Global Knowledge™

Expert Reference Series of White Papers

The Value of Virtualization

The Value of Virtualization

David Egan, Global Knowledge Course Director, RHCE, MCSE, PMP®

Introduction

Using virtualization of hardware can save a bundle, when having to test new software releases or software patches. But software that provides virtualized hardware can also save a bundle if you know how licensing works on virtual hosts. There are some very attractive license bundle pricing options that make virtualization even more attractive.

There are many ways you can save money and staff time. Virtualization offers multiple ways to save money including hardware purchase, power and cooling costs over time and staff time needed when providing a test and analysis environment to name a few.

How Did Computing Get to Virtualization Today?

Early computers ran essentially only one program at a time. As time advanced, new computing ideas abounded. The idea of running multiple programs at the same time evolved and multi-tasking was conceived. Then the idea of securing files from other users was developed. The idea of sending files and messages between machines evolved into the networking services and related utilities we know today.

This constant development of these support features developed into a special set of programs, also just software, that is commonly referred to as the operating system or OS. The OS starts with a main program called a kernel that handles memory management and access to hardware plus utilities to manage how multiple user and system programs could run concurrently, independently, and secure from each other in memory. At this point in the development of the OS, each user and system program (also known as an application) was given access to a virtual memory space of its own. By providing separate memory space, these programs performed their functions independently of every other program in memory. Programs in memory, called apps or applications, were unable to "see" the other programs in memory. They "co-existed," unaware of the other programs.

Data Centers and Management

The standard in the early days of computing was keep control of your mission critical systems (computers that ran the really important functions like inventory and payroll) in a data center. This meant up until the 1980s there was (usually) one location for the computer. All users connected from desktop terminals that had no real computing power.

These data centers housed all the computing resources for the company in one climate-controlled and power-managed area. Users rarely actually touched or even saw the computer. Users ran programs (apps or applications) and printed output using a terminal device from their desktops. This was how the mainframes, UNIX- and VMS-based operating systems functioned. Everything was centrally maintained.

The mainframes were usually physically big and required a lot of power and air-conditioning. The next breed was called mini-computers. They were smaller in physical size and supposedly weaker in computing power, according to some pundits. The mini-computer era saw the development of dedicated workstations separate from the normal central main computers, which were commonly referred to as the backend servers. Workstations were computers dedicated to one major task like Computer Aided Design / Computer Aided Manufacturing (CAD/CAM) software programs, and audio or video applications. The other "back office" servers provided email, file, and print services in general.

Take Back Control

In the 1980s and 1990s, a new battle cry was trumpeted by the new kids on the block, "take back control." The new thinking was to get away from central computing environments where there was a single point of failure. Rather than buy one big computer and share the use, why not have your own departmental-sized server?

This was in an effort to wrestle control away from the big computer companies by selling against the folly of a single, central location. They offered businesses the opportunity to take back control, to build their own departmental servers that were just the right size for their use. But what is the "right" size for a departmental server? And why are servers more expensive than desktops generally?

This often led to the sale of over-provisioned central file and print servers in the PC market. Many companies bought the best computer, stuck it literally in a closet that had a network connection, and used it to singularly store all user files. Having all files on this server (or more than one) provided for centralized backup.

Checked your server farm, data closet, under your admin's desk or the datacenter lately?

Cookie Cutter Departmental Servers

Times changed. The needs for a department usually continuously grew.

The new sales battle cry? Back to the centralized, secure datacenter. One location to better manage the disparity and diversity of servers, make them similar!

Out came the blade servers, racks of headless machines that were hardware-compatible. Corporate policy required administrators to make clones of similar servers to handle more service requests at the same time. Each machine was essentially a cloned duplicate of the entire operating system and sometimes even the entire file system. In some cases, where specific applications would "not play together", additional clones were created

that were then customized for specific needs, thereby separating services like Domain Name Service (DNS), Dynamic Host Configuration Protocol (DHCP), Terminal Services, and many others onto their own server.

Why the retreat back into the fortress, behind the big walls of security, management and safety? **Management control and server sizing!**

With the scattering of the servers came the scattering of the resources needed to manage, update, maintain, backup, restore, recover, upgrade, and monitor. Each department usually bought a separate backup service, expanded the hardware with additional disk and RAM as desired, changed system configuration willy-nilly and then wanted centralized support.

We've Created a Monster!

The "monster" grew. Separate file and print services were followed by separate email, web, user accounts, security and access controls, and a plethora of other small services in some cases. The networks became even more of a tangled web of inconsistent, and in many cases, independent machine designs.

Why the separation? Why the isolation? **Why not!**

Mission-critical servers were deemed production servers. Development servers were dedicated to the development of the applications before pushing these applications to the production servers. Test servers were added to test the development server applications or upgrades before they went into production.

