# Demystifying Virtualization:

**Ty Schwab**

**10/31/2006**

## Introduction:

If you take the broad definition of virtualization and break it down to its core meaning, it can be summed up as follows: Virtualization is a method of managing computer systems and resources in a functional matter, regardless of where they are located (geographically). This concept of virtualization is growing in importance to all areas of Information Technology (IT). Some areas of virtualization are growing in their maturity, like storage virtualization and server virtualization; some others, for example, application virtualization and streaming as well as Operating System virtualization are just beginning to spark the technology market and revive the concept of virtualization. Taken as a whole, the above listed technologies are providing the industry with some very specific benefits in a wide variety of areas. However, virtualization, regardless of which type is used, is also presenting the IT world with some new and unique challenges, especially in the area of systems management.

This brief report introduces the reader to the concept of virtualization and attempts to define the various areas of virtualization technologies, the primary business driving forces behind virtualization, the benefits of virtualization and the challenges that virtualization brings.

## Virtualization, a Definition:

Virtualization, in computing, is the process of presenting a logical grouping or subset of computing resources so that they can be accessed by users or systems in ways which give benefits over the original configuration. This virtual view of the resources is not restricted or limited by implementation, geographic location or physical configuration of its underlying resources. This can include, making a physical resource, such as a server, an Operating System(OS) or a storage device(Storage Area Network) device appear to function as multiple logical resources; or, it can include making multiple physical resources, such as, storage devices or servers, appear as a single logical resource.

Armed with this very technical definition of virtualization, one can see that it offers a wide variety of tangible benefits to those wanting to consolidate their computing resources and at the same time save money.

## The Various Types of Virtualization:

One of the major problems with virtualization is that there are so many different types of virtualization technologies and confusion over their definitions. In this section of the report, I will define a majority of the key terms which are associated with virtualization.

## Operating System Virtualization:

Operating System (OS) virtualization is a method of creating one or more virtual instances of a guest OS either on top of a host OS or directly on top of a piece of specialized software layer called a "hypervisor ".

In either case, the host system's virtualization of other Operating Systems is accomplished by software, which is proprietary to the vendor, such as VMware Workstation and Microsoft Virtual PC, which resides between the physical hardware, CPU, memory, or chipset, and the "Guest" Operating Systems. Each host or guest OS runs its own applications independently, as if it were the only system operating on the piece of hardware.

In this host/guest environment, each instance of a guest Operating System stores a file called a "Virtual Disk" on the host system. This is a very common implementation of Machine Virtualization.

With the implementation of the Hypervisor architecture, the requirement for a host system is removed. With a hypervisor, virtual machines run on a thin layer of hardware abstraction software. That software layer, the hypervisor, addresses all hardware communications for all of the virtual systems on that particular machine. An example of this is VmWare ESX server.

A typical OS virtualization setup would be to have one physical system, complete with CPU, memory, hard drive, etc., and have a default installation of Windows XP Professional as the host OS and three instances of guest Operating Systems running, such as, Red Hat Linux 9, Windows Server 2003 and Debian Linux. The guest Operating Systems would live in their virtual file and co-exist and use the existing hardware on the machine.

Server Virtualization:

Server virtualization is where the base hardware of a system is virtualized, allowing multiple guest operating environments to run directly on top of the hardware, without requiring a complete host Operating System.
Basically, an administrator uses a software application which is used to divide up one physical server into multiple isolated virtual environments. These virtual environments are also known as Virtual Private Servers, however, they are also known as partitions, guests, instances, containers or immulations. There are three popular approaches to server virtualization: the Virtual Machine Model, the paravirtual Machine Model and virtualization at the OS level.

The Virtual Machine model is based on the host/guest concept. Each guest runs on a virtual imitation of the hardware layer. This approach allows the guest Operating System to run without modifications. It also allows the administrator to create guests which use different Operating Systems. The guest does not have any knowledge of the hosts Operating System because it is not aware that it is not running on real hardware, it is virtual hardware. It does, however, require real computing resources form the host Operating System. It makes use of the Virtual Machine monitor software that coordinates instructions to the CPU. Basically, the hypervisor validates all of the guest-issued PU instructions and manages all of the executed code which requires additional privileges. VMware Workstation and Microsoft Virtual PC are products that make use of the Virtual Machine model.

