# Virtualization: Systems, Applications, and Beyond

EDUCAUSE Evolving Technologies Committee
John S. Moses, The University of Chicago

October 2007

## ***What is it?***

In simple terms, *virtualization* refers to a layer of abstraction between hardware and any application or software system that seeks to use that hardware. SearchServerVirtualization.com describes it well when it says "Virtualization is the creation of a virtual (rather than actual) version of something, such as an operating system, a server, a storage device, or network resources."[1] While this concept has received significant buzz in the industry recently, it is not new. Early implementations included the Compatible Time-Sharing System (CTSS) developed by MIT on the IBM 700/7000 series in the 1960s. While still evolving, virtualization is now an established part of many Information Technology (IT) environments.

Virtualization can be applied to many areas of computing. One area receiving the most attention at present is hardware virtualization. This involves the concept of "host" and "guest" systems. There are typically several guest operating systems (example: Windows, Linux, Solaris) on a single host. This host could be running either a "hypervisor" software layer, an operating system, or an embedded virtualization layer. Vendors in the hypervisor space include VMware (ESX Server), SWsoft (OpenVZ), Citrix(XenSource), and Virtual Iron (Open source Xen). Operating system (also referred to as hosted virtualization) vendors include Novell, Red Hat, and Microsoft's upcoming Viridian. Embedded virtualization vendors include Hitachi (Virtage). One other twist on hardware virtualization is what is referred to as Operating System (OS) Virtualization. In this case, a single OS is virtualized for use – not to run other operating systems on it, but to have virtualized shells that depend on the underlying operating system. Vendors in this space include SWsoft (Virtuozzo) and Sun Solaris (using zones and containers).

Virtualization is prevalent in other areas of IT as well. For example, file virtualization[2] (the abstraction of the actual location of files from where they appear to be located), application virtualization (the abstraction of the application and data from the operating system, thereby making the application portable.[3]), storage virtualization (the abstraction of storage components from the operating system or network), and network virtualization (the existence of multiple virtual networks on top of a single physical network). It should be noted that there are several interpretations of these terms, and it would not be surprising if these differences continue to persist as this space evolves. There is also a general confusion in some quarters about the difference between emulation and virtualization. While emulation seeks to intercept and interpret each and every call, and is done entirely in software, virtualization deals with copies of the real system. As a result, emulation tends to be much slower than virtualization. Virtual PC 2004 for the Mac was one example of emulation.

---

1 http://www.searchservervirtualization.com

2 http://searchstorage.techtarget.com/generic/0,295582,sid5_gci1260549,00.html

3 http://www.byteandswitch.com/document.asp?doc_id=128177

***Why does virtualization matter?***

The concept of a single piece of hardware hosting several virtual machines is attractive from a management, flexibility, risk reduction, and cost perspective. This matters to Higher Ed where lowering costs while increasing efficiency is often difficult to achieve. What should not be lost however is the improved service availability that virtualization can provide. Below, we shall examine the impact of virtualization in Data Centers, Learning Technology, and General Computing environments.

**Data Centers:**
Virtualization is most effective in the Data Center when a multi-dimensional approach is considered, as show in the figure below. It is easy to look at the different virtualization technologies in their respective silos, and not at the big picture. In addition to hardware (server) virtualization, one needs to examine I/O, network, storage, and application virtualization.

