

Chapter 12

Performance Optimization

VMware ESX is and has been one of the highest performing server virtualization platforms on the market for some time now. So what that means is VMware typically does a really good job of providing high performance out of the box; however, that doesn't mean that performance tuning and tweaks aren't necessary. In fact, there are still plenty of opportunities for you to get your hands dirty by doing a little bit of work to help further optimize the performance of your virtual infrastructure's environment.

You might be one of those users of VMware ESX that is happy with the performance of your virtualization platform or the performance of your individual or collective virtual machines. But what happens tomorrow if something in the environment changes and it negatively affects your performance? What happens if you need to achieve a 12:1 server consolidation ratio on an existing VMware ESX that today only hosts eight virtual machines? At that point, it may become extremely important to fine tune the environment in order to squeeze out every ounce of performance that your environment can offer.

While this chapter won't go into any great detail on any one single area of performance, it will cover a number of different options that apply to most implementations or discuss other practices that might cause performance problems.

In order to achieve performance optimization, we will focus on three main target areas where improvements can be made. The first target area covered is the VMware ESX host server machine. The configuration of the host server should be carefully examined at both the physical hardware layer and the virtualization platform layer. Decisions made here will ultimately affect the overall performance

of every virtual machine contained on the host server. The second target area is the virtual machine layer. Configuration options are one such factor that can affect the performance of the virtual machines. And finally, the third target area that can be optimized is the configuration of the guest operating systems themselves that operates within each virtual machine. Tuning the guest operating system is often times overlooked simply because it operates within a virtual machine. If you normally optimize the operating system installed on a physical server, why wouldn't you try to perform similar tuning techniques to your virtual machines?

Configuration of the Host Server

Virtualizing your infrastructure or even a small number of your physical servers can offer enormous benefits to your organization. But it is important to realize that virtualization can also negatively affect the performance of your server, operating system and/or the applications if not properly setup, configured or planned. Therefore, you should understand the trade-offs that occur at the hardware layer with virtualization and to make sure to properly configure the host server on a component-by-component basis. In this section, we will focus our optimization efforts on the physical host server including the virtualization platform chosen, its processor, memory, storage and its network.

TIP

Make sure you check all the installed hardware items in your host server against the most recent supported hardware compatibility list (HCL) from VMware for whatever version of ESX you are installing. Adding unsupported hardware components can cause performance problems, or worse.

Upgrade your Version of VMware ESX

If your environment is still operating with some earlier version of VMware ESX Server 2.x, now might be the time to start thinking about upgrading, or better yet migrating it to the latest version. In case your VMware sales representative hasn't told you yet, there are significant performance improvements that can be made by simply running the latest version of VMware ESX. By upgrading your environment to VI3 3.5, you not only gain out-of-the-box performance improvements, but you also get significant improvements in scalability, usability, hardware compatibility and a host of new features and benefits described throughout this book.

VMware has optimized several of the component areas in VMware VI3. Doing so allows organizations to see an increase in overall performance, and it also offers

the ability to now virtualize high demanding and intensive workloads that might have been previously difficult to do.

With the latest version of VMware ESX, virtual machines can enjoy an increase in the number of virtual CPUs that can be assigned to them. Early versions of VMware ESX restricted virtual machines to a single-processor environment. Since then, VMware has created Virtual SMP technology to give virtual machines the ability to operate with multiple processors. With the latest version of VI3, virtual machines can now enjoy the freedom of using up to four virtual processors. ESX 3.5 now makes use of 32 logical processors on a single host, with experimental support of as many as 64 logical processors. And it can also take advantage of ACPI power saving mode to better handle the wasted cycles consumed by idle virtual machines. Virtual memory limitations have also been increased. Virtual machines can now enable Physical Address Extensions (PAE) to access up to 64GB of memory where the early limitations had been set at 3.6GB of memory for quite a number of years. The memory limit for VMware ESX host servers has also been increased to 256GB. ESX 3.5 has also improved its Non-uniform Memory Access (NUMA) scheduling algorithms which should greatly benefit the virtual machines running on NUMA platforms.

