# The Rise of Server Virtualization

## Learn about the benefits of virtualization, current solutions and what's coming down the pike.

With demands that exceed the available resources, many organizations are looking to leverage server virtualization to maximize the return on investment for their server hardware and to operate more efficiently. This white paper provides a framework for understanding what virtualization is and includes a side-by-side comparison of the Microsoft and VMware solutions. The topics of alternative server virtualization applications and the next generation of server virtualization solutions are also discussed.

### Doing More With Less

Today's world is fast paced, and the world of IT is even faster. Organizations continue to leverage their computer systems to migrate functions and tasks and increase their efficiency and productivity. Systems that were previously separate, such as telephone and e-mail, are now converging, placing added burden on the IT infrastructure to provide the capabilities and quality of service that their users expect.

IT departments don't operate with unlimited resources, though. The natural growth of an organization's computer infrastructure needs, combined with the accelerated reliance on technology to manage more functions, means that the IT department needs to add hardware and software at a pace that many organizations can't afford to maintain. There are also physical limitations, such as running out of rack space in the data center and having nowhere to physically place a new server.

Server virtualization can solve these issues and then some. Server virtualization enables organizations to run multiple "virtual" computers on a single physical server, allowing them to operate more efficiently and maximize their investment in computer and networking hardware. Each virtual computer runs as a separate instance and is, for all intents and purposes, a complete and self-contained computer system. Virtual server images can be installed quickly as new resources are needed, and in the event of a hardware failure, a backup of the server image can be moved to new hardware and be back online in a fraction of the time that normal server disaster recovery takes.

## Two Tips That Pay Dividends

- Look at all aspects of server virtualization potential to define ROI benefits that might be overlooked if they are not specifically enumerated. For example, virtualization might let you toss out old, outdated equipment that's too expensive and difficult to maintain, while concurrently allowing new functions that previously had not been considered because of the legacy equipment's limitations.
- Although virtualization can help squeeze more output from less equipment, it won't do anything to alleviate the reluctance of some organizations to share space on the same equipment. Bottom line: To reap ROI, you will need to resolve any culture, turf or security issues.

### Benefits of Virtualization

There are a number of benefits to server or desktop virtualization. Here are some of the most common usage scenarios to give you an idea of how virtualization can be leveraged for your organization:

- **Legacy support:** IT does not stand still for anyone. Hardware continues to evolve with exponential leaps in functionality. Software, on the other hand, does not always keep up, particularly specialized, industry-specific software solutions. Your organization may want to upgrade all server hardware to operate on Windows Server 2003 to take advantage of the advances in server management and security, but you may have legacy applications that will only function on Windows 2000. Rather than keeping an outdated Windows 2000 Server on your network, you can update your servers to Windows Server 2003 and execute your legacy application within a virtual server Windows 2000 environment.
- **Cost reduction:** Using virtual servers, organizations can deploy new servers without the associated investment in hardware. As needs arise, a new virtual server can be implemented for the cost of the operating system license.
- **Efficient deployment:** Rather than installing hardware on a rack in a data center and then installing an operating system from scratch, organizations can create base images of standard-build servers and have a new server up and running in a virtual environment in minutes rather than hours.
- **Consolidate data center servers:** Servers use power and generate heat. The heat then requires even more electricity to be devoted to cooling the data center. Servers also take up space. By running multiple virtual machines on each physical server, organizations can get the same net productivity and user experience, but with a fraction of the space, power and cooling required to run separate individual servers.
- **Disaster recovery:** Even a good disaster recovery plan may take hours to execute and return a server to functionality, and that assumes that equivalent replacement hardware is available. Attempting to restore a backup or even an image to a different hardware architecture is tricky at best, if it is possible at all. A virtual machine is an entire server environment basically contained within a file. Periodically backing up the virtual machine data provides an organization with the ability to simply run the virtual machine from a different server and return to functionality in minutes.
- **Reduced maintenance overhead:** Virtualization can also save you significant time when you have to bring up a new server. By staging server images with the different operating systems you may need, the time to bring up a new server can be reduced dramatically. For example, an IT administrator can build an image of Windows Server 2003 with the latest service pack and hot fixes, but not add it to the existing domain. When a new server is needed, an administrator can copy this virtual server disk file on the host to a different location, rename it and then create a new virtual server on the host and add in the virtual disk that was just copied.

