Virtualization
What is it and why is useful?

Qureshi Mohammad Umer
University of East London

Abstract:
With all the advancement to consume resources, all sorts of methodologies are being adopted by big companies like Intel and AMD that are investing millions of dollars to come up with software that resides on their machines to bring out the most of them.
Software giants like Microsoft took over Connectix Corporation to develop virtualization software for Windows and Macintosh systems. This led to EMC acquiring VM Ware whereas VERITAS acquired Ejascent all to develop Virtualization software.
IBM being pioneer in virtualization thought of this in 60's by bringing out their time sharing system concepts in M44 44X machine.
If hardware companies make wise use of virtualization they can save quite a lot on the research and development of the machines they intend to bring out in future.
With a widespread of this technology being adopted small servers are also relying on it.

This report is being presented after a research conducted by the author to explain what virtualization is and why is it useful.

Introduction

Many papers and reviews have been published to bring about awareness of what virtualization is and what makes it so useful (sources for example Kernelthread.com, Softricity, Ideas International Inc. and Network world).

Something that can be considered as an endeavor to start off with what we now call virtualization began with Manchester University and Ferranti Ltd.'s joint venture on the Atlas project. It used one virtual machine to handle supervisor calls activated either by interrupt routines or by extra code instructions and the other to run user programs.
With further information gained from kernelthread.com it came to my knowledge that IBM then developed IBM 7044 (M44) that had an address space of 44X residing in the M44's memory hierarchy implemented via virtual memory and multiprogramming.
The MAC project by MIT later did not require IBM even with their Compatible Time Sharing System (CTSS) as they required a better time sharing system based on the CTSS. Even though IBM is still considered to be the best today with respect to robust computing platforms as they

came up a virtual machine that had a component called the VMM Virtual Machine Monitor that can run multiple virtual machines having their own operating systems.

What is Virtualization?

"Virtualization is a framework or methodology of dividing the resources of a computer into multiple execution environments, by applying one or more concepts or technologies such as hardware and software partitioning, time-sharing, partial or complete machine simulation, emulation, quality of service, and many others."
----http://www.kernelthread.com/publications/virtualization/ ***(1)***

To obtain an environment that can support virtualization the machine has to be partitioned so that multiple operating systems are concurrently executed.
This however requires three considerations as pointed out by the Systems Research Group in University of Cambridge's Computer Laboratory ***(2)***.

- The virtual machines should be isolated from each other in order to avoid affect of one execution over the performance of another.
- A variety of different operating systems should be supported to accommodate the heterogeneity of popular applications.
- The performance overhead introduced by virtualization should be small.

Types of Virtualization

Application Virtualization:

Softricity Technologies ***(3)*** having work done in application virtualization came up System Guard a virtual application environment. They claim application virtualization outstands machine virtualization
Application virtualization takes the concept of managing and simultaneously operating multiple environments on a single machine.
It acts an abstraction layer lying between the operating system and the applications running within by virtualizing all the features of an application, yet not affecting the operating system or other applications running on that machine.
It does this by creating a virtual IT environment where all computing resources can be dynamically allocated in real-time based on real-time needs.
Applications using this software execute with full performance, functionality and access to local services
This reduces the total cost of deploying and maintaining applications and systems.

However Xen Source ***(4)*** does not consider it as a contender in the data center as application virtualization loads all processes, threads, and application related state for each application in a

virtual area to provide balanced performance to all virtual areas by using the operating systems scheduler.

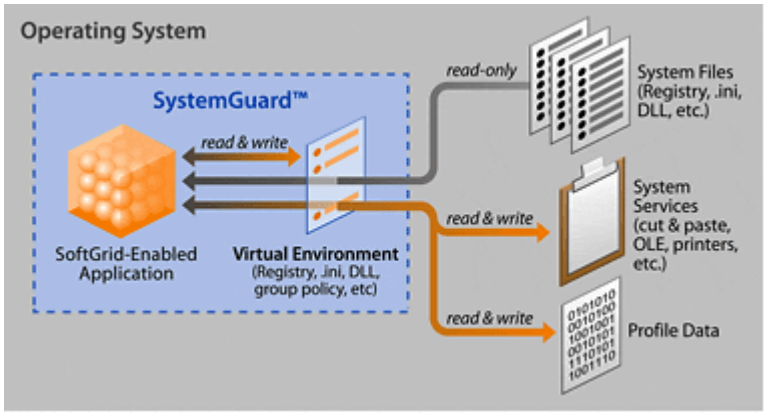

Application Virtualization ***(3)***

Using Hybrid:
Many customers, such as Fidelity National Financial, use both Softricity and machine virtualization to maximize server and IT management efficiencies

OS Virtualization:

As mentioned in the white paper by Xen Source ***(4)*** OS virtualization can be described as a layer of software (the hypervisor) that allows multiple OS instances along with their running applications to share resources from a single server. The hypervisor is embedded between the operating system and the sever hardware. This virtualization software ensures that all resources are transparently provided to all operating systems whereas each operating system has an illusion that it has all the resources from the server.
The way this is achieved is that the hypervisor manages all hardware structures that is; the page tables, the I/O devices, DMA controllers and ensure that the underlying hardware is visible by each OS that is running. It also preserves any state that the currently running OS will expect to be in place in the hardware if its execution is later resumed. Loss in performance may still occur when data structures such as Translation Lookaside Buffer (TLB) are invalidated on a context switch. This causes the incoming OS image to fault on each page reference until the TLB is refreshed. The existing operating system binary is therefore required to be re-written ensuring that the hypervisor regains control on the system from time to time. This type virtualization is known as **Emulated Virtualization.**
An understanding that can further be gained of what OS virtualization is can be derived from a review by Linux format's June 2005 issue. ***(5)***
It explains that privileges can be classified into four levels or 'rings', with ring 0 having the most privileges and ring 3 having the least.

The Table below shows how a normal OS runs directly on hardware in ring 0

3. Applications 2. 1.	Applications run in the least privileged ring, which stops them running amok.
0. Host OS	The host OS runs in the most privileged ring, where it controls programs running in other rings

The following table shows how a normal OS runs with a virtualized guest on top which is considered to an ideal scenario for now.

3. Host OS Applications Guest OS Applications 2. 1. Guest OS	Host and Guest OS applications all run in ring 3, unprivileged. Guest OS's are allocated to ring 1, making them more privileged than applications but less than the host
0. Host OS	The host OS continues to run in ring 0

Time Slicing of CPU time is applied by the parent system and each ring is allocated sufficient time for its process with out tampering processes in other rings.

Para Virtualization:

This technique will be discussed in accordance to Xen the Para-virtualization software ***(4)***.
All hardware interfaces are completely virtualized by the guest OS which is ported to an idealized hardware layer. Calls into the API that are offered by the hypervisor are made when hardware data structures are updated by the OS. The hypervisor can then decide how to modify the hardware on any context as well as it keeps track of any change made by the OS. The hypervisor has full access over the OS and can enable it to be virtualized.
As claimed by the Xen Source Para Virtualization has benefits over other techniques in terms of device drivers and device interfaces. It divides the OS drivers into two halves running one as separate domain with memory, CPU and other resources. As drivers are run as separate protection domains from the core hypervisor, it is protected from bugs and crashes in device drivers thus enabling its hypervisor to be faster than competitors.
Two achieve fast virtualization one of the two approaches adopted by Xen hypervisor:

1) The hypervisor makes the guest OS aware that it does not have full access and control of the CPU
2) It provides hardware based CPU support for I/O virtualization and multiple guest OSes.

Image below taken from Xen Source White Paper ***(4)***

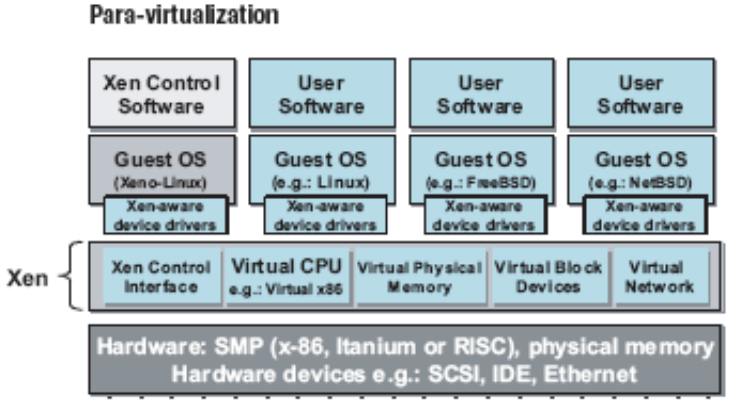

The guest OS is ported to the Xen virtual hardware interface. All guest OS modifications of hardware data structures are performed via the API. The hypervisor is mapped into the guest OS address space, avoiding a TLB flush on a context switch into the hypervisor. Guest OSes are optimimzed for virtualization.

Some Virtualization Software developed:

(1)

- **Wine:** enables Windows to be run on Linux, FreeBSD and Solaris. It does not emulate a processor and is X86 only.

- **Disco:** a virtual memory manager that was implemented as multithreaded shared memory program allowed multiple virtual machines was developed in the Stanford University.

- **HP-UX Virtual Partitions:** developed by Hewlett-Packard runs applications in isolation from the operating system.

- **LPAR:** Logical Partitioning was developed by IBM that allowed multiple independent operating system images of AIX and Linux to run on a single server.

- **Mac-on-Linux:** most virtualization functionality is implemented as a kernel module. Runs under Linux on most PowerPC hardware.

- **Microsoft Virtual Server:** Includes: Microsoft SQL Server 2000, Microsoft Exchange Server, File Print Servers, Terminal Server and e.t.c.