Virtual networking has also gone through a significant overhaul and refresh with the latest version of VI3. Improvements have been made with both the virtual network adapter and the virtual switch. Virtual adapters are capable of GigE speeds, the number of virtual switches has increased as did the number of ports available, and an additional load balancing method called port ID is also now offered. A new version of the VMXNET virtual device called Enhanced VMXNET is available, and it includes several new networking I/O enhancements such as support for TCP/IP Segmentation Offload (TSO) and jumbo frames. ESX 3.5 hosts now fully support 10 GigE NICs which offer huge performance improvements compared to traditional 100 Mb Ethernet cards. Experimental support for Intel I/O Acceleration Technology (I/OATv1) has also been added. If your system's chipset has this feature, you should be able to see improvements in your networking performance.

A number of architectural improvements were made to the ESX 3.5 storage subsystem. VI3 now supports Infiniband-based Host Channel Adapters (HCA). Instead of using NICs and Fibre Channel adapters, you can install Infiniband HCA adapters which would look like Fibre Channel adapters and Ethernet NICs to the virtual machines. VMFS-3 is a new version of the VMFS family of clustered file systems that offers enhanced performance and scalability over previous versions. Virtual machines are less dependent upon the Service Console and their user-level virtualization components can now run on any available processor core which gives the ability to scale up to 50% more virtual machines per host server.

Other features available in VI3 that enhance performance and scalability include VMware Distributed Resource Scheduler (DRS), Resource Pools, High Availability (HA), VMotion, Storage VMotion and VMware Consolidated Backup

(VCB) to name but a few. These and other VI3 enhancements are discussed in more detail throughout this book.

Host Server Processor Performance

The physical hardware of your host server greatly affects the performance of all of the virtual machines that reside upon it. And the processor is considered one of the most vital resources contained within your host server.

As cliché as this may sound, a general rule of thumb to follow is the more physical processors in your server—the better. Likewise, most new servers are now shipping with multiple cores to increase the processing power found in the box. Packing in multiple core processors can help achieve higher densities per server while at the same time increasing the performance of existing virtual machines because of reduced CPU contention. And the newer versions of VMware ESX have official support for multi-core processors.

You should also be aware that although VMware ESX is compatible with a broad range of server hardware, including older server equipment, it doesn't necessarily pay to virtualize those aging servers. Without getting into a cost game here, let's just stay focused on performance. Yesterday's servers have single-core processors that are significantly slower in GHz speed and have slower or less cache. You can greatly improve the performance of your guest environments by simply upgrading your servers and improving the CPU used. If the performance of the virtual environment calls for it, use the highest performing processors you can afford.

In addition to being faster, these newer and more modern processors are also more efficient. Be aware, virtualization environments are typically making more efficient use of the server by fully utilizing the processor, in many cases, reaching and sustaining between 60-80% utilization. That's fine, except if you are dealing with older processor technology that is less efficient with heat and power consumption. Many processors automatically step down their performance when a certain thermal threshold is reached. Remember, these older processors were designed with the notion that they would probably only reach 80% utilization for a small period of time during a spike in usage; once they reach the thermal threshold, they step down and downgrade their performance, which will negatively affect every virtual machine hosted on that server.

It is also important to optimize the ratio between the number of active virtual machines hosted on a server and the number of physical processors or cores available on the host. Usually, the more cores the higher the density. But watch out, operating too many virtual machines per processor core can adversely affect the overall performance of the server and its virtual machines. There is no magic formula that gets the ratio mix right every time and for each scenario. For years, virtualization administrators have been trying to come up with this formula but to no avail. Instead, VMware created DRS and resource pools to try and take much of

the guess work out of the equation. While it may not be perfect, it certainly helps keep the performance of the environment at a certain level and without the need for constant human attention.

TIP

Avoid running programs in the service console that consume large amounts of the processor. By allowing a third-party service console application to excessively consume CPU resources, the virtual machines and the host server performance will be adversely affected. Use `esxtop` to monitor service console CPU usage.