When you start the virtual server, all you have to do is rename the server, join the domain and install the latest hot fixes. This entire process can be completed within 15 minutes, compared with an hour or more to build the server from scratch.

- **Multiple hosts will improve fault tolerance:** If your organization deploys numerous servers that run various applications, it may be possible to consolidate these servers onto a smaller number of physical hosts and still provide fault tolerance if one of the hosts is lost. If the servers are due for upgrades, instead of purchasing one new server per operating environment and application, you may be able to consolidate virtual server guests onto one physical host. Of course, these servers will have more processor, memory and disk capacity than single-purpose servers. In this multiple-server scenario, you can even configure the virtual guests as "warm" guests on a different host server.

### How Virtualization Works

Server virtualization is a computer "shell game" in which virtual machine monitors allocate system resources to the various virtual computers to trick each of them into thinking that they are the only operating system running on that hardware. The virtual machine monitor exists as a component of the hypervisor, a layer between the true physical hardware and the various virtual computers running on the system (see Introduction to Hypervisors).

On a standard, single operating system server, the hardware resources such as memory, hard drive space and CPU often sit idle waiting for something to do. Virtualization allows that idle time to be distributed across multiple virtual computers and maximize the productivity of a given physical computer.

x86 architecture provides four levels of privileges called rings. Operating systems typically require direct access to interact with the core hardware architecture, otherwise known as Ring 0. The applications that users install and run within the operating system generally function with the lowest privileges, or Ring 3.

The CPU architecture is not inherently designed to be controlled by more than one operating system at a time, and the operating systems themselves do not have a mechanism for sharing control. The hypervisor layer does the sharing of control on behalf of the CPU and other hardware resources, and it does the arbitration on behalf of the operating systems trying to access those resources.

The challenge with x86 architecture was that certain CPU instructions resulted in errors or system crashes when virtualized. The original solution for virtual computing on x86 architecture, devised by VMware in 1998, was to trap the problem instructions, convert them into instructions that can execute safely in a virtual environment and forward them to the CPU.

Currently, there are three basic approaches to virtualization: full virtualization, paravirtualization and hardware assisted virtualization. Here is a brief explanation of each:

#### Full Virtualization

Full virtualization completely abstracts, or separates, the virtual operating system from the physical hardware architecture. This type of virtualization uses binary translation to trap specific CPU instructions and convert them to code that will work in a virtual machine, while most instructions are passed straight through for better performance.

This type of virtualization can be implemented directly between the hardware and the virtual machines, or as an application running in an operating system installed on the hardware. In either event, the virtual, or guest, operating system is not aware that it has been virtualized and does not require any updates or modifications.

Full virtualization provides better security and portability for virtual machines than other virtualization solutions. The most common virtualization solutions, such as Microsoft Virtual Server and VMware Server, use full virtualization.

#### Paravirtualization

A virtualization platform that utilizes paravirtualization may offer better performance than its full virtualization cousins, but with some significant caveats as well. Unlike full virtualization, paravirtualization does require that the guest operating system be modified. Rather than using binary translation

to trap and convert CPU instructions, paravirtualization modifies the guest operating system kernel to replace those instruction requests with hypercalls that communicate directly with the hypervisor layer.

The catch is that the guest operating system has to be capable of being modified. That means the closed-source, proprietary code. Some versions of Windows have been engineered to work with paravirtualization, but Windows 2000 and Windows XP, both still used heavily, are not compatible with paravirtualization.

Paravirtualization is more closely tied to both the kernels of the guest operating systems, as well as to the underlying hardware. These factors make upgrading or maintaining the operating systems and the physical server more complex and sensitive than with full virtualization.

## Hardware Assisted Virtualization

On the opposite end of the spectrum from modifying the operating system to enable virtualization is modifying the CPU to enable virtualization. This is called hardware assisted virtualization. Both Intel and AMD have developed new processor technology with specialized instruction sets to enhance virtualization capabilities.

Intel has developed Virtualization Technology (VT-x), and AMD has developed their AMD-V solution. While vendors have their own approach, the goal is the same. The CPU architecture itself has been enhanced to allow virtual machines to actually run with higher privileges than even Ring 0. The CPU traps virtual machine CPU instructions and passes them directly to the hypervisor.