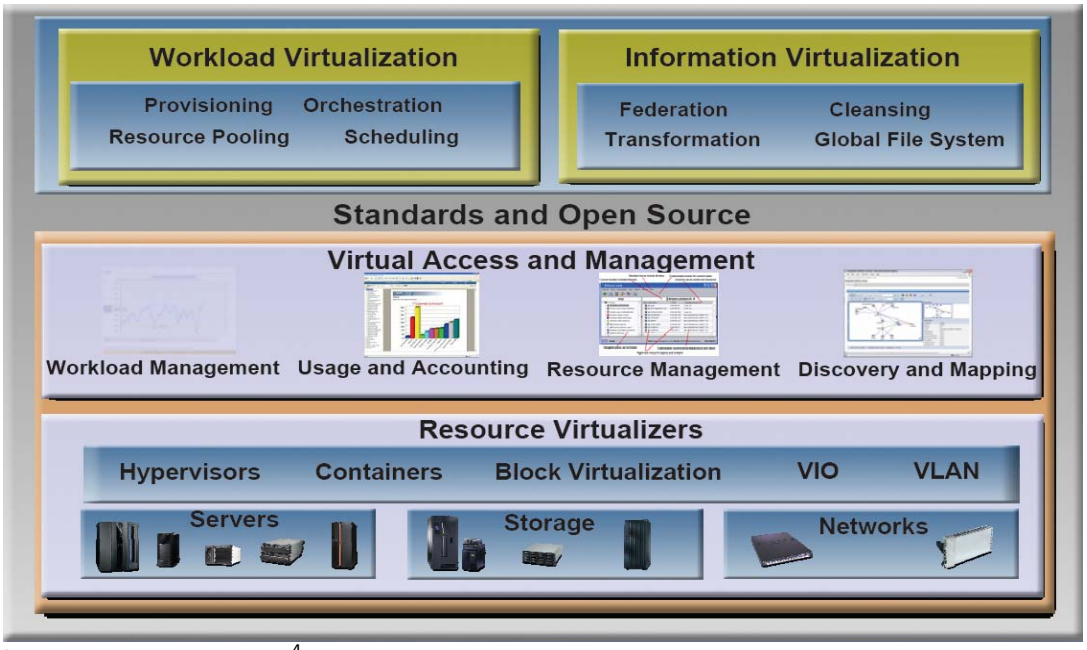

©2006 IBM Corporation.[4]

Some of the known benefits of data center virtualization are listed below.

1. Server and Service Redundancy via fail-over mechanisms and other means: Providing Business Continuity for services. Service availability is seen as a key driver for virtualization implementations.[5] To truly address redundancy of entire data centers, a unified approach including network, storage, and application virtualization is needed.

---

[4] http://www.championsg.com/Champions_InnSite_v10.nsf/ObjectsView/Symp_ETSIII_IBM/$File/IBM%20-%20Executive%20Track%20Session%20III%20Tampa.pdf

[5] Virtualization goes mainstream, Jason Compton, Smart Enterprise,Vol. 1, No. 3, 2007 (http://www.smartenterprisemag.com/articles/2007fall/researchwatch.jhtml)

2. Reduced downtime for equipment servicing: When servers need to be taken down for hardware maintenance, virtual servers can be moved to other hosts either within the same data center or in another site while maintaining service availability.
3. Reduction in the quantity of physical hardware, and thereby a reduction in the costs associated with space, power, and cooling.
4. Reduction in the time necessary to provision new services. What was once a process of ordering/racking a physical server, building the operating system, installing an application, and then deploying it into production, can now be reduced to a few mouse-clicks (when standardized templates are used).
5. Increase in the utilization of each unit of server hardware. Taking a few servers that had utilization figures in the single-digit percentages, and optimizing their use with a virtualization approach.
6. Standardization of server builds and reducing errors in the server build process for new servers. This may not seem as clear in the beginning, but the old adage of "garbage in, garbage out" is apt for virtualized environments. If your existing processes and procedures need help, it pays to fix them before embarking on virtualization.
7. Testing of services was traditionally tied to the setup of a physical machine. Virtualization has forced this thinking to change. According to Chad J. Kainz of the University of Chicago, "Virtualization allows one to take a snapshot of the entire environment (multiple servers) and do an in-situ style test without impacting production. One could have full rollbacks of a whole environment, not just a server." This is a fundamental change in the way we view and operate test environments.

**Learning Technology Environments:**
Learning environments could use both hardware (desktop/server) virtualization and application virtualization effectively, in addition to leveraging any existing virtualization technologies within the data center. A framework in use within Academic Technologies at the University of Chicago has done just that. An initiative code-named Scholarly Projects Infrastructure Initiative (SPII) aims to provide much needed infrastructure for piloting projects, hosting emerging applications, and sharing data in support of research and education at the University of Chicago[6]. It utilizes open-source and commercial software in a utility model. Desktop, Server, and Application virtualization are core components of the implementation, as are data center storage and networking.