How do you support so many different platform changes and client needs as the company grows, adding more computers and servers but not necessarily matching that growth in the Information Technology (IT) department staffing? The stretching of the IT staff resources only got worse and continues to get "worser."

Standardization

The new cry of management, **We need "standardization."**

The IT staff were given the power to take back the control of their network servers. A common dictum was to standardize the equipment and standardize the desktop according to a company wide policy. Usually, this meant moving everything back into one area or data center facility while still retaining the independence of the departments to have their own dedicated computer resources.

Virtualization on the Desktop

Along came a few stalwarts who cleverly noticed what the mainframes were doing and decided to offer this clever leveraging of hardware to the even more massive PC industry. Mainframes had gotten very big indeed, running hundreds of processors and gargantuan amounts of memory, but by using Virtual Memory management, they sliced and diced the machines into virtual partitions of hardware and RAM. They could run Linux in

one partition with XX CPUs and YY RAM, as well as run a Windows OS in another partition using NN CPUs and RR RAM from the host.

Mainframe Ideas for the PC

As the PC became faster and more powerful, it also became another candidate for virtualizing hardware so that multiple operating systems could be run at the same time. The early versions provided a client that sat on top of the operating system and presented a virtualized world to a guest operating system.

This was called a "thick" client design as it required an OS as well as a virtual guest management utility, a lot of layers for the guest OS to go through before it got to the real hardware. Thick client virtualization products include VMWare's Workstation and Server products, Microsoft's Virt PC product, and a host of others. These thick clients ran on top of a general-purpose Linux or Microsoft OS.

Thick clients are good for testing but are less commonly used in a production environment; they are just not robust enough. But they do offer an easy and cheap method of testing new versions of OS software without having to give up the current one. But, and this is a big but, every new OS usually requires even more hardware resources to run.

Back at the Data Center

Meanwhile, every department and every sub-project asks for their own servers dedicated to their needs and isolated from everyone else. Data centers respond by putting in rows and rows of rack-mount towers, followed by upgrades to racks and racks of "blade" servers (look like a pizza box with lots of lights, see Figure 1 below).

Racks have standardized slot sizes. Blade servers come in various widths, based on the standardized rack design, measured in rack units known as U size. Single CPU, single-disk blade servers took up one width and were referred to as 1U in size. Bigger units with multiple CPUs, multiple disk drives, multiple network cards and multiple power supplies might be 2U or more in height. Racks only hold so many of these things before you hit the ceiling—not to mention the power consumption each server requires—whether it is busy or not!

Figure 1. Rack Server – (In white above) Fifteen 1U blade servers on top of two 2U blade servers and two 1U UPS

Data Center Capacity Limits Reached

So, here we sit, today. Datacenters filled to over-capacity with blade servers back-to-back, wall-to-wall in many cases - and we need more! More computing power for these database servers, more horsepower for these web and transaction servers and we notice that most of these other smaller service computers are barely utilizing the CPU or disk space resources they were provisioned with. A new dilemma appears, we have oodles of horse power in all these machines, but we cannot always get to it or shift it quickly or easily.

The New Super Hero Called Virtualization

Enter the new super hero of modern computer systems: hardware virtualization software.

Yes, this is another layer of software. It installs, in essence, as the lower-most layer on the hardware, as a specialized and usually highly customized version of an operating system.

Two Variants of Virtualization

Many versions of virtualization software for the x86 machines were layered on top of an existing operating system like Linux or any Microsoft OS. These are hereinafter referred to as thick clients because there are quite a few layers of software between the guest OS and the real hardware.

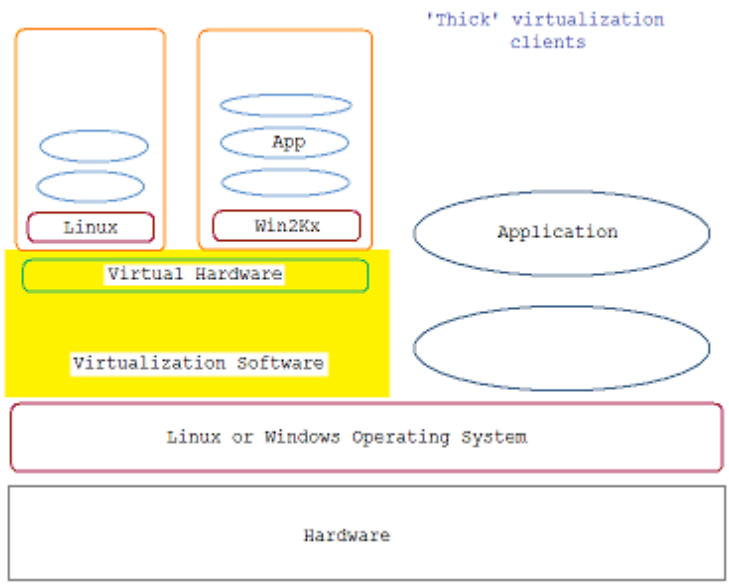

Figure 2. Linux or Windows Operating System

Thick clients, as mentioned earlier, include VMWare's Workstation and Server products, Microsoft's Virt PC product, and others. Each of these products is installed onto another already installed desktop or server operating system on just about any modern computer. The virtualization software is just another application on this operating system and competes for the same CPU and RAM resources with every other application.

