older server that you're replacing.

A technology refresh may also be the right time to consider more capable networking technologies such as 10 GbE or Fibre Channel over Ethernet, which can provide high bandwidth and low-latency I/O critical to highly consolidated virtual servers. ■

Get Ready for a Private Cloud

ANALYZE YOUR WHOLE INFRASTRUCTURE BEFORE DECIDING TO MAKE THE MOVE TO A PRIVATE CLOUD ENVIRONMENT. **BY ERIC SIEBERT**

EVERYBODY IS talking about cloud computing as an alternative to traditional computing models, but many are leery of relinquishing control to a public cloud provider. As a result, companies looking to embrace cloud architecture are creating their own private cloud environments. What defines a private cloud infrastructure compared to a traditional data center? Can you just declare your existing data center a private cloud and be done with it?

It's not as simple as that. Private clouds are those that a company manages itself and deploys in its own data center, but they have several characteristics that distinguish them from traditional data centers, including:

DYNAMIC SCALABILITY. The amount of available resources should be able to dynamically increase or decrease based on demand.

HIGH AVAILABILITY. The cloud infrastructure should have as much uptime as possible with minimal unplanned outages.

CHARGEBACK MODEL. Cloud resource consumption should be monitored and measured to calculate user chargeback.

SELF-SERVICE. Users themselves should be able to request and provision cloud resources as needed.

AUTOMATION. Processes within the cloud infrastructure should be as automated as possible to ensure timely deliverables to users.

You should also assess your environment's readiness to become a private cloud. Virtualization is one of the key requirements for building a private cloud infrastructure, so first determine if your virtual environment can meet the demands of a cloud architecture—and what changes you'll have to make if it can't.

Because clouds are dynamic environments, they must be able to handle peak workloads when they occur

—not just meet the demands of average workloads. You will most likely have to add features to your virtual environment to provide greater availability and automation, measure resource usage for chargeback and offer self-service capabilities.

Purchasing a storage system that meets the requirements of your private cloud but that isn't adequate for virtualizing the rest of your infrastructure can be an expensive mistake.

PRIVATE CLOUD INFRASTRUCTURE PRODUCTS

These features are sometimes available from hypervisor vendors, either as companion products or in more expensive editions. For example, some VMware vSphere editions offer Dynamic Resource Scheduler, VMotion, High Availability and Fault Tolerance, which help balance host resources and also provide high or continuous availability for virtual machines. The VMware vCenter Server product line also has LifeCycle Manager, Orchestrator and Chargeback for automation, chargeback and provisioning.

Otherwise, there are many third-party vendors that offer the features you will need in a private cloud infrastructure. VKernel has a Capacity Analyzer product for Microsoft Hyper-V and VMware environments and also a Chargeback product for VMware. Akorri Inc.'s BalancePoint has many features that are critical to cloud environments, including chargeback, capacity planning and reporting. In addition to supporting both Hyper-V and VMware environments, Akorri also supports physical servers—in case parts of your private cloud infrastructure are not virtualized.

Purchasing a storage system that meets the requirements of your private cloud but that isn't adequate for virtualizing the rest of your infrastructure can be an expensive mistake. Likewise, when it comes to networks, you must consider maximum redundancy and sufficient bandwidth capacity to ensure that your cloud is highly available.

PRIVATE CLOUD STORAGE AND NETWORK VIRTUALIZATION

There is more to private clouds than server virtualization. You should also look to deploy storage and network virtualization to ensure maximum interoperability between all the layers of your private cloud.

Storage virtualization consolidates multiple storage devices into a single resource that reduces the back-end complexities and provides a more simplified view. In addition, many storage products integrate directly with server virtualization platforms to provide automation, intelligent multi-

pathing and failover, which are beneficial to cloud architectures.

Network virtualization combines the physical and virtual networks so you can manage and administer them as one network. Products such as the Cisco Nexus 1000v extend the advanced features commonly found in physical switches to virtual environments and have better security, management and quality of service for your private cloud.