The paravirtual Machine model is based on the host/guest concept and it uses a hypervisor. In the paravirtual machine model, paravirtual machine model, the hypervisor actually modifies the guest Operating Systems code, also known as porting. Porting supports the hypervisor so that it is able to utilize privileged system calls. While this requires modifying the OS kernel, it minimizes execution overhead and increases speed. Just like virtual machines, paravirtual machines are capable of running multiple Operating Systems. The Linux Xen virtualization product from XenSource makes use of this model.

Virtualization at the Operating System level is a little different from the above two models. It is not based on the host/guest concept. In the Operating System level, the host runs a single Operating System kernel as its core and then exports Operating System functionality to each of the guests. The guest has to use the same Operating System as the host, although different distributions of the same system are allowed. This particular distributed architecture eliminates the system calls between layers, which can greatly reduce CPU usage overhead. It also requires that each partition remains strictly isolated from its neighbors, so that a failure or security break in one partition is not able to affect any of the other partitions. In this model, common binaries and libraries on the same physical machine can be shared, allowing an Operating System level virtual server to host thousands of guests at the same time. The OS virtualization offering from SWSoft, Virtuozzo, makes use of this OS model.

Application Virtualization:

To put it simply, application virtualization breaks down the barrier between the physical hardware, OS and the program which runs on top of them. To this day, a software application has been installed on the physical machine that it is intended to run on.

This application is then provided to a user from a remote location, a central server, without needing to completely install the application on the user’s local system. Unlike the traditional client/server operations, the application itself is not necessarily designed to be used my multiple users at one time and it is likely to be shared in the same way. Each user has their own, fully functional application environment, with few or no components actually being shared in the runtime environment.

Some of the advantages that application virtualization can bring are as follows:

- Faster delivery of higher quality applications
- Lower total cost of software ownership
- Reduction of the number of servers needed
- Increased system reliability
- Geometric scalability
- Better performance
- Reduced administration, deployment, monitoring, and maintenance costs

Desktop Virtualization:

Basically, desktop virtualization provides the end-user with a desktop environment which in turn allows the user access to any authorized application, regardless of where the application is physically located. This allows the end-user to have a single interface from which they can start or access their web, local, and server-based applications, without having to look through web pages, the Windows Start menu, or their Terminal Services interface. The virtual desktop might be hosted remotely on a central server, giving the end-user access only to the applications which are available remotely; or it might be hosted locally, giving the end-user access to local, as well as remote, applications.

Streaming:

Streaming is a subset of other virtualization technologies, which provides a way for software components, including applications, desktops, and complete Operating Systems, to be dynamically delivered from a central location to the end-user over the network. Unlike traditional methods of software delivery, however, the software component is usually available for use by the end-user before the entire download has completed. There is no complex and time-consuming installation process. Some of the companies that offer application, desktop and Operating System streaming are AppStream (Application streaming) and Neoware (Operating System streaming).

Some of the benefits of application streaming are:

- On-demand application delivery
- Efficient bandwidth utilization
- Self-service software distribution
- Efficient utilization of software licenses
- Simplifies application provisioning and updates
- Enables end-user roaming and free seating – users are not tied to specific machines
- Simplifies business continuity planning

## *Storage Virtualization:*

Storage virtualization is the union of multiple storage devices into what appears to be a single storage unit. Storage virtualization is often used in Storage Area Network (SAN), a high-speed sub-network of shared storage devices, and makes tasks such as archiving, back-ups, and recovery easier and faster. Storage virtualization is usually implemented in software. The users and the applications that make use of storage virtualization do not have to be concerned with where the storage is physically located.

## *Data Virtualization:*

To put it simply, data virtualization abstracts the source of individual data items, including entire files, database contents, document metadata, and messaging information. It provides for a common data access layer for different data access methods, such as SQL, XML, file access, and JDBC. This common data access layer interprets calls from any application which is using a single protocol, and translates the application request to the specific protocols which are required to store and retrieve the data from any supported data storage method. This basically allows applications to access data with a single methodology, regardless of how or where the data is located.

## *Clustering:*

Basically, a cluster is a form of virtualization which makes multiple physical systems appear to the application and end-users as a single processing resource. This is different from other virtualization technologies, which do the opposite, making a single physical system appear as multiple independent operating environments. One such case for clustering is to group a number of identical physical servers to provide distributed processing power for high-volume applications, or as a "web farm"--which is a collection of web servers which can all handle the load for a web application.