The other type of virtualization software stripped the OS down and then added the Application Programming Interfaces (APIs), libraries, and hardware access software drivers specifically designed for virtualization. This produced a much leaner design and is referred to as a lean hypervisor. There are a couple of lean design virtualization products, known as Type I and II. This lean design had as few layers as possible between the guest OS virtualized hardware and the real physical hardware, required far less overhead to run and usually left 95% or more of the CPU time available for the guest operating system VMs.

Hypervisor Competitors in Virtualization

VMWare's ESX 3.5 commercial server product leads the pack with a large market share and is a heavily customized RHEL 3 version of Linux.: http://www.vmware.com/

Parallel's (Formerly SWsoft) Parallels Virtuosos Containers is another popular product: http://www.parallels.com/

Virtual Iron provides a server and enterprise version: http://www.virtualiron.com

Xen is an open source-based, free to use lean hypervisor included with RHEL 5, Fedora Core and some other Linux versions: Http://www.xensource.com/ ==> http://citrix.com/

Microsoft has recently come out with HyperV bundled into some of their Windows 2008 Server editions. This is also a lean hypervisor design according to the press releases, but it does not have all the functionality of the other versions. Most of these extra features are due out in 2010 or later. http://www.microsoft.com/windowsserver2008/en/us/hyperv.aspx

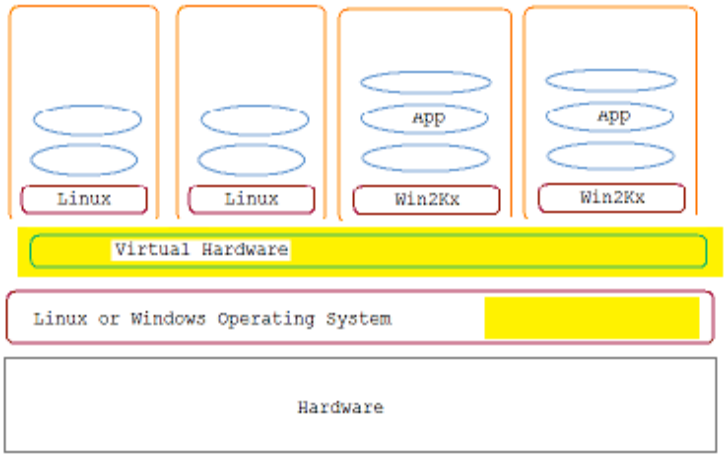

Figure 3. Depicting multiple operating system instances sharing the resources of one server machine.

Virtualization in a nutshell: Provide virtual hardware equivalents to the guest OS software in a secure and shared environment. Fairly manage access to real physical hardware.

In a virtualized computing environment, the physical input and output to real devices from the virtual devices is managed by the hypervisor. The hypervisor tricks the guest OS into believing it has real physical hardware to run on. The real hardware is divided up by the hypervisor, and only the needed virtual hardware is presented to the running guest OS.

This is just plain "wizardry" to the average user. Virtualized hardware, from a managed pool of resources (CPUs, RAM, NICs, disk space) is made available to one or more guest operating systems.

What Does The VM Consist Of?

When running in memory, the guest OS is referred to as a Virtual Machine or VM. The hypervisor is a lean base OS version that provides a virtualized subset of the real hardware to each instance of an operating system.

The set of hardware presented to any given guest OS would be just like the physical hardware of a PC as far as the guest OS is concerned. The virtual hardware would be used by the guest OS to boot-up on and continue running on. The typical virtual hardware presented is any number and combination of (virtual) Network Devices, Video Controller, RAM, raw disk space, keyboard, mouse, CD/DVD, floppy disk, and/or USB ports. The number and choice of these virtual devices is left up to whomever creates the VM to house this guest OS.

Price per License for Each Edition		
Standard*	$	719
Enterprise*	$	2334
Datacenter	$	2381

Figure 4. Windows Server 2003 Pricing per Edition.

Hypervisor Tricks

The hypervisor does many things. It stands between the guest OS virtual hardware and the real hardware, and it manages all input and output requests from all VMs to the various real hardware. The hypervisor manages scheduling of the CPU to each VM and the allocation of needed RAM.