- **Programming Language Virtual Machines include:** Java Virtual Machine (JVM), Microsoft .NET CLI and Parrot.

- **Solaris:** Introduced by Sun Technologies includes virtualized components of: system calls, NFS, IP address space, syslog facility and overall system management layer.

- **UMLinux:** framework made to evaluate the behavior of networked Linux machines in the presence of faults.

- **VMware:** uses the principles of dynamic, logical partitioning by enabling physical computers to be available for operating systems to be run as a pool of secure virtual servers.

Why is Virtualization Useful?

A number of reasons were found out that do make virtualization very useful. Sources include reviews mainly from kernelthread.com ***(1)*** and Technology Trends ***(6)***.

The following are reasons which make virtualization useful:

- With only few machines the work load of several under utilized servers can be consolidated, benefiting in terms of monitorial savings of hardware, management and hardware.
- Secure Computing platforms are provided. Harmful applications could be separated from the OS and processed in a different ring or a container depending on the type of virtualization technique being used. It also enables existing operating systems to run on shared memory multiprocessors.
- Resource limits can be set that give advantage of setting execution environment with allocated time slices, memory and other resources.
- Simulated networks of independent computers can be created as virtualization can projects an illusion of hardware or hardware configuration that is not present.
- Some old files require formats that are not currently used but with virtualization different operating systems can be run even different versions.
- Virtual machine monitors can debug operating systems without losing productivity.
- Faults and errors can be contained as virtual machines can isolate what they run. Entire Research work can also be isolated, saved separately and later loaded to run.
- With hypervisor working virtualization can be made effective to provide binary compatibility.
- Server utilization is being increased with virtualization as multiple applications can be run without each intervening with another.
- Multiple virtual computers can be hosted on a single server as each application runs as if it was on a real server and on its own operating system.

Conclusion

This report is based on different views searched from different sources that have vital role in the evolution of virtualization. As a conclusion virtualization can be considered as resource partitioning where every process be that of an operating system or of an application is made to believe that it is of the utmost importance.
This report presented different perspective of professionals working in different types of virtualization fields. Xen Source used the words so called for both emulated and application virtualization and in some way there software may be better but that can be found with respect to industrial application and implementation. I believe Xen is an improved version of Hewlett-Packards HP-UX and IBM LPAR.
As mentioned earlier a specific technique can not outcast another virtualization technique as it is a matter of requirement and therefore a solution may become a hybrid technique.

In the end I would Quote Anne Skamarock of Enterprise Management Associates in Boulder, Colo.*(7)*

"At this point, virtualization may seem like smoke and mirrors, but it will be an important technique for solving many of the complexity issues we are faced with in managing storage, especially as storage networks become prevalent"

References:

(1) Introduction to Virtualization by Amit Singh
http://www.kernelthread.com/publications/virtualization

(2) Xen virtual machine monitor, available online at
http://www.cl.cam.ac.uk/Research/SRG/netos/xen/

(3) Softricity-- Application Virtualization technology
http://www.softricity.com/products/virtualization.asp

(4) Xen – the Art of Virtualization, White Paper available online at
http://www.xensource.com/files/xensource_wp.pdf

(5) LINUX Format magazine; June 2005 issue:
Article: "We are on the brink of computing revolution" by Paul Hudson

(6) Technology Trends—Operating System virtualization simplifies server resource management:
http://www.swsoft.com/r/upload/Ideas_International.pdf

(7) Network World—Virtualization and storage:
http://www.networkworld.com/newsletters/stor/2000/1106stor1.html

Bibliography:

- Linux Kernel Development, 2nd edition by Robert Lowe (Library Book Shelf Number: 005.432 LOV)
- Understanding the Linux Virtual Memory Manager by Mel Gorman (Library Book Shelf Number: 005.432 Gor)
- Essential System Administration by AELeen Frisch, 3rd Edition (Library Book Shelf Number: 005.432 FRI)
- Operating Systems: Internals and design principles by William Stallings, 3rd Edition (Library Book Shelf Number: 005.432 UNI STA)
- Introduction to Virtualization by Amit Singh
  http://www.kernelthread.com/publications/virtualization
- Xen virtual machine monitor, available online at
  http://www.cl.cam.ac.uk/Research/SRG/netos/xen/
- Softricity-- Application Virtualization technology
  http://www.softricity.com/products/virtualization.asp
- Xen – the Art of Virtualization, White Paper available online at
  http://www.xensource.com/files/xensource_wp.pdf
- Network World—Virtualization and storage:
  http://www.networkworld.com/newsletters/stor/2000/1106stor1.html
- Technology Trends—Operating System virtualization simplifies server resource management:
  http://www.swsoft.com/r/upload/Ideas_International.pdf
- LINUX Format magazine; June 2005 issue:
  Article: “We are on the brink of computing revolution” by Paul Hudson