Whether in Media classrooms or Public Computing facilities, whether it be applications that don't behave with one another, or complete operating system environments with custom applications to support a particular course , virtualization technologies have a role to play. A framework that utilizes virtualization to provision applications and systems for faculty and students is an asset. In addition to the fact that it allows the speedy provisioning of services, it also reduces the cost of support of the desktops traditionally used to run these applications. There are challenges for sure, such as how high-end video editing is accommodated, but this should not detract from the promise. The flexibility, manageability, and cost savings that can be gained from the use of these tools is worth the time and investment necessary to set it up.

**General Computing:**

[6] Scholarly Projects Infrastructure Initiative, Academic Technologies, The University of Chicago, Chad J. Kainz.

Administrative desktops can be virtualized – either at the desktop level, or at the application/file level. Applications and data can be virtualized into a wrapper, then delivered to the desktop via application streaming technologies. Microsoft SoftGrid and Citrix Presentation Server are two such streaming technologies in wide use. One of the reasons for virtualizing applications is to manage multiple versions of the same application – Microsoft Access for example. In a broader sense, an application virtualization strategy could be employed for the deployment and management of all applications. Nirvana in this space is the concept of having an application available to any person on any device in any location on any storage. Sounds like a tall order, but it is no longer a dream.

## ***How is virtualization evolving?***

Certain components of virtualization have become mainstream, while others continue to evolve. Certain components that were available for a fee are now available free-of-charge – e.g. VMware GSX server. This is a reflection of the changes in the marketplace – commoditization of mature technologies, and differentiation in newer areas. We'll take a look at some of the changes below.

*Native Virtualization*: Operating System vendors are beginning to incorporate virtualization into their products natively. Sun Solaris v10 ships with native sandbox-style virtualization. RedHat Fedora ships with Xen.

*Embedded Hypervisors*: Hardware vendors (Dell, HP, IBM) are backing embedded hypervisors. The "operating system" that the servers ship with is the hypervisor, either VMware ESX Server 3i or XenSource XenExpress OEM Edition. Systems with embedded hypervisors don't have a hard drive, and generally run off a flash drive that contains nothing but the hypervisor. This could lead to a trend of purchasing systems pre-built with VM hypervisors direct from the factory and doing away with operating system installs on bare-metal once equipment is received at customer sites. Plug and play simplicity seems to be the trend here. VMware is readying a new VMware ESX Lite embedded hypervisor for use in server blades. In addition, chips that are specifically designed for virtualization are here in the form of Intel with Intel VT, and AMD with AMD-V chips.

*OS-to-hardware virtualization*, where the hypervisor is loaded directly onto the hardware as an operating system, thereby removing the need for an underlying operating system like Windows, Linux, or Solaris. According to Gartner[7], "The number of PCs using OS-to-hardware virtualization technology ("virtualized PCs") will more than double during 2007."

*Virtualization vendors offering management tools*: According to the research firm Gartner[8], VMware will need to become more of a management company. VMware manages multiple guests across multiple nodes of hardware. Integration with existing management tools will be embedded in the basic virtualization solution. In some sense, this is already happening. The latest version of VMware's ESX Server includes additional tools and integration with network,

[7] Gartner Research, Forecast for PC Virtualization, 8 August 2007, Brian Gammage, George Shiffler III

[8] Gartner Research, Why VMware Must Morph Into a Management Company, 31 August 2007, Cameron Haight

storage, and clustering environments. This begins to encroach on the traditional systems management space that is dominated by Symantec, IBM, BMC, CA, HP, and others.

*Vendor consolidation* may drive changes in the product landscape. Citrix has recently acquired XenSource. Although currently dominated by VMware, the hardware virtualization market will begin to coalesce around VMware, Citrix, and Microsoft. There is a strong sense that virtualization will drive us towards a model where large centrally run groups of computing devices provide distributed and customized services to a diverse user community. In other words, flexible centralization through the leveraging of the network, storage, systems, and eventually applications.

## ***Challenges and Opportunities:***

Implementing virtualization poses a few challenges that are worth examining. Some of these are standard challenges that need to be addressed with the arrival of disruptive technologies. Others are more focused on virtualization.