## *Grid Computing:*

Just like a cluster, a grid provides a way to abstract multiple physical servers from the application they are running on. The major difference here is that, the computing resources are spread out over a Wide Area Network, potentially across the Internet, and the physical servers which make up a grid do not have to be identical. Unlike a cluster, where each server is locally connected, and is likely to be identical, while handling the same processing requirements, a grid is made up of heterogeneous systems, in diverse locations, each of which might specialize in a particular processing capability. In order to make use of grid computing, a great deal of coordination will be needed to allocate the resources to the appropriate workloads.

## *Software as a Service:*

Software As A  Service(SaaS) is an implementation of virtualization, where the software is being provided by an external Application Service Provider(ASP), generally on a rental or subscription basis. Usually, the end-user will access the software service through a web browser and, in some cases, specialized software might need to be used. The complete software application is not hosted locally, or even within the network, but it is hosted at a third-party service provider. Microsoft plans to implement this service in late 2007. Application streaming can also be an enabling technology for SaaS, by providing access to applications on-demand through any web browser.

## *Thin Client Computing:*

A thin client is a local system which has limited or no independent processing, storage, or peripherals of its own, relying entirely on a remote system for all of its operations. This kind of computing takes its concept from the old-style mainframes and dumb terminals. Usually, a thin client will have limited local processing power which allows it to send and receive I/O to a central server, which hosts the Operating System, desktop, and applications.

## Benefits of Virtualization:

There are a number of driving forces behind virtualization, both business needs and availability needs. Some of the most significant drivers behind this technology are as follows:

- Flexibility and agility
- Server consolidation
- Business continuity and disaster recovery
- Reduction in downtime
- Reduction in administrative costs

### *Flexibility and Agility:*

Deploying virtualization technologies can greatly increase business flexibility and agility. By decoupling business processing from the physical hardware, virtualization enables IT departments to respond rapidly to growing changes in demand.

Virtualization technologies also allow businesses to quickly deploy new products and services, to offsite premises, remote offices and contract personnel. This also enables expansion into new markets. There is also a much lower hardware requirement for testing out new applications. Software developers can develop and test their code on multiple Operating Systems, which will reduce the development and testing time. They can also instantly reload their test systems from an image, which can result in a faster build/test/rebuild cycle. Virtualization can also reduce the routine deployment processes for production implementation from minutes instead of days or even weeks. It also gives the enterprise a central location for all application updates, instead of having to touch a number of dispersed systems for each new application upgrade or update.

### *Server Consolidation:*

Server consolidation and improved server utilization is yet another driving force for the adoption of virtualization technologies. Virtualization allows businesses to combine workload from multiple underutilized physical machines into a single physical system. This can greatly reduce the overall hardware spending, because it requires far fewer physical systems for the same application load. It also has a greater effect on the overhead costs, including, cooling, power, storage, and physical administration.

### *Business continuity and disaster recovery:*

Virtualization and streaming allow for easier software migration, including system backup and recovery, which can make it extremely valuable as a disaster recovery or a business continuity planning solution. Virtualization can duplicate critical servers, so that it does not need to maintain expensive physical duplicates of every piece of hardware for disaster recovery purposes. Disaster recovery systems can even run on dissimilar hardware. In addition to this, virtualization can reduce downtime for maintenance, as a virtual image can be migrated from one physical device to another to maintain availability while maintenance is performed on the original physical server. This applies to both severs and desktops, and even mobile devices. Virtualization allows end-users to remain productive and get back to work faster when their hardware fails.

### *Reduction in Downtime:*

The reduction in downtime is another key driving force behind virtualization. Virtual images are much easier to restore after a failure, either an operational failure or a hardware failure. The portability of virtual images allows new and different hardware configurations to be used for recovery purposes, thereby reducing downtime. Likewise, from an end-user perspective, desktop failures are critical, but with application virtualization and streaming, end-users are not tied to a specific failing desktop or location, and as a result, can get back to work on any machine in no time, reducing the impact of any downtime.