Host Server Memory Performance

Much like the processor in a VMware ESX host server, the memory of the host server is also considered a significant bottleneck. And just like the processor scenario, it is important to add as much memory to your VMware ESX host server as possible. Having a sufficient amount of memory for all of your virtual machines is important to achieving good performance. A sufficient amount of memory is roughly equivalent to the amount of memory you would have assigned to each virtual machine if they were physical. As an example, to effectively run a Windows XP Professional virtual machine you might allocate 512MB of memory to it.

It is important to note here that system memory has quickly become one of the most expensive components found in today's modern server; so you need to make sure that you properly match up the amount of memory in the system to the virtual machine density that can be achieved with the amount of processing power available. In other words, if you have enough processor resources in your host server to support ten virtual machines, don't overpopulate your host server with expensive additional memory beyond what that density is capable of consuming if you cannot achieve a higher consolidation ratio due to processor limitations.

VMware ESX server hosts are required to have a minimum of 1GB of memory. In reality, 4GB is probably more of a normal minimum requirement. Today's modern virtualization host servers are more typical with 8 or 16GB of memory, and larger environments are seeing host servers with 32 or even 64GB. ESX server hosts require more RAM than a typical application server. It must be equipped with enough memory to concurrently power and operate a number of virtual machines plus its service console and any third-party tools that may be installed.

When raw virtual machine performance is the most important metric, avoid over committing memory. Make sure the host server has more memory than the total amount of memory needed by ESX plus the sum of the memory sizes assigned to each virtual machine. Although ESX can handle over committing

memory, that doesn't mean you should do it. Doing so creates a swap file per virtual machine that is located on disk in the same location where the virtual machine is located. Make sure you have enough disk space available for the swap. And remember, disk is much slower than RAM, so for performance reasons, it is best to avoid swapping to disk.

TIP

If swapping cannot be avoided, make sure you at least assign the swap file to a location on a high speed storage system.

Host Server Storage Performance

When searching for a performance bottleneck in your virtualization system, the CPU and the RAM are usually the first two suspects. But by increasing the density and usage of a physical machine with numerous virtual machines, the amount of disk I/O activity greatly increases across that server's disk subsystem. And certain workloads are also very sensitive to the latency of I/O operations. Because of that, the host server's storage system becomes the third bottleneck that needs to be addressed in a virtual world.

Storage performance issues are usually the result of a configuration problem. With ESX, you have the choice of using either local or remote storage. Either selection comes with its own set of configuration choices to be made that can affect performance. And there are also a number of other dependencies such as workload, RAID level, cache size, stripe size, drive speed and more.

Without sounding too obvious, using the fastest and highest performing disks and controllers will greatly improve the performance of your virtual environment. If you have decided to go the route of local or direct-attached disk storage in your environment, you should go with 15K RPM disks to improve the I/O performance across all of your virtual machines on that host server.

In addition to the speed of the disk drive, you should also consider the type of disk drive. While SATA drives are now possible to use with the latest version of ESX, you are better off for performance reasons spending the money and selecting either Ultra320 SCSI drives, or even better if your system supports it, SAS disks.

Many disk controllers can support multiple channels on the card, and by splitting your disks across multiple channels, you can achieve a performance improvement. For example, if you have six SCSI disks and a two-channel controller, your hardware might allow you to place three disks on each channel and configure them in a six-disk RAID-5 array. This would allow you to effectively split your I/O across two channels.

You might also be able to install multiple disk controllers and additional disks within your host server. This would allow you to split up your file system and strategically place your virtual machines according to I/O needs. In other words, if you have ten virtual machines on a host server and two of those virtual machines are running disk I/O intensive applications, you have a choice on how to configure your environment for the best possible performance. You can either locate the eight low intensive virtual machines on one controller with the two high intensive virtual machines on the other controller, or you can configure five virtual machines on each controller where only one disk I/O intensive virtual machine is allocated to each file system.

If your organization is lucky enough to afford one, it is recommended to use SAN storage rather than using locally attached disks so that you can achieve much better performance. SAN storage technology allows VMware ESX to shine with its added capabilities. Using a SAN will offload I/O operations from the ESX host server which leaves more resources available to the virtual machines. To further optimize performance, spread the I/O loads across multiple 4Gbps Fibre Channel SAN HBAs where possible. And make sure that heavily used virtual machines aren't accessing the same VMFS volume at the same time. Spread them across multiple VMFS volumes so that disk performance won't suffer.