Hardware assisted virtualization is a relatively new technology. In its current implementation, the processing overhead that the CPU must manage in order to trap sensitive calls impacts performance. The result is that full virtualization and paravirtualization offer superior performance in many cases. Hardware assisted virtualization has promise, and the performance will improve as Intel and AMD refine their CPU architectures, and newer operating system and virtualization software is developed with the enhanced virtualization hardware in mind.

## Introduction to Hypervisors

The concept of server virtualization dates back more than 40 years to early IBM mainframe days. IBM was the first to develop a system capable of running multiple operating systems within a virtual environment with their CP-40 system created in 1967.

The term hypervisor, commonly used to refer to virtual machine monitor applications, is also derived from the IBM mainframe days. The operating system kernel was called the "supervisor." In order to access the kernel-level system resources, the virtual server software had to be able to communicate with the system supervisor. Hypervisor was used to refer to the ability of the virtual operating system to make system calls to the supervisor and gain direct access to system-level resources.

In modern terms, a hypervisor is a layer that runs directly on the "bare-metal" x86 architecture without an underlying operating system and provides virtualization services. This approach provides more direct and more efficient control of hardware resources, and less system overhead due to the lack of resources being drained by a running operating system.

Implementations include VMware's ESX Server and the upcoming Microsoft Hyper-V product. Virtual server platforms such as these enable you to abstract the hardware resources of the underlying physical computer and virtualize it so that a complete virtual computer — with its own CPU, RAM memory, hard drive, network connectivity, etc. — is created.



The job of the hypervisor is to play liaison between the virtual operating systems and the physical hardware. The hypervisor intercepts requests and dynamically allocates the real hardware resources so that the different virtual machines can run simultaneously on a single physical device without hardware conflicts.

Normally, a computer is only capable of running a single operating system. Operating systems are not designed to cooperate, or co-exist, with other operating systems. With a hypervisor, multiple operating systems can run on a given computer system without having any awareness that other operating systems are running on the same hardware.

The hypervisor even enables heterogonous platforms to co-exist, such as running Windows and Linux virtual computers on the same hardware. Rather than having to allocate and set up multiple computers, a hypervisor can be used to leverage the hardware resources and run disparate operating systems side by side.

In addition, hypervisors are generally able to emulate different instruction set architectures (ISA). The ISA is the underlying instruction set that defines, at a low level, how the computer system works. This includes the interrupt handling, memory architecture, CPU instructions, addressing modes and more.

The ability for the hypervisor to emulate different instruction set architectures provides even greater flexibility to run various operating systems within reason regardless of their specific hardware requirements.

The hypervisor is the engine that drives server virtualization and the brain that allows it to function. Separating the true hardware resources from the virtual machines, and allocating resources as needed, the hypervisor enables the multiple operating systems to function simultaneously so that more efficient server and desktop virtualization are possible.

## Microsoft vs. VMware

Virtualization is a hot industry, and many vendors are looking to capitalize on its rising popularity. The major players in the arena are Microsoft and VMware. While VMware pioneered the x86 virtualization industry and has offered products for nearly 10 years, Microsoft entered the virtualization scene in 2003 with the purchase of Connectix, makers of Virtual PC.

In 2006, Microsoft released Microsoft Virtual PC 2004 as a free product. The current version, Virtual PC 2007, remains free, while VMware's VMware Workstation 6 for Windows sells for approximately $189.

Microsoft carried the competition to the server level as well, releasing Virtual Server 2005 for free. VMware followed suit, making VMware Server available for free. VMware does provide support for VMware Server for a fee.

Both Microsoft and VMware also provide tools for centralized management and administration of multiple virtual servers.

Microsoft's System Center Virtual Machine Manager provides organizations with a familiar interface and user experience, as well as solid integration with other Microsoft products. The System Center Virtual Machine Manager has tools to assist with planning and architecting the virtual infrastructure and provides the ability to convert physical servers to virtual servers and centralize deployment for greater efficiency.

VMware has much of the same functionality built into their VirtualCenter for VMware Server product. VMware benefits from their maturity in the market with wizards and templates that make administration and deployment simpler. VMware also provides optional tools, such as VMotion, which can migrate a live, running virtual machine from one VMware Server to another with no downtime or impact to users.

This functionality does come at a price. The VMotion option is only available to users of high-end VMware solutions such as ESX and requires both VirtualCenter Server and a robust fibre-optic LAN infrastructure.