WHAT NOT TO VIRTUALIZE IN A PRIVATE CLOUD

A final consideration is whether or not you need to virtualize everything in your private cloud. Cloud applications are typically multi-tiered, and you should consider whether or not you need to virtualize all those tiers to satisfy the requirements of your cloud.

The nature of a cloud may require some of the tiers to dynamically increase resources when needed, but it may not require that from all tiers. Determine which applications will benefit most from virtualization and which ones you may not need to virtualize.

Even though virtualization does have additional benefits unrelated to cloud computing, it may not make sense to virtualize all existing applications—especially those that would be expensive.

On the other side of that argument, however, not virtualizing all your tiers can make chargeback more difficult,

Assessing your needs and requirements early on during the planning process is the key to deciding what you will virtualize in your cloud and what you will not.

because your resources will span both virtual and physical servers. Assessing your needs and requirements early on during the planning process is the key to deciding what you will virtualize in your cloud and what you will not.

Be sure to include your whole infrastructure in your analysis. Servers, storage and networks all have specific requirements and dependencies, and neglecting any one area can negatively affect the other areas. ■

How to Optimize Network Performance to Improve Server Consolidation

LEARN TACTICS TO MONITOR STRESSED HARDWARE RESOURCES THAT CAN AFFECT NETWORK CONNECTIVITY. **BY MIKE LAVERICK**

AS BUSINESSES MOVE up into the next generation of virtualization, consolidation ratios will rise. Adding more virtual machines to servers can easily stress the underlying physical hardware resources and potentially affect network connectivity.

A number of factors contribute to this situation. First, hardware vendors are increasing the amount of supported physical RAM in the host. In turn, hypervisor vendors are updating their software to address this memory.

Additionally, as businesses embrace the virtual desktop approach, they will begin to see a greater density of virtual machines (VMs) per physical box because virtual desktops are generally less memory-hungry than server-based VMs. Memory has typically been the biggest single constraint in achieving increased consolidation ratios. As that hurdle slowly drifts away, concern is turning to the other areas that could be performance bottlenecks—namely the IOPS generated by network traffic.

VIRTUAL MACHINE OPTIMIZATION

The key to optimizing the network layer is making sure your configuration doesn't generate unnecessary contention, which is the general term that many virtualization vendors use to describe a situation where any critical resource becomes oversubscribed. And, there's little to be gained from two network-intensive VMs being placed on the same hypervisor using the same NICs. This would potentially create contention as they compete for an underlying resource. The same applies to such resources as the CPU or memory.

In some cases, however, such a configuration would be appropriate. For example, if there are two VMs

that communicate with each other frequently and transfer large amounts of data, it might be the case that they *should* reside on the same host.

When two VMs speak to each other on the same vSwitch and on the same host, the physical data layer is not touched at all. All network communication happens within the hypervisor.

In this case, the system would not be throttled by the speed of your network but rather by the speed of the physical hosts, CPU and bus. A good example of this scenario would be the communications between a front-end Web server and a back-end database.

So as you can see, the simple adage of "avoid contention" works in most cases, but true network optimization requires that you understand the relationships *between* your VMs as well. Most virtualization management software allows you to express these relationships in the form of anti-affinity and affinity rules (see **Figure 1**).

You can create rules that say, for example, that a database and Web server must reside on the same host but that each of your Microsoft Active Directory VMs must never reside on the same host. This allows you to

When two VMs speak to each other on the same vSwitch and on the same host, the physical data layer is not touched at all.

maintain the application scalability and availability that most vendors achieve by scaling out—as opposed to scaling up—their technologies.

The next step in optimizing your VMs for network connectivity is to

Figure 1

Anti-affinity and Affinity

This shows anti-affinity and affinity rules in VMware's Distributed Resource Manager configuration.