TIP

Make sure I/O traffic isn't queuing up in the VMkernel. To verify, monitor the number of queued commands with `esxtop` over a period of time.

VI3 has added new options for you to take advantage of in your remote storage. In addition to Fibre Channel SAN, you can now use iSCSI and NFS to take advantage of cheaper storage solutions using your existing IP networking technology. For iSCSI and NFS, it is important to make sure that your configuration and design does not result in an oversubscribed Ethernet link. The TCP session control will ensure that there is recovery from packet loss, but frequently recovering from dropped network packets will result in huge performance problems. Virtual machines with intensive I/O applications should not share Ethernet links to a storage device. And they will perform even better if they have multiple connection paths to its storage.

When your virtual machines share access to the ESX host server's I/O subsystem, use the I/O share allocation for each virtual machine to adjust the amount of I/O resources that the virtual machine is given. For virtual machines running applications that aren't very I/O intensive, you can set their resource shares to something low, like 500. And for the more resource intensive virtual machines that require more priority to I/O resources, set their shares to something higher, like 2000.

Host Server Network Performance

Network utilization can also present bottleneck issues in your environment much like CPU, memory and storage. But in most virtualized environments, you will find that the network is probably the least likely culprit of performance problems. However with that said, the host server still needs to be supplied with the appropriate amount of network bandwidth and network resources so that the virtual machines don't add any significant amount of network latency into the equation.

If you haven't already done so, upgrade your network environment with Gigabit Ethernet. With 10GigE waiting to take over, Gigabit Ethernet network adapters and switches should be affordable. Using Gigabit network adapters allows more virtual machines to share each physical network adapter and it greatly improves the amount of network bandwidth made available to network intensive virtual machines.

When configuring your physical network adapters, the speed and duplex settings on each card must match the speed and duplex setting used on the switch port to which it is connected. The VMkernel network device drivers start with a default speed and duplex setting of auto-negotiate. The auto-negotiate setting is fine and should work correctly with network switches that are set to auto-negotiate. This is the default and preferred setting for gigabit connections. When using 100Mbit Fast Ethernet adapters, you should set the network adapter and the switch port speed and duplex settings to match at 100/Full. If you have conflicting settings between the network adapter and the switch, it can not only cause a performance problem but in some cases a connectivity issue as well.

TIP

If you encounter bandwidth issues, check to make sure the NIC auto-negotiated properly. If not, change the speed and duplex settings manually by hard coding them to match. Do so at the switch or the VMkernel networking device using the VI Client.

You can also increase the available network bandwidth and increase network fault tolerance by teaming multiple Gigabit network adapters into a bond. This will also simplify the number of virtual switches being mapped to physical network adapters as well. You can also use separate physical network adapter and vSwitches to avoid network contention between the service console, the VMkernel and the virtual machines. It can also be used to segment network I/O intensive virtual machines from one another.

You might want to leverage the new VMware ESX 3.5 networking enhancements that have been integrated into its networking code. Jumbo Frames are now supported. Supporting Ethernet frames up to 9000 bytes (as opposed to standard Ethernet frames supporting a Maximum Transfer Unit of 1500 bytes) allows guest

operating systems using Jumbo Frames to require fewer packets to transfer large amounts of data. Not only can they achieve a higher throughput, but they also use less CPU than a standard Ethernet frame.

Another new feature in 3.5 is the support for TCP Segmentation Offload (TSO). TSO is widely used and supported by today's network cards. It allows for the expensive task of segmenting large TCP packets of up to 64KB to be offloaded from the CPU to the NIC hardware. ESX 3.5 utilizes this concept to provide virtual NICs with TSO support even when the underlying hardware doesn't have the special TSO capabilities. Because the guest operating system can now send packets that are larger than the MTU to the ESX server, processing overheads on the transmit path are reduced. TSO improves performance for TCP data coming from a virtual machine and for network traffic that is sent out from the server.