### *Reduction in Administrative Costs:*

With virtualization technologies, administration becomes a lot easier, faster and cost-effective. Visits to the user's place of work can almost be eliminated through application virtualization and streaming, since business applications are being maintained centrally. Any failure in the end-users environment can be fixed quickly and easily. In addition to this, virtual server Operating Systems can for the most part be managed remotely using standard tools and network interfaces, rather than needing physical attention.

## Challenges in Managing a Virtual Environment:

While virtualization offers a plethora of benefits, it also introduces some new management challenges which will have to be considered and planned for by businesses which are considering a virtual stance.

Some of the key business challenges that virtualization brings to the table is as follows:

- Bandwidth implications
- Policy management
- Image propagation
- Security

Some of the other challenges that virtualization bring to the table include: a new level of complexity for capacity planning; a lack of vendor support for applications which are running on a virtual environment; increased reliance on hardware availability; an additional layer of monitoring complexity; and an overall increase in the complexity of the IT environment.

### *Bandwidth Implications:*

Businesses will need to ensure that they have the appropriate network bandwidth for their Server virtualization requirements. For example, instead of having one server which uses 100MBPS Ethernet cable, now 10 or even 100 virtual servers will have to share the same physical connection. However, Application streaming actually minimizes load on the network since there is no need to download the entire application. Typically, only 10-15% of the application is required initially and blocks are delivered as needed after that based on usage.

### *Policy Management:*

Businesses will need to look to deploy a form of automated policy based management stance together with their virtualization strategy. For example, resource management should include some form of automated policy tools for disk allocation and usage, I/O rates, CPU utilization, memory allocation and usage, and network I/O. These management tools will need to be able to push resources in shared environments, to maintain service levels and response times which are appropriate to each virtual environment. The administrators will need to be able to set maximum limits, and allocate resources across virtual environments. Allocations will need to have the capability to change dynamically to respond to peaks and drops in load balancing characteristics.

### *Image propagation:*

OS and server virtualization can and will lead to rapid propagation of system images. This happens for the simple fact that it is much easier and faster to deploy a new virtual image than it is to deploy a new physical server, without management approval or hardware procurement. This can bring with it a high-degree of management and maintenance requirements, and potentially lead to some significant licensing issues, including higher costs and compliance issues. This propagation can also lead to a significant storage challenges, such as competing I/O and extreme fragmentation, multi-disk access, and increased maintenance time, effort, and cost. Businesses will need to manage their environment with the same level of discipline as their physical infrastructure, making use of discovery tools to detect and prevent new systems from being created without following the appropriate processes.

### *Security Considerations:*

While virtualization can bring a lot of security benefits, security also becomes a management issue in a virtualized environment. There will be more systems to secure, more points of entry, more vulnerabilities to patch, and more interconnection points to exploit, across virtual systems, as well as across physical systems. Access to the host environment becomes more critical, because it will allow access to multiple guest images and applications. Businesses will need to be diligent in securing their virtual images just as well as they secure their physical systems.

## Conclusion:

Virtualization, in all of its glory and forms, is a highly beneficial and at the same time, disruptive technology. Businesses are deploying virtualization for a number of benefits. The strongest and by far in my mind the most realistic, is the area of business continuity, aside from all of the other benefits, is the most beneficial and utilized.

The management of virtual environments, by all means, is a lot easier than it might otherwise be. In the stagnant disciplines, it is rather surprising to see how most businesses believe that it will be easier to manage a virtual environment than a physical environment. Even with all of the benefits, there are some significant challenges which need to be dealt with and resolved before a business can even think about deploying and managing virtual environments. Among these challenges, is training and staff development. Some of the other management challenges are less apparent. Policy-based resource management and capacity planning are a fairly new challenge in the virtual world, which require new technologies and processes.

Overall, it is clear to me and other IT professionals that despite these challenges, virtualization and streaming is a mature technology and it is here to stay. Businesses who are considering this technology should review all of their options and decide on the right course for them. They should also carefully plan their deployment, taking into account the potential costs, disruptions, and skill level challenges.

## About Fire Dragon Air Defense:

Fire Dragon Air Defense is a computer and network security business which offers its customers the most wide range and innovative processes from which to secure their data and computing resources. With Fire Dragon Air Defense services and products, your business can secure your computing resources and network environment through innovative and flexible network products, technologies, and services. We invite you to explore these network technologies and security services and understand how they can play a vital role in your business.  For more information about Fire Dragon Air Defense and its products and services, please contact us at 541-501-1943