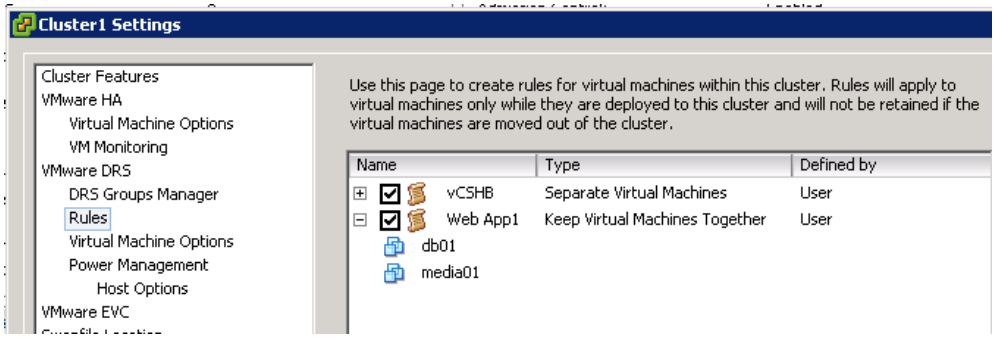

ensure that the guest operating system that resides within the VM has been correctly configured. Most virtualization vendors allow you to configure the VM with an optimized network driver that sits inside the VM.

For example, VMware has the vmxnet2 and vmxnet3 driver sets, whereas Microsoft Hyper-V has the concept of a synthetic network device (see **Figure 2**). For these to work, you have to install the integration software that ships with the virtualization vendor.

Enhanced network drivers reduce the number of CPU cycles required on the physical machine to move network packets from the virtual world to the physical world. Without them, you would probably see reduced network performance and an increase in CPU use on the physical server.

These drivers offer a discrete amount of paravirtualization to the guest operating system making it more VM-aware than it would be without them. Paravirtualization is a general attempt to make any part of the system more tuned to a virtualized environment.

Additionally, these drivers are often a requirement to leverage more advanced network enhancements. So, for example, if you want to offer larger maximum transmission unit (MTU) sizes to the guest operating system, the installation of their enhanced network driver type will often be a requirement.

Remember: On its own, virtualization doesn't necessarily make your guest operating system run any quicker. If a VM outperforms a physical server, it will be largely because

Figure 2

Microsoft's Hyper-V Technology

This shows the ability to add a synthetic network adapter, as opposed to the Legacy Network Adapter.

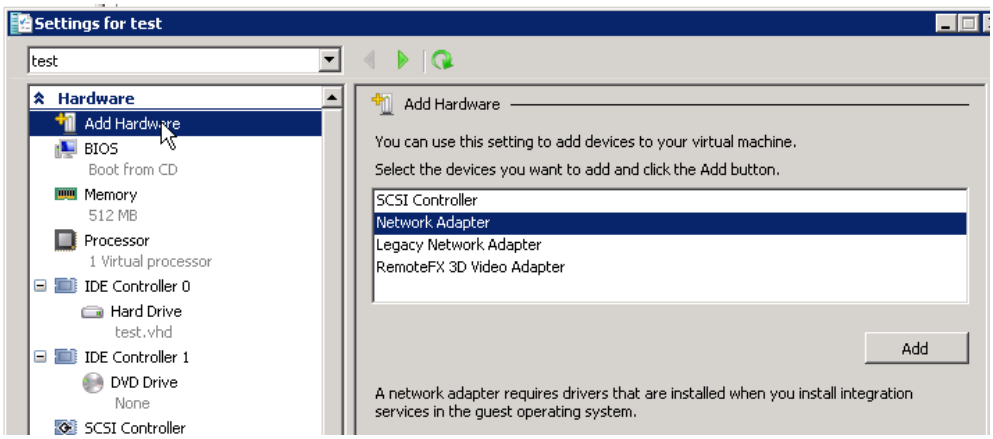

of external forces, such as having a faster physical server, more resources to the VM than the operating system previously endured in the physical world or a storage layer that has been upgraded.