The Microsoft solution is a solid option, particularly for organizations that rely primarily on Microsoft for their operating systems and server infrastructure. VMware provides more market experience and a more diverse capability to host diverse operating systems, including NetWare and Solaris x86.

## The Alternatives

While Microsoft and VMware are the marquee players, they are not the only players in the market. The open-source Xen project, the popular Mac OS application Parallels and the relative newcomer Virtual Iron stand out as alternatives to these dominant players.

## Being Xen

Xen was founded as an open-source project by Ian Pratt at the University of Cambridge. In October of 2007, XenSource was purchased by Citrix for $500 million, providing some added legitimacy for the open-source application.

Being open source, Xen is available at no cost. It can be used for server virtualization on a variety of hardware architectures, including x86, x86-64 bit, PowerPC and others. Xen is capable of hosting virtual machines for Windows, Linux, Solaris and some BSD operating system versions.

Xen uses a slightly different virtualization model than VMware and Microsoft. The Xen virtualization relies on paravirtualization, which modifies the guest operating system to use hypercalls that allow it to access the core hardware functionality. On an x86 platform, the Xen kernel runs in Ring 0, while the guest operating system runs in Ring 1. Through paravirtualization, Xen can run more efficiently and achieve higher performance. Paravirtualization requires that the operating system is capable of being modified. Closed source code is not open to such modification and can not run in the native Xen paravirtualization environment. In order for Xen to be able to host Windows in a virtual machine, you must be running at least Xen 3.0, and the underlying hardware must be capable of hardware assisted virtualization.

Intel and AMD have both developed extensions to their instruction sets that provide enhanced functionality for virtualization. In order to host a Windows operating system as a virtual machine, Xen must be running on a system that supports either Intel's VT-x or AMD's AMD-V technologies.

There are a couple of commercial products available that rely on Xen for their virtualization foundation. Citrix and Oracle are both recognized names in the world of IT that are marketing Xen-based virtualization solutions.

## Citrix XenServer

After the purchase of XenSource by Citrix, Xen has also been the basis for their server virtualization offerings. Citrix offers three versions of their virtualization software: Citrix XenServer Enterprise Edition, Citrix XenServer Standard Edition and Citrix XenServer Express Edition.

As an open-source project that is evolving into a market player, it still lacks some of the luster of the VMware and Microsoft offerings, but it makes up for it in cost savings. A Citrix XenServer Enterprise v4 license is about $2,500, plus $500 a month for support, compared with $5,750, plus $1,200 a month for support if you want VMware Virtual Infrastructure 3. Organizations that are willing to get the virtualization power without some of the bells and whistles may find the difference in investment compelling.

Even as a relatively new offering, Xen has some advanced functionality, such as the ability to migrate a live virtual machine across the LAN to a different physical machine. VMware provides this functionality as well through their VMotion tool, but for a fee.

## Oracle VM Manager

The Oracle VM Manager utilizes the open-source Xen server virtualization software, combined with a proprietary Web front end. The Oracle VM Manager provides the same Xen capabilities to host Linux, Windows or other operating systems. In addition, Oracle provides a Web interface with similar functionality to the Microsoft System Center Virtual Machine Manager or VMware VirtualCenter. Users can create, clone, share, boot, configure or migrate virtual machines using the Web interface.

## Parallel Universe

The move by Apple to build Mac hardware on the ubiquitous Intel chipset also allowed the Mac OS crowd to join the world of virtualization. The makers of Parallels Workstation for Windows and Linux released a Mac version dubbed Parallels Desktop for Mac.

Parallels Desktop for Mac is a popular utility for Mac users to be able to run Windows and/or Linux operating systems from within their Mac system. Parallels uses full virtualization to fully abstract the guest operating systems from the underlying hardware. It also provides access for the guest operating systems to additional hardware resources such as USB and parallel ports.

## Virtual Iron

Virtual Iron is one of the first virtualization vendors to jump into the deep end of the pool when it comes to adopting hardware assisted virtualization. Virtual Iron is based on the open-source Xen code, but Virtual Iron has chosen to forego paravirtualization and adapt Xen to run completely in "native virtualization."

Native virtualization is the term Virtual Iron uses to describe their process of leveraging next-generation CU hardware from Intel and AMD to run unmodified guest operating systems. Virtual Iron requires that the underlying CPU support either Intel VT or AMD-V hardware virtualization enhancements.

