

White Paper



Virtualization Technologies

Hariprasad Nellitheertha

Abstract

Over-provisioning of hardware has led to high TCO and low ROI on IT infrastructure. Virtualization provides an efficient solution to address these issues. It is a technique to transform hardware resources viz. processors, storage, I/O and networks on one or more machines into multiple execution environments, each of which can act as a complete system by itself [2]. This is done by hardware or software partitioning, time sharing and simulation/emulation of target machine. Virtualization brings in benefits such