The optimization you learned in the physical world with Windows or Linux still applies to the virtual world. So if you learned how to modify the Windows registry to tweak the performance of TCP/IP, those optimizations are still likely to pay dividends in the virtual world. The same applies to other tweaks such as changing the firewall settings and disabling unwanted services that might generate a network load that is not needed for your application to function.

HYPERVISOR OPTIMIZATION

Once you are happy that the right VMs are running on the right hosts and that the correct guest operating system customization and optimization has been carried out, the next area you should focus on is the underlying vSwitch configuration of the hypervisor.

A couple of key configuration settings exist in any hypervisor that can be used to guarantee optimized network performance. Once again it's about avoiding contention and oversubscription to the underlying resource.

First, it's important to create NIC teams for dedicated network traffic types. There are generally six traffic types in most hypervisors:

- Management
- IP-based storage (optional)
- Live migration
- High-availability heartbeat
- Fault-tolerance (specific to VMware)
- Virtual machines

In the ideal world, each traffic type should be serviced on a dedicated physical NIC. If redundancy is deemed a requirement for the traffic type in question, then more than one NIC could easily be bonded together to create a NIC team. That is often a best practice or recommendation from the virtualization vendor.

However, there is room for maneuver if your hardware doesn't have a healthy number of NICs or the platform prohibits the configuration—for example, on older blades that support only two NICs per blade.

Some administrators place their live-migrate traffic on their management network because historically this has been an underused chunk of bandwidth. Despite the use of VLAN tagging, you may want to separate this traffic for security compliance as well.

The key here is to make sure the underlying network traffic that allows the hypervisor to function—management, IP-based storage and live-migrate traffic—does not affect the VMs. So storage and live-migrate traffic, which can be particularly bandwidth-intensive, should be on separate physical NICs.

Another area of vSwitch optimization can be achieved by enabling

(Continues on page 17)

SETTING MAXIMUM MTU

THE MAXIMUM TRANSMISSION unit (MTU) setting within the hypervisor is configurable in many different ways and locations. For instance, if you are using Standard vSwitches in VMware, you would use the command line tool esxcfg-vswitch to set the MTU value. And if you are using VMware Distributed vSwitches (see Figure 3), you find it as a setting in the dvSwitch itself.

FIGURE 3

MTU IN DVSWITCH

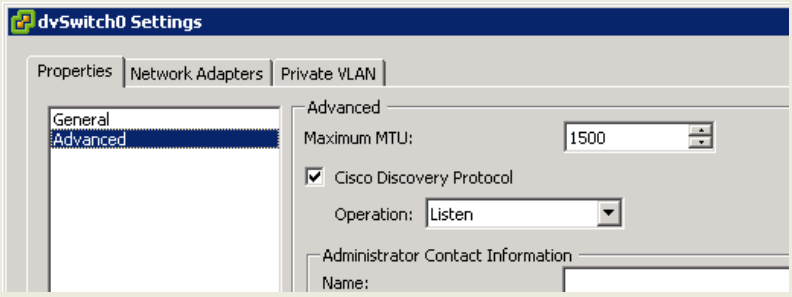

In Microsoft's Hyper-V, the configuration is held on the properties of the local area connection or network team, and it is controlled essentially by the vendor of your network cards (see Figure 4).■