Configuration of the Virtual Machine

When you create virtual machines for your environment, accepting default responses aren't necessarily the right choice. Will it work? Sure. Will it be optimized for performance? Perhaps not. There are things that you should be aware of and things that you can do with the configuration of your virtual machines to squeeze out additional performance.

Remove Unneeded Virtual Hardware

One of the nice things about working with a virtual machine is the ease with which you can add or remove hardware components. You don't need any tools, and you don't need to open the hood of the server. A quick and easy way to gain a small amount of performance back in your guest environment is to disconnect or remove any unused or unnecessary devices. If your virtual machine doesn't need a CD-ROM drive, Floppy drive, Network Adapter or COM and LPT ports, get rid of them. When you do need them, it is extremely easy to either enable them or add them back. This can free up IRQ resources and it eliminates IRQ sharing conflicts.

Power off Idle Virtual Machines

Because virtual machines can so quickly and easily be created, a common problem called virtual machine sprawl has erupted. If it isn't managed properly with processes and controls, unmanaged and forgotten virtual machines can spring up and consume datacenter resources. It is important to identify virtual machines that are no longer necessary or no longer being used, because these virtual machines should be powered off or suspended to keep them from wasting valuable resources. Even if

a virtual machine is idle, it still uses memory and CPU that other virtual machines could be leveraging for their own performance.

Virtual Machine Processors

When a virtual machine's performance seems sluggish, many people automatically assume that it just needs more processing power. And since VMware has upped the ante on the number of virtual processors that can be assigned to any given virtual machine, why not just add more virtual CPUs to the virtual machine to increase performance? Unfortunately, that doesn't always work. In fact, in many cases, doing so may adversely affect the entire environment and further degrade performance.

Adding additional processors to a virtual machine by using VMware's Virtual SMP is not always going to solve the performance problem. Why? First off, your performance problem may not even be caused by a lack of processing power. Your bottleneck could be memory, disk I/O or network I/O. If the guest operating system or applications aren't being starved for CPU resources, adding more CPUs may only negatively affect the rest of the environment on that host server.

And second, not all operating systems and applications are able to take advantage of having multiple CPUs. If your application is single threaded and will not benefit from the additional VCPUs, you shouldn't use Virtual SMP. In this case, adding Virtual SMP may consume valuable physical processor resources without actually offering the virtual machine any added performance benefits, and as a result, take away resources from other virtual machines on the same physical host server. Virtual SMP should therefore be used sparingly and only after determining whether or not an application is indeed multi-threaded. Running a virtual machine with a single virtual CPU, in many cases, will outperform the same virtual machine with Virtual SMP turned on. This determination should be treated on a case by case scenario, and you should test each virtual machine environment to see if Virtual SMP will help or not.

Make sure to configure your virtual machines with the correct HAL or kernel. Single processor virtual machines should be configured with a UP HAL or kernel while a multi-processor virtual machine should be configured with an SMP HAL or kernel. SMP versions are required in order to fully utilize multiprocessor systems but can also be used on single-processor systems. Because of their synchronization code, SMP operating systems used on single-processor systems will operate slower than a UP operating system.

Another way to affect virtual machine performance is by altering the virtual machine's minimum and maximum CPU resource allocation percentages. If you want to avoid CPU starvation for a virtual machine, set its minimum percentage to something other than zero. Conversely, to keep low priority virtual machines from consuming too many CPU cycles, you can set its maximum percentage to some-

thing lower, like 50%, thereby effectively throttling down that virtual machine and allowing other more pertinent virtual machines to make use of those valuable CPU cycles.

In addition to setting thresholds with MIN and MAX settings, you also have the ability to control which physical processor or processors that each virtual machine can use. This control is called processor affinity. The default setting is to use no affinity, and this is usually the best choice for most situations. You should really only set a virtual machine's CPU affinity when absolutely necessary. As an example, if you have a very resource intensive virtual machine running on a host server, you might want to set its CPU affinity to isolate that virtual machine and to protect its performance. Doing so will also protect the performance of all other virtual machines running on that same host server by changing each of their affinity settings to a different processor from that of the resource intensive virtual machine.