## Running Bare Metal

Over the past 10 years, almost every instance of virtual computing has involved installing a virtual machine monitor application on top of an existing operating system. In order to run virtual machines on a computer, you first have to install an operating system, then a virtual computing solution that is compatible with that operating system.

This adds processor and resource overhead, as the underlying hardware architecture must support the host operating system as well as the virtual machine monitor and each of its guest operating systems. It also adds an unnecessary implementation and maintenance burden in order to install and update the underlying operating system.

The next generation of virtualization is to run a pure hypervisor directly on the hardware. VMware's ESX Server and Microsoft's upcoming Hyper-V are two examples of virtualization applications that can be installed directly on the physical hardware without the added overhead of an underlying operating system.

VMware is taking the bare metal approach even further with the development of VMware ESX Server 3i. They have reduced the hypervisor code to 32MB of space. As a hard drive install, the smaller hypervisor simplifies the complexity of the code and minimizes the opportunity for vulnerabilities, creating a more secure environment. It is also faster and simpler to install. Server vendors are working to introduce new hardware with the VMware ESX Server 3i embedded in firmware, creating a hypervisor-ready machine right out of the box.

## Maturing to Enterprise Class

Organizations may appreciate the benefits that server virtualization can provide, but they also need the rest of the enterprise-level functionality they currently have with their physical hardware. Things like load balancing and fault tolerance are vital to mission-critical server applications, and the virtual server environment has to be at least as robust as its physical counterpart. As you explore the various virtual server options, here is a list of features or functions you should consider:

### Automatic Restart

The virtual machine manager should be capable of automatically restarting a virtual machine that has stopped or crashed for some reason. Servers that are essential to productivity can not afford to be down, and the virtual computing infrastructure should be able to manage incidents like this.

### Failover

For truly critical systems, organizations rely on clustering and automatic failover to ensure that the system or application is always available. Make sure that the virtual infrastructure solution you implement is capable of providing the type of clustering and server failover functionality you need.

### Load Balancing

Similar to failover, but with a slightly different application. Heavily used systems such as Web servers are generally load-balanced across multiple servers. If you need to virtualize an application or server farm that needs to be load-balanced, consider the features of each virtual infrastructure solution and ensure that the vendor you choose can deliver the functions you require.

### Migrating Virtual Machines

Hardware will eventually need to be upgraded or replaced. You may also wish to move virtual machines from one physical device to another just for architectural or organizational purposes. Your ability to easily migrate a virtual machine will depend on whether your virtualization solution is based on full virtualization, paravirtualization or hardware assisted virtualization.

The efficiency of backup and disaster recovery of virtual servers is also impacted by the type of virtualization being used. Virtualization can make disaster recovery significantly faster and easier than traditional methods, but the ability of your virtualization solution to easily move from one physical server to another will affect your ability to restore a backed up virtual server image to new hardware in the case of a disaster.

Full virtualization offers the best portability between physical hardware, but even with a full virtualization solution you need to make sure that the vendor you choose has the tools and technology to migrate a live, running virtual machine without any downtime or impact to your users.

### Accessing Hardware Resources

All of the virtual computing solutions virtualize hardware components such as the CPU, memory, hard drive and network connectivity. Each virtual machine is, in essence, its own self-contained computer. However, not all virtualization solutions provide access to other hardware resources on the physical machine such as the USB or firewire ports.

### Virtual Infrastructure Management

For an organization deploying a virtual infrastructure, this is perhaps the most important feature. Vendors have their own approach to centralized management of virtual machines. Some vendors provide virtual infrastructure management tools that are restricted to managing their own virtual machines, while others provide management functionality capable of overseeing and managing virtual machines from all major vendors.

## The Heat Is On

VMware pioneered x86 virtualization, and they have led the way for a decade. They enjoy marketshare and wide acceptance as the de facto virtualization solution. Previous Microsoft virtual computing applications have put a dent in VMware's market, but have not been a significant threat. With Hyper-V, and the upcoming Virtual Machine Manager 2, Microsoft is aggressively pursuing a bigger piece of that pie.

The increased competition is a positive thing for customers looking to leverage the benefits of virtual computing. Microsoft Hyper-V, combined with Citrix adding enterprise-class tools and legitimacy to the open-source Xen product, mean the landscape is widening and there are more options to explore.