FIGURE 4

MICROSOFT HYPER-V CONFIGURATION

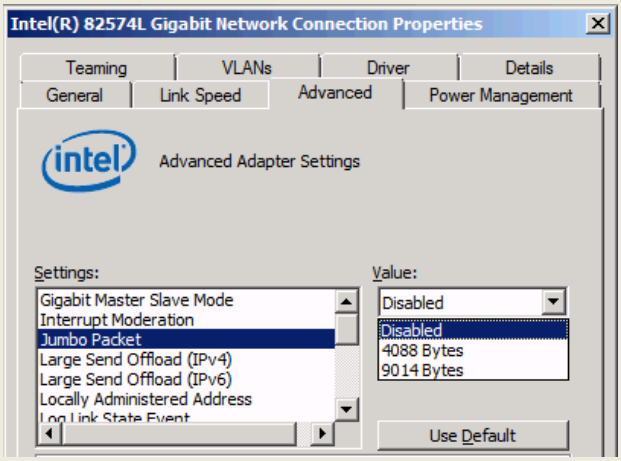

EDITOR'S LETTER

ACHIEVING THE RIGHT LEVEL OF SERVER CONSOLIDATION

GET READY FOR A PRIVATE CLOUD

HOW TO OPTIMIZE NETWORK PERFORMANCE TO IMPROVE SERVER CONSOLIDATION

(Continued from page 15)

jumbo frames. This packet type has a much larger MTU size of anywhere up to a maximum of 9,000 on most modern networks. Increasing the MTU value from the standard value of 1,500 to a larger value improves performance across the board for all traffic types whether they are virtual or physical.

This happens because fewer frames are used to send the same volume of data. Because fewer frames are being sent, overhead is reduced. This results in fewer TCP/IP acknowledgement packets.

Implementing a larger MTU value is not a task to be undertaken lightly. For the benefits to be tangible, every component—the physical switch, the vSwitch and the VM—will need to be correctly configured. In cases with secure communication using protocols like SSL, a poor configuration of the MTU value could stop communication altogether.

Other optimization areas to investigate are any settings you can enable to improve the algorithms associated with the use of teamed network cards. The most common use of teamed network cards is to offer redundancy to the network, especially if NICs are split across different Layer-2 switches.

But many hypervisors adopt teaming policies by default that are designed more for compatibility than for better optimization of the network layer. Most people assume that merely by teaming the NICs together they will double or triple their available bandwidth, but for many hypervisors this is not the case.

A lot depends on the way the hypervisor has been developed, and the default settings may indeed need modifying. For example, the default in VMware is to use a policy called Originating Port ID. This carries out a round-robin on the network cards in an effort to distribute traffic load.

Although this offers excellent compatibility in many different network environments, it is not the most optimized policy to adopt. Switching to a load-balancing policy that uses IP data is usually the best approach.

MONITORING VM TRAFFIC

Moving away from optimization, it's worth dwelling on the methods available for monitoring your network I/O and investigating the possible causes of slow networking. Just how do the symptoms show themselves in virtualization software.

Although the approaches described thus far are those that can be made to optimize network configuration, the reality is that poor network performance can often be attributed to sources outside virtualization.

If an application owner reports slow networking then it is worth making sure that what they are experiencing is not caused by a bottleneck in the wide-area network. There may be an outage or routing problem that has yet to be reported or discovered.

Another area to check is in the IP configuration. Simple tools like ping,

(Continues on page 19)

OPTIMIZING PHYSICAL SWITCHES

CARE MUST BE taken to confirm that the physical switches support the correct protocols and standards. For example, in the case of VMware's IP Hash policy, the physical switch needs to be enabled for the IEEE 802.3 Link Aggregation feature. In the case of VMware, the configuration is carried out on the properties of the vSwitch.

Microsoft Hyper-V (see **Figure 5**) is the function of your vendor's network settings. The interesting thing here is that you may have more vendor-specific proprietary options from the NIC vendor. For example, in the case of Intel network cards, there is a specific option for Virtual Machine Load Balancing.