Virtual Machine Memory

If you are trying to optimize the performance of your virtual machine and find that the performance degradation isn't being caused by the processor, examine the virtual machine's use of memory. If the guest operating system is paging or swapping memory too much, performance will suffer since writing to disk is much slower than writing to memory. To identify if your virtual machine is paging, use `vmstat` from the command line on Linux or the Performance tool found under Administrative Tools on Windows to check the value for pages/second. If the number is high, such as 1000 pages/second, increase the amount of memory assigned to the virtual machine to eliminate excessive paging.

Only allocate as much memory to the virtual machine that is needed to allow enough memory to hold the working set of applications being asked to operate in the guest environment. Some amount of testing is needed here since giving a virtual machine too much memory will reach diminishing returns and simply be wasteful. The wasted memory could have been used as an additional resource elsewhere on the host server.

If possible, configure Linux virtual machines to use less than 896MB of memory. Linux uses different techniques to map memory in the kernel if the amount of physical memory is greater than 896MB. Every physical page of memory up to 896MB is mapped directly into the kernel space. This memory section is faster and more efficient and can keep your guests running optimally. Any amount of memory over that is no longer permanently mapped but is instead temporarily mapped. These techniques add additional overhead on the virtual machine monitor and can lower performance.

NOTE

It usually depends on whether you are talking to the VMware ESX administrator or the virtual machine end-user as to what the “proper” configuration amount of memory is for each virtual machine. The amount of memory a virtual machine actually needs and the amount of memory the end-user wants normally doesn’t match up. Finding a proper balance is the key to a successful infrastructure implementation.

Virtual Machine Networking

If you are trying to create a networking connection between virtual machines that live on the same host server, connect the virtual machines to the same vSwitch. While it isn’t mandatory that these virtual machines be connected to the same vSwitch, doing so will keep the networking traffic from going out and across the wire. At the same time, it also keeps CPU and network overhead down and increases the performance of the network communications that take place between those virtual machines.

When you create a new virtual machine, the default network adapter that is emulated on a 32-bit guest is the AMD PCnet32 device and it is configured with VMware’s vlance driver. For performance reasons, in a GigE environment, you should change the emulated network adapter to use either VMware’s vmxnet driver or e1000. The vmxnet driver passes through network traffic from the virtual machine to the physical network adapter with minimal overhead. And the latest version of VMware ESX provides a new version of the vmxnet virtual device called Enhanced VMXNET. It includes several new networking I/O performance improvements such as support for TCP/IP Segmentation Offload (TSO) and Jumbo Frames. It also includes 64-bit guest operating system support. The vmxnet driver comes with the VMware Tools.

Configuration of the Guest Operating System

As mentioned earlier in this chapter, if you are trying to increase the performance of your virtualization environment, why wouldn’t you attempt to tune your guest operating system? Just because this is a virtual environment, don’t think that many of your normal operating system tweaks used in your physical servers isn’t applicable here. They very well could be the case.

There are numerous books, Web sites and other informational sources available to you that help with common operating system tweaks. This section of the chapter cannot possibly cover them all. Instead, it will discuss some of the tuning options available that are specific to a virtual machine environment.

TIP

Make sure to use supported guest operating systems. VMware Tools may not be available for an unsupported guest.

Updating VMware Tools

As simple as this may seem, it is important to remember to keep your VMware Tools up-to-date inside your guest operating system.

You should always install the latest version of the tools in your virtual machine, and you should ensure that it gets updated after each VMware ESX upgrade or patch. There are times when update patches include fixes for components of VMware Tools, which would make these updates very important to your virtual machines. If you migrate or convert virtual machines from an earlier version of ESX or from another VMware virtualization platform, make sure to remove the old VMware Tools and then install the latest version for your platform. There are differences in the editions of VMware Tools from one VMware platform to the next. If you migrate or convert a virtual machine from another vendor's virtualization platform, make sure to remove that vendor's tools and replace them with VMware's.