*The right scope, the right fit*: How much are you willing to virtualize? There are several choices when it comes to selecting a virtualization solution. Which solution is right for you? Vendors like VMware and Platespin[9] offer assessment tools to help. Most vendors in this space (ex: IBM, HP, etc.) offer either minimal virtualization assessments or complete data center optimization assessments. Even if the organization has a good idea as to what can and cannot be virtualized, an independent assessment would help verify internal findings, and would help towards developing appropriate funding models for future years.

*Internal Processes and Procedures*: Developing strong processes and procedures to ensure that the VM environment is manageable and sustainable is a pre-requisite to embarking on a virtualization initiative. Poor controls will wreak havoc on a new virtualization environment. Time needs to be spent to carefully take apart existing processes and adapt them for the new, or replace them altogether. If your data center has poor processes and procedures already, virtualization will make it worse.

*Server sprawl*: The traditional paradigm of setting up servers involved the procurement and setup of physical servers, the installation and configuration of an operating system, the installation and configuration of an application, and then the deployment of the service. In the virtualized data center, a few mouse clicks will accomplish this – provided of course that your environment has the necessary capacity to add new virtual servers as needed. A recent survey by the research firm Gartner stated "Virtual machine (VM) "sprawl" is a major concern that appears to lack a comprehensive solution."[10] In the survey, VM sprawl was listed the number one concern among managers of virtualized environments.

*VM management*: There are two challenges here, as listed below.

---

9 http://www.platespin.com/

10 Gartner Research, Survey Identifies Server Virtualization Needs and Trends, 29 August 2007, Cameron Haight

1. The management of virtual machines is in need of better tools. Virtualized environments have a threshold for either processors or memory. The number of virtual servers has to be tracked against the available processors and memory in that system.
2. A common interface is lacking to manage different types of virtualized environments – such as Xen and VMware.

*Open Virtual Machine Format:* Does the particular virtualization technology that you are considering have a roadmap for Open Virtual Machine Format (OVF) compatibility? If not, consider the consequences of investing in a particular technology that may or may not allow a transition to a different platform if and when the need arose.

*Application vendor restrictions:* Are there specific applications that will not run in a Virtual Machine (VM) environment, or that the application vendor refuses to support in a VM environment? If such a restriction exists, then consider placing test systems in a virtualized environment.

*Security:* In addition to taking care of security on operating systems such as Windows or Linux, we now have hypervisors to contend with.

*Software Licensing Models:* In response to the growing reality of virtualized IT infrastructures, application licensing has begun to evolve. For example, depending on the version of Windows Server 2003 that is used, Microsoft allows either unlimited virtual machines, or up to four. If using OS virtualization, as in SWsoft's Virtuozzo, there may not be the need to purchase OS licenses for each virtual machine instance – at least for now. The problem is that vendor response to virtualization is a moving target, and remains one of the more difficult expenses to plan for. Overall, there is the potential for a cost reduction in this space, especially when items such as Backup software licenses are considered.

While there are many challenges posed by virtualization, it can be deployed effectively now, and it can be done either organically within the organization, or quickly using external partners. In either case, there are benefits to be had.

## ***What does the future hold?***

Academic institutions are often faced with random demands for new systems and storage. The traditional approach to meeting these needs would be via the purchase and setup of new workstations or servers within laboratories or data centers. A virtualization strategy allows for a different model to meet these needs. For example, a virtualized data center would provide a scalable and cost-effective platform from which to provision new services to customers in addition to increasing service availability and recovery. These environments offer the flexibility to provision new services with a few mouse-clicks – from pre-defined templates. The human effort involved in purchasing, racking, installing, and configuring new systems would be replaced with a process that utilizes pre-defined templates and a few mouse clicks.

The increased efficiency offered by virtualized environments should be welcomed. Institutions looking for new models to address application provisioning and hosting, data storage, and data backup/archiving should consider a virtualization strategy. A virtualization strategy that takes

into account the collective needs of the institution may greatly benefit the community as higher education transitions to a new way of allocating and managing IT resources.