FIGURE 5

MICROSOFT HYPER-V PHYSICAL SWITCHES

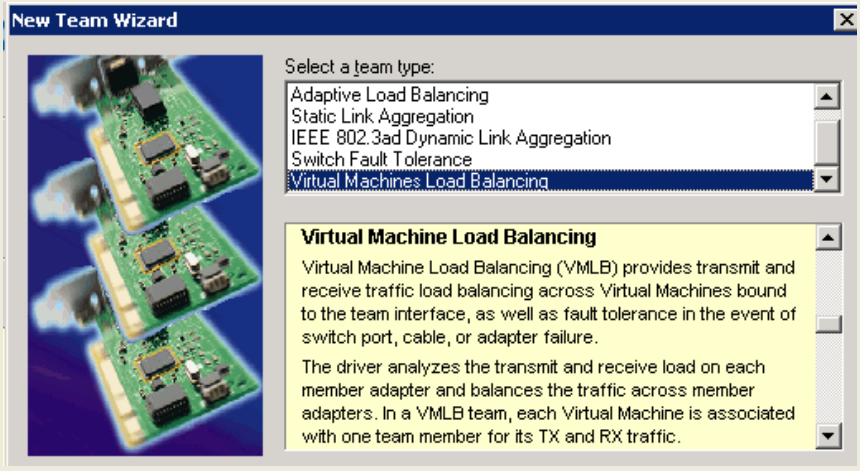

For the new release of VMware vSphere 4.1 and Citrix XenServer 5.0, you will find that modern hypervisors ship with some kind of Network I/O settings that allow you to control bandwidth as it leaves and enters the hypervisor. These controls are new, and it remains unclear how significant this development will be.

Some customers may prefer to handle the IOPS controls by using methods available outside of the physical server—and embedded in the new 10/20 Gps hardware that allows the administrator to control bandwidth allocations independent of hypervisor version or vendor. Remember the configuration changes don't necessarily double or triple your bandwidth available, but they should increase the overall I/O capabilities of the hypervisor in question. ■

EDITOR'S LETTER

ACHIEVING THE RIGHT LEVEL OF SERVER CONSOLIDATION

GET READY FOR A PRIVATE CLOUD

HOW TO OPTIMIZE NETWORK PERFORMANCE TO IMPROVE SERVER CONSOLIDATION

(Continued from page 17)

pathping, tracert and nslookup can still be useful in diagnosing network problems. For example one of the most common problems is poor or incorrect DNS configuration.

Another place to check is the configuration of the application within the VM. If there a setting or an option that could significantly degrade network performance—perhaps the application polls the network for availability of external network components—then this can lead to unnecessary traffic.

Once you have excluded these as potential problems, it's worth confirming if the components that are optimized have been configured correctly. Next, check whether the network problems affect just the VM in question or all the VMs on the same host, which is a good way of figuring out whether the problem is specific to the application owner's VM or whether it represents a systemwide problem. Most hypervisor vendors offer network tools that allow you to monitor traffic coming in and out of the VM.

In the case of VMware, you can use a utility called *esxtop* to see network statistics (see **Figure 6**) and troubleshoot performance problems. Hitting *n* on the keyboard toggles esxtop to a network mode, and *f* on the keyboard allows the administrator to add additional fields.

Most hypervisor vendors offer network tools that allow you to monitor traffic coming in and out of the VM.

Figure 6

Checking Network Stats Using the Esxtop Utility

VMware's esxtop utility allows administrators to view network statistics.

```
root@esx1:~
7:15:31pm up 51 days  5:13, 133 worlds; CPU load average: 0.02, 0.02, 0.03
```