By installing the VMware Tools in your guest operating system, you also get these performance enhancement capabilities:

- **VMXNET driver**—As discussed earlier, VMware Tools provides a new updated and enhanced high speed networking driver to help improve networking performance.
- **Improved video graphics driver**—VMware Tools installs an improved graphics driver which gives you better mouse, keyboard and screen performance. In a Windows virtual machine, enable the hardware acceleration feature under the advanced graphics display settings on the Troubleshooting tab. This will help smooth out the mouse when remoting into the virtual machine.
- **BusLogic driver**—Installing VMware Tools updates the BusLogic SCSI driver within the guest operating system to the VMware supplied driver. This VMware version of the driver has certain optimizations that the guest operating system supplied driver does not.
- **Memory balloon driver**—The balloon driver is part of VMware Tools and is used for memory reclamation on ESX. It helps to minimize ESX swapping by better managing guest memory. Memory ballooning will not work without Tools installed.
- **Idler program**—An idler program is added for Netware guests to help de-schedule these guests when they go idle.

- Timer sponge—An experimental timer sponge has been added to give a correct accounting of time within the guests.
- Time sync—When VMware Tools are installed, a time synchronization feature is added to improve the time keeping function within the guests.

TIP

The time synchronization feature is not turned on automatically. Be sure to enable time sync between the ESX host server and the virtual machine from within VMware Tools. When using VMware Tools to do the time synchronization, you should not use another form of time sync within the guest operating system. Do not enable VMware Tools time sync for Windows virtual machines that are members of an Active Directory because they should time sync with a domain controller instead.

Microsoft Windows® Guest Operating System Performance

Many normal tweaks and tuning techniques used to increase the performance of a Microsoft Windows environment on a physical server can translate over to a virtual machine. And just like a physical server, not every virtual machine is alike—so the tuning method used on one machine may not work the same on another.

The latest version of VMware ESX now supports 64-bit Windows guest operating systems as well as 64-bit applications. These 64-bit versions will typically have better performance than their corresponding 32-bit counterparts.

Defragmenting the contents of a virtual machine's hard disk can have a positive effect on its disk I/O performance. Using a third-party tool like Diskeeper can keep your disk structure well organized. Be sure not to schedule the defrag task during critical or normal business hours or else you might cause more performance problems than you fix.

You can free up resources within a virtual machine by stopping and disabling unnecessary services and background tasks. Make sure however that you don't disable a service that is needed by one of your applications. The following list contains a subset of common Windows services that can usually be stopped and disabled in a Windows virtual machine. This list should only be used as a suggestion until you know if a particular service is needed or not.

- Alerter
- Clip Book
- Computer Browser
- DHCP Client (Unless using DHCP IP addresses)
- Fast User Switching Compatibility (Windows XP)
- IMAPI CD-Burning COM Service (Windows XP)

- Indexing Service (Unless needed)
- Internet Connection Firewall (ICF) / Internet Connection Sharing
- IPSEC Services
- Messenger
- Network DDE
- Network DDE DSDM
- Network Location Awareness (NLA)
- Print Spooler (May be needed in some cases)
- Remote Desktop Help Session Manager
- Remote Registry (May be needed)
- Routing and Remote Access (May be needed)
- Smart Card
- SSDP Discovery Service
- System Restore Service (Windows XP)
- Telnet (May be needed)
- Themes
- Uninterruptible Power Supply
- Windows Audio (Windows XP)
- Windows Image Acquisition (WIA)
- Windows Time (May be needed)
- Wireless Zero Configuration

If your guest operating system is Windows XP, you may want to disable the System Restore feature. Doing so will free up system resources like CPU, disk space and I/O. Instead, use the more powerful VMware snapshot feature.

Keep your Windows virtual machine as lean as possible. Uninstall any of the Windows components that aren't going to be used. Doing so will reduce the amount of memory and CPU consumed within the operating system, and will offer more resources to the applications installed in the guest. Power features such as hibernation and hardware power management (turning off hard drives, monitors, etc) really don't have much meaning in a virtual machine or add any value. Likewise, screen savers, visual effects and animations consume additional CPU resources unnecessarily and should be disabled. In most cases, you probably don't need desktop wallpaper on your virtual machine either. One exception to that might be to use something like Microsoft's Sysinternal BgInfo tool which can be very useful by offering system metrics and information as the desktop wallpaper.