The hypervisor has many tricks up its management sleeve. If a VM is idling, the hypervisor will quickly notice that the VM is killing CPU time in an idle process and will instead disconnect the CPU resource and pass it on to any other VM that really can use those cycles. The VMs do not waste CPU cycles idling or doing nothing unless there really is nothing to do in any other VM.

And yet another trick is to steal back RAM from virtual machines that are not currently needing it. The hypervisor only allocates as much memory as the OS actually uses, not what it was provisioned with. This means that if a VM is provisioned with a maximum of 2 GB of RAM but rarely uses more than 1.5 GB of RAM, the extra .5 GB of RAM is kept back. The VM can at any time ask for the additional RAM or any part of it, and the hypervisor will then provide the extra RAM, always on an as-needed basis.

One more memory management trick is to reduce redundancy between the VM memory maps. If each VM uses the same OS, then we are duplicating every shared library and all the OS files in memory in each VM. The hypervisor looks for duplicate memory pages between all the VMs and removes the duplicate blocks from all other VMs. Only one copy of the blocks that are duplicated is kept and all VMs share a read-only status to that block. Should one of these VMs update the block, then a new block is allocated for the update and that VM retains access to the updated block.

The Savings?

If we have to provision a new machine, it will require its own memory and CPU resources. If we do not know how much memory it will really need, then we usually overprovision the machine. This means we put in as much memory as we can afford at the time and the fastest CPU at the time that makes sense. After spending the money, if we monitor this system and find it is only about 15% to 25% busy on a daily basis and only 73% of RAM is needed, the extra CPU cycles and RAM on this dedicated machine will never be available for any other system.

If instead, we build the machine in this virtual world, we can utilize just what this system would need from a larger pool of resources, and the hypervisor can add or subtract resources as the need changes or arises.

Brilliant! Provide plenty of horse power to all of these now-virtualized machines and let them use what they need, sharing the resources as evenly as possible. The sharing evenly of these resources is the job of the hypervisor.

But, and this is a big but, how else can you save money?

In licensing the old, independent servers, you had a nightmare of licensing issues to keep everything up-to-date, and the independence of each server added to the overall costs of maintaining them.

Microsoft has special deals for datacenter-oriented versions of its many operating systems that allow for unlimited VMs with a singular license. Instead of having to license every machine independently, you can use a centralized license in the datacenter and not have to provision each server with every service that needs a client or server license.

This is where the real money savings occurs. And, you can go to a Microsoft developed web site to prove it! The URL for Windows Server 2003 is:

HTTP://www.microsoft.com/windows/windowsserver2003/howtobuy/licensing/calc_2.htm

Microsoft has created this web site specifically for the USA and Canadian markets as a quick overview of which version of their OS is the best value depending on the number of machines to license. This web site starts with a disclaimer due to the fact that pricing models vary among countries.

Our Scenario: Licensing 200 Win2K3 Machines

Microsoft Licensing

The retail estimated pricing for a standard Win2K3 license is $719 according to this calculating website. It includes the statement: "Standard and Enterprise editions include CALs. A single Standard Edition livence does not grant rights to run both a host OS and guest OS, so an additional license is required for the host OC. Stancard Edition is able to run on 1-4 processor servers and Enterprise Edition is able to run on 1-8 processor Servers."

Let us assume there are 200 blade servers to be virtualized. Each would require a Win2K3 license. The savings start to add up when you use different versions of Win2K.

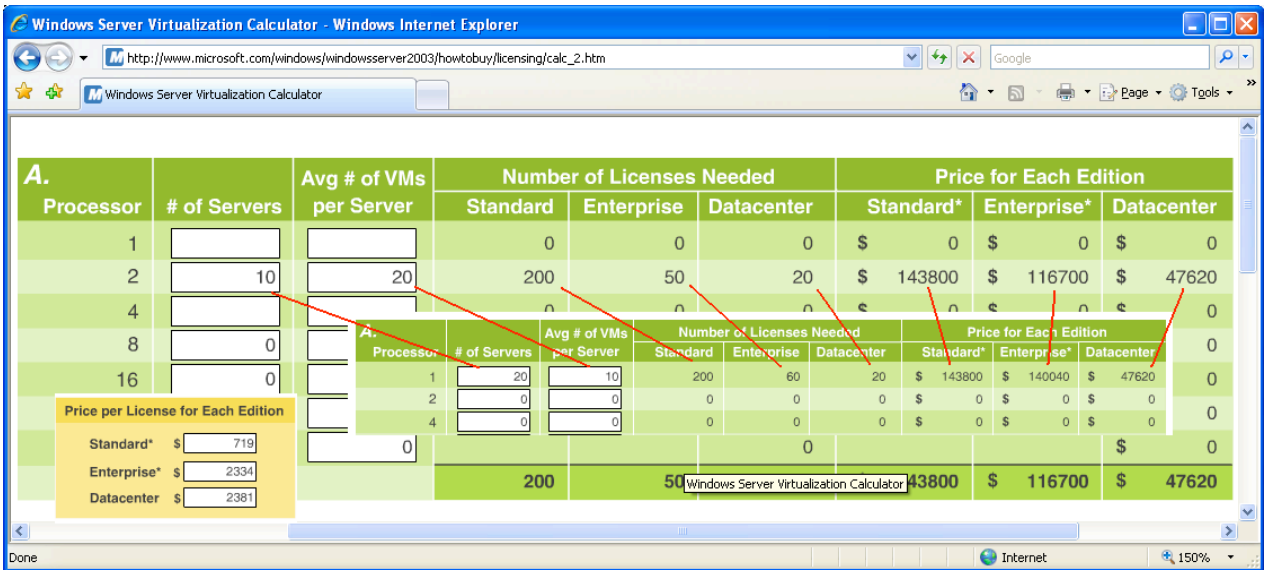

Figure 5. Sample Pricing Analysis

If you had to license 200 independently installed Win2K3 machines, the cost is $143,800 for individual standard Win2k3 licenses. Each of these machines could be single or dual processor, the price does not vary.

You can also compare the pricing based on the number of processors each CPU on the server contains as in Figure 6 below.

B.	Server	# of Processors	# of VMs on Server	Number of Licenses Needed			Price for Each Edition		
				Standard	Enterprise	Datacenter	Standard*	Enterprise*	Datacenter
Ex:	1	1	4	4	1	1	$ 2876	$ 2334	$ 2381
A	2	2	4	4	1	2	$ 2876	$ 2334	$ 4762
	3	4	4	4	1	4	$ 2876	$ 2334	$ 9524
	4	1	5	5	2	1	$ 3595	$ 4668	$ 2381
B	5	2	5	5	2	2	$ 3595	$ 4668	$ 4762
	6	4	5	5	2	4	$ 3595	$ 4668	$ 9524
	7	1	8	8	2	1	$ 5752	$ 4668	$ 2381
C	8	2	8	8	2	2	$ 5752	$ 4668	$ 4762
	9	4	8	8	2	4	$ 5752	$ 4668	$ 9524
	10	1	20	20	5	1	$ 14380	$ 11670	$ 2381
	11	2	20	20	5	2	$ 14380	$ 11670	$ 4762
D	12	4	20	20	5	4	$ 14380	$ 11670	$ 9524
	13	8	20	NA	5	8	$ NA	$ 11670	$ 19048
	14	16	20	NA	NA	16	$ NA	$ NA	$ 38096
	15	32	20	NA	NA	32	$ NA	$ NA	$ 76192

Figure 6. Example Pricing For Individual Servers Based on Processors

In the figure above, when the number of VMs is 1-4, same price; 5-8, price doubles for datacenter version. But 20 VMs on a dual or quad processor is a good deal with regards to licensing costs.

A.		Avg # of VMs	Number of Licenses Needed			Price for Each Edition		
Processor	# of Servers	per Server	Standard	Enterprise	Datacenter	Standard*	Enterprise*	Datacenter
1	10	20	200	50	10	$ 143800	$ 116700	$ 23810
2	10	20	200	50	20	$ 143800	$ 116700	$ 47620
4	10	20	200	50	40	$ 143800	$ 116700	$ 95240

Figure 7. Comparison of Servers with 1, 2 or 4 Processor Sockets

If you bought just 10 dual processor (probably quad core would be your best bang for the buck) machines and could then run on each of these Dual Quad servers 20 VMs having Win2K3 installed, a total of 200 instances of Win2K3 would be available. Unfortunately, you would also either have already paid or will have to pay $143,800 in licensing if you were to have to independently license each instance of Win2K3.

However, if you bought the Enterprise version of Win2K3, you would only have to pay $116,700 for the same 20 VMs on each of the 10 Dual CPU server machines. Enterprise edition of Win2K3 allows for a restricted number of VMs free of charge (4 free Win2K3 VMs). The additional Win2K3 VMs, 16 per server, would require an additional 160 licenses for the rest of the 200 total machines you want to run.

The last alternative is even better, only pay $47,620 if you were to use the Datacenter version on the servers. Why the disparity? The Datacenter version, and this applies to Win2K3 and Win2K8 versions, allows for unlimited VMs to be run within it. You only have to pay for 20 Datacenter licenses.

Pricing Analysis Microsoft

In the 3 pricing scenarios of Figure 11 above, you are getting 200 hosts running Win2K3 as their OS, all are licensed. This scenario should also apply to Win2K8, although the pricing may vary somewhat. You are saving a bundle with the Datacenter versions. Realistically, 20 hosted VMs on one machine would require an "awesome" machine and relatively low CPU and RAM usage per VM, spread out evenly during the day, if there were only one CPU, even if it had 4 processors. Alternatively, a dual Quad-Core server (total of 8 CPUs) with 64 GB RAM would mean each hosted VM would not need more than about 25% of any one CPU all day, and it would not need much more than about 3.5 GB of RAM. A better solution for this scenario of 20 VMs on one host server would be to have a motherboard with four Quad-core CPUs and able to access more than 64 GB RAM. With a Quad Quad-Core Server, each hosted VM would have access to about 80% time usage of a single CPU and not be restricted to, at most, 3.5 GB RAM on average each day.

Pricing Analysis RHEL Support

The Microsoft pricing is just for the operating system. MS-SQL Server is licensed separately. Client Access Licenses for staff usage are also individually licensed. This is not how the Linux world works. Red Hat Enterprise Linux is licensed yearly for support of ALL of its operating systems and all included network services.

RHEL Support Pricing Model

This includes the ful suite of open source server applicaitons, inlcuding Apache, Samba, nfs, ftp, Tomcat, MySQL, PostgreSQL, and network servers. Many other operating sysems charge extra for some of these services.

Red Hat Enterprise Linux			**Red Hat Enterprise Linux Advanced Platform**	
Basic	Standard	Premium	Standard	Premium
$349	$1,299	$1,299	$1,499	$2,499

Full details are at http://www.redhat.com/rhel/server/compare/.

Imagine the savings if you could replace any of these Microsoft licensed servers with RHEL 5, at just about $1,300 to $2,500 per server machine for licensed support, and as little as $350 for individual instances, you immediately see significant savings.

Testing Environment Savings

Notice that your best bang for your buck is in quad-core processors on each server machine (at least one but 2 is better) when licensing for Windows Operating Systems. With a Datacenter license on any machine, you have the ability to create an unlimited number of Virtual Machines providing your staff with the ability to test systems without requiring additional licensing of these VMs.

Every time your support staff have to test a new product, new upgrade, new application or a patch, you will need another machine that is identically provisioned to the one in production. This is not a problem in the virtual world as all virtual machines "see" the same general hardware. Testing is a breeze in comparison to trying to find another identical host or having to have purchased a backup of your production server for testing purposes.

Sample Cost Savings

The cost of 200 server computers from a reputable manufacturer, with support contract and installed. is estimated to be $3,000 to $5,000 each. Even at just $3,000 per server, 200 servers is $600,000 worth of hardware.

Hardware Savings Can Also Be HUGE!

What if you could replace these 200 machines with just 10 bigger machines? The cost of the bigger machines would be much higher, somewhere between $10,000 to $20,000. Even at $20,000 for a top-of-the-line machine with multiple quad-core CPUs, tons of RAM, and SAN connected storage, for 10 of these servers, the cost is only $200,000 for the same 200 running instances of Win2K3 (or Linux for that matter). Even if 20 VMs seems like too many eggs in one basket, at 10 VMs per server, means you need 20 servers at $20,000 each or $400,000 total. Still a significant savings, assuming you have a standard corporate policy to replace machines at the end of their useful life span of 2 to 3 years.

If you wanted to spend the same amount of money, $600,000, you could get 10 servers at $40,000 and then spend $200,000 on extras like SAN or NAS storage!

Imagine how much disk space is available on all 200 of these machines. Is it likely to be fully used or less than half full? Consolidation of disk space to half or less will also save money in the long run. And what if you need more disk space on just a few and less on others? Over-provisioning will soon be less of an issue than ever before.

Where Else Can You Save Money?

Do not forget about power savings - hardware and air conditioning! A typical rack of servers can consume 40 to 60 amps of power 24x7. Reducing the rack space to half the number of physical boxes will save power requirements significantly and immediately provide savings. If you house your own servers, you also provide your own cooling of these servers or datacenters. Reduce the number of machines needing to be cooled, and you again gain immediate savings.

The math is easy on this one. 200 servers down to 20 machines that are unlikely to draw 10 times the power of the older servers, more likely 2 to 3 times the power which is still at least a 50% savings in power consumption and cooling costs. This alone is huge!

Proof of Concept

With Xen, you can use almost any hardware as a starting point. Unlike VMWare's very specific hardware requirements, Xen is happy to work on almost any type of PC that you can install RHEL 5 on.

You may have some retired or current hardware given to some staff or departments that were pretty beefy models at the time. Maybe these older machines, say 1 to 2 years old, are or were someone's desktop. Are they really using all those cycles? If not, you may be able to use them for some lower-need VMs, if you were to take them back and put these boxes into service.

Caveat Emptor – Do You Need This?

The hype around virtualization software is incredible, but do you need it? The consensus among the fortune 1000 is they have no choice, they desperately need to reduce the space, reduce the power needs, and improve the reliability of their systems.

If you were to look at your own environment, at 20 to 200 servers sitting in a rack, would each of these hosts be a candidate for virtualization of the hardware? Yes and no is the best answer. Yes, technically you could virtualize everything, or just the weakest or simplest of servers. A thorough analysis of your current computing services and needs now and in the future may lead you towards virtualization as the ultimate answer because of the load-balancing, management, and reduced sizes alone.

Summary

Imagine if you could actually combine ALL the CPU cycles on every machine in your office into one big computing device, that you could then allocate out these resources on an as-needed basis instead of giving each individual user more machine than they needed!

Imagine the possibilities!

Sorry, you just cannot do that yet, but you just know someone in the open source world is thinking about it. Remember that RAID and SCSI, someone thought "outside the box!"

It is hard not to hear all the hype. It is harder to believe some of the outlandish savings that are being advertised, and yet, there has to be a reason all the big hardware suppliers, like IBM, Sun Microsystems, Dell, and HP, are jumping on the bandwagon of virtualization. Yes, it sells more hardware, but you are constantly looking for more hardware anyways. Now you can reduce some of the clutter too!

If nothing else, the added management options, ease of setup and maintenance, ease of centralized management, the security of really good, redundant hardware with low mean times to failure, the high availability aspect, and the general cost savings of embracing virtualization is hard to ignore!

Just remember, the entry cost may seem very steep to begin with, but it does not have to be. Starting slowly is never a bad idea. Proof of concept should not be too expensive, nor too time-consuming, nor costly. Even if you can just reduce your server farm a few machines now and then, every little bit helps save you money in the long run!

Learn More

Learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge. Check out the following Global Knowledge courses:

Citrix + Xen Oriented Classes

Citrix Access Gateway 4.5 Advanced Edition: Administration – CTX-1308AI

Citrix Access Gateway 8.0 Enterprise Editions: Administration – CTX-1309AI

Citrix Access Suite 4.0: Build/Test Workshop – CTX-1456AI

Citrix Password Manager 4.5: Administration – CTX-1327AI

Citrix XenApp (Presentation Server 4.5): Administration – CTX-1259

Citrix XenApp (Presentation Server 4.5): Support – CTX-1264

IBM Training

Managing IBM System x and BladeCenter Servers with IBM Director 5.20 – XTR43

System p LPAR and Virtualization I: Planning and Cofiguration – AU730

System p LPAR and Virtualization II: Implementing Advance Configurations – AU780

System p LPAR and Virtualization III: Virtualization Performance Management – AU750

Microsoft Training

Managing a Virtual Environment

Red Hat + Xen Offerings

RH401 Red Hat Enterprise Deployment, Virtualization, and Systems Management

RH401 Red Had Enterprise Deployment, Virtualization, and System Mgmt Exam – EX401

VMWare Training

VMWare Infrastructure 3: Deploy, Secure, And Analyze v3.5

VMWare Infrastructure 3: Fast Track v3.5

VMWare Infrastructure 3: Install and Configure v3.5

For more information or to register, visit **www.globalknowledge.com** or call **1-800-COURSES** to speak with a sales representative.

Our courses and enhanced, hands-on labs offer practical skills and tips that you can immediately put to use. Our expert instructors draw upon their experiences to help you understand key concepts and how to apply them to your specific work situation. Choose from our more than 700 courses, delivered through Classrooms, e-Learning, and On-site sessions, to meet your IT and management training needs.

About the Author

David Egan is an engineer by education but a true computer geek by experience. He has over 30 years' experience in the computer industry, working with large and small companies and on three continents. He has risen through the ranks of programmer, systems analyst, systems developer, systems architect, project manager, and now virtualization evangelist. He has been an avid fan and advocate of Hardware Virtualization ever since the early versions appeared for the Intel/AMD platforms.

He has been using virtualization software on PC hardware for over seven years while consulting for Red Hat, Inc. and Global Knowledge in order to deliver cutting-edge services to Fortune 1000 companies. His virtualization experience includes working with and supporting VMWare, XEN, and Virtual-PC. His systems' experience and certifications include VMS from Digital Equipment Corporation, many Windows versions, and Red Hat Enterprise Linux.

Organizations he has helped include various levels of government and military, FBI, CIA, NSA, NASA, as well as Fortune 1000 companies such as Goldman Sachs, Chase Manhattan, Schlumberger, Google, AOL, and Boeing.

He can be reached at DavidW@Vaddac.com

Historically Speaking

IBM

IBM was one of the founding computer companies of the 1950s and is still around today. IBM is also the original developer of virtual machines with their various operating systems and equipment.

IBM has been partitioning their systems, dividing the CPU+RAM, into chunks or partitions that are made available to individual operating systems. Many a PC-based company was created based on ideas and or people from IBM, and this includes Citrix.

Citrix + Xen

Citrix was specifically developing a service of centralized hosts running general applications that were available anywhere to anyone securely. This lessened the need for regular PC upgrades. The central hosts were constantly being upgraded and provided all the horsepower needed. The PC was literally just a fancy TV-like terminal appli-

ance. Any such terminal from anywhere could access the desktop of any user from essentially anywhere. Lots of competition has evolved from this, including the latest offering of Windows 2008 from Microsoft having Terminal Server Services that rival the main features of the Citrix WinFrame product.

Citrix is one of the companies that started all this PC-oriented virtualization way back in the late 1980s. Back then, the Citrix (from Citrus and UNIX) design was to provide applications from a central host to any desktops via the network without having to change the desktops so often. Most PCs were replaced every 3 years because the applications required more horsepower every year. Instead, with Citrix, you upgraded the central host and simply sent the display to any old or new computer. This extended the life of older computers and provided savings related to Total Cost of Ownership (TCO).

Citrix started in 1989, but until 1993, when Citrix negotiated an agreement not to compete with Windows Terminal Services (which was based on Citrix products), and instead sell an add-on product Winframe for Windows, they earned nothing. Winframe turned their fortunes around. Companies soon found the benefit of virtualizing all their applications on a central host and then using any remote desktop to access them securely, whether at any company office or at home, helped extend the life of older PC's, reducing TCO.

Citrix has been expanding their services ever since and are still offering centralized computing services that are network available on client machines via their own highly successful Independent Computing Architecture (ICA) protocol. ICA set down rules for server-to-client communications without restrictions to what the server or client had to be.

Citrix Winframe ran on almost any OS. Citrix was one of the main proponents of thin clients, desktop or portable computers with no local resources, just the ability to connect to a network. Everything else came from the Citrix server. The dream was to replace all the hardware at all the desktops with color terminals having no disk drives, no CDROM/DVD, no USB ports, no floppies - no way to add devices. The thin client had nothing extra, just an input device (keyboard, maybe a mouse, if not touch screen) and a screen. That was the thin aspect of the client. It could physically be made thin as well since there was little hardware needed, a motherboard with RAM and video display.

The latest MAC computer offerings are thin client design, but are not really thin clients; they are far from it. Microsoft, MAC, and Linux desktops all provide Remote Desktop Protocol (RDP) clients that can connect to the many central host computing environments including Citrix.

The biggest drawback of these centralized host computing environments is that the network connection becomes critical to everything. Graphical applications require a lot of pixel manipulation, all of which must come via the network connection and need the highest level of compression possible.

VMWare

VMWare is specifically just virtualization, either as a layer on top of an existing general OS (Workstation and Server versions) or as a full OS-included version (ESX runs on a slimmed down Red Hat Enterprise Linux OS version).

VMWare only does full virtualization, nothing else. Their whole existence is to reduce the required footprints of hardware by virtualizing the physical hardware such that almost every known OS can be installed as a VM within their flagship products: VMWare Workstation (single host license), VMWare ESX 3.5 (Datacenter host license) that is normally coupled with Virtual Infrastructure Client product (adds High Availability and Clustering services to ESX).

RHEL + Xen

Red Hat Enterpise Linux has included Xen-enabled kernels with the last few versions and is adding another version, Kernel Virtual Machine (KVM) in fall of 2008. Xen offers both paravirtual and fully virtualized environments.

Microsoft Virtual PC, Virtual Server, and HyperV

Microsoft introduced their desktop-oriented Virtual PC product a few years ago in competition with VMWare Workstation and equivalents. This is a thick client that layers on top of an existing Windows OS already installed on the host machine. Microsoft also introduced their server product MS Virtual Server with Windows 2003 and have recently announced Hyper V virtual server with Windows 2008. Although relatively late to the game, they are still a formidable opponent in any arena they join.

Licensing

The biggest difference between all these virtualization offerings is licensing. All but Xen require a license for the product itself and or the supporting OS. Licensing for the OS encased in the VM is on a per-OS-version basis in all virtualization software. For instance, if you run Windows as a guest OS VM on a Xen-enabled Red Hat Enterprise Linux host, you still need the server and CAL licenses for the Windows OS as well as the RHEL support licensing.