PORT-ID	USED-BY	TEAM-PNIC	DNAME	PKTTX/s	MbTX/s	PKTRX/s	MbRX/s	%DRPTX	%DRPRX
16777217	Management	n/a	vSwitch0	0.00	0.00	0.00	0.00	0.00	0.00
16777218	vmnic0	-	vSwitch0	8.69	0.01	9.32	0.01	0.00	0.00
16777219	4096:vswif0	vmnic0	vSwitch0	5.93	0.01	6.15	0.01	0.00	0.00
16777432	18540:vc4nyc	vmnic0	vSwitch0	2.54	0.00	4.87	0.00	0.00	0.00
16777435	20226:scvmm	vmnic0	vSwitch0	0.21	0.00	1.70	0.00	0.00	0.00
33554433	Management	n/a	vSwitch1	0.00	0.00	0.00	0.00	0.00	0.00
50331649	Management	n/a	DvsPortset-0	0.00	0.00	0.00	0.00	0.00	0.00
50331650	vmnic1	-	DvsPortset-0	2.12	0.00	4.24	0.01	0.00	0.00
50331651	vmnic2	-	DvsPortset-0	1.48	0.01	3.18	0.00	0.00	0.00
50331652	vmnic3	-	DvsPortset-0	2.97	0.01	4.45	0.00	0.00	0.00
50331653	vmk0	vmnic3	DvsPortset-0	2.97	0.01	4.03	0.00	0.00	0.00
50331654	vmk1	vmnic1	DvsPortset-0	0.00	0.00	1.70	0.00	0.00	0.00
50331655	4096:vswif1	vmnic1	DvsPortset-0	2.12	0.00	4.03	0.01	0.00	0.00
50331656	vmk2	vmnic2	DvsPortset-0	0.00	0.00	1.70	0.00	0.00	0.00
50331675	18540:vc4nyc.eth1	vmnic2	DvsPortset-0	1.48	0.01	2.97	0.00	0.00	0.00

These utilities allow you to see how much bandwidth is actually being used by the VM and if the physical system is seeing a significant number of dropped network packets. They also show the rate of transmit and receive on the system.

When packets are sent out but are not receiving acknowledgement, it can indicate a problem with NIC Teaming algorithms referred to as the Reverse NIC Team problem. In this scenario, advanced NIC Teaming has been enabled, and although packets leave the physical host via one network layer, they arrive back at the host via the wrong physical switch and to the wrong NIC.

Serious problems such as this may need wider investigation. In some cases, it can result in the abandonment of particular NIC teaming policy that has been deemed as not reliable enough for the wider network.

Plenty can be done to improve and monitor network performance for VMs as your consolidation ratios grow. The key to the best optimization is following your virtualization vendor's best practices while modifying them to suit the unique traffic characteristics of your network. The most critical part is to understand the relationships between your virtual machines and the wider physical world. ■

The key to the best optimization is following your virtualization vendor's best practices while modifying them to suit the unique traffic characteristics of your network.

ABOUT THE AUTHOR

Stephen J. Bigelow, a senior technology writer in the Data Center and Virtualization Media Group at TechTarget, has more than 20 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.

Eric Siebert is a 25-year IT veteran whose primary focus is VMware virtualization and Windows Server administration. He maintains vsphere-land.com, a VMware information site. Siebert was one of the 300 VMware vExperts for 2009, is the author of the book VI3 Implementation and Administration and is a frequent TechTarget contributor.

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

Matt Stansberry
Executive Editor
mstansberry@techtarget.com

Christine Casatelli
Editor
ccasatelli@techtarget.com

Marty Moore
Copy Editor
mmoore@techtarget.com

Linda Koury
Art Director of Digital Content
lkoury@techtarget.com

Jonathan Brown
Publisher
jebrown@techtarget.com

Peter Larkin
Associate Publisher
plarkin@techtarget.com

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

SPONSOR/LINKS

→ **Brocade Data Center Solutions**

About Brocade: Brocade® (Nasdaq: BRCD) networking solutions help the world's leading organizations transition smoothly to a virtualized world where applications and information reside anywhere. Industry-leading offerings for Ethernet, storage, and converged networking environments help these organizations achieve their key business objectives through unmatched simplicity, non-stop networking, optimized applications, and investment protection. To ensure a complete business solution, Brocade partners with world-class IT companies and provides a full range of services and support offerings. (www.brocade.com)