Linux Guest Operating Systems Performance

Many normal tweaks and tuning techniques that are used to increase the performance of a Linux environment on a physical server can translate over to a virtual

machine. And just like a physical server, not every virtual machine is alike—so the tuning method used on one machine may not work the same on another.

The latest version of VMware ESX now supports 64-bit guest operating systems as well as 64-bit applications. These 64-bit versions will typically have better performance than their corresponding 32-bit counterparts.

The guest operating system timer rate found in the Linux kernel can impact the performance of a Linux virtual machine. Linux operating systems keep time by counting its timer interrupts. Unpatched 2.4 and earlier Linux kernels request clock interrupts at 100Hz or 100 interrupts per second. By default, some 2.6 Linux kernels request interrupts at 1000Hz while others do so at 250Hz. The overhead involved with delivering so many virtual clock interrupts can unnecessarily stress the virtual hardware and negatively impact performance. The only way to change the behavior is to reduce the number of ticks per second and then recompile the kernel. Where possible, it is best to try and use a Linux kernel that uses lower timer rates to avoid this problem.

If your virtual machine is running a server-class Linux guest operating system and X server is not required for the environment, disable it. Unnecessarily running X server and screen savers in your guest environment will consume extra resources that can negatively affect performance of the virtual machine or its host server. When you do need a graphical desktop, make sure you select a light-weight window manager.

When optimizing your Linux virtual machines, make sure that you disable or remove any unnecessary daemons, services and background tasks. At the same time, be sure to remove any unneeded packages as well. This will free up processor and memory resources that can be redirected to your applications.

Backups and Anti-Virus Configuration

Backup and anti-virus configurations will probably vary based on the use of your virtual environment. In other words, how you configure backups and anti-virus to improve performance in your environment will depend on if your virtual infrastructure is a production environment or a non-production environment such as test and development. A non-production environment may not need to have anti-virus software installed (unless you are testing the anti-virus solution or the effects of anti-virus solutions). However, if this is a production environment running a Windows guest operating system, you will probably be required to have an anti-virus package installed. In that case, it should be configured in the most optimal way in order to help maximize the performance of the virtual machine.

One of the first things you can do is to setup scheduled virus scans to take place during off-peak hours so that the applications in the virtual machine do not compete for resources with the anti-virus solution. Unless your organization specifies otherwise, in most cases, a single, daily virus scan should be adequate. Additionally,

most servers typically do not need real-time virus scanning turned on since this can greatly impact server performance. Disable the real-time virus scanning features when possible for production virtual machines, especially those used as database or web servers. Also, most anti-virus solutions have the ability to exclude certain files, file types, and directories from being scanned. Some application file types, such as database data files, do not need to be scanned, and in fact doing so could destroy performance. You should also configure the anti-virus solution to exclude mission critical application files that are not high risk for virus infection. Swap files are also good candidates for exclusion.

If you are running a backup agent in your virtual machines, you should also schedule the backups to take place during off-peak hours, and equally important, not during a scheduled virus scan. Again, just as with anti-virus software, scheduling backups to take place during off-peak hours will help alleviate performance degradation from applications competing for resources with the backup solution. If your organization is using VMware VI3, it's also a really good idea to make use of VMware's Consolidated Backup feature or use a third-party application that leverages the technology. Doing so will help eliminate a lot of performance overhead within the virtual machines and on your network because the technology will offload much of the overhead onto your SAN and off of your host server.

Summary

If you decide that you don't want to perform any other optimization or tuning suggestion found in this chapter, at least consider upgrading your environment to the latest version of VMware ESX. Doing so will offer you immediate environmental optimization and give you additional tools with which to better operate and control your virtual infrastructure. Upgrading can make it even easier for organizations to virtualize their most resource intensive and demanding applications. Otherwise, each of these best practice performance enhancements can be viewed on a case by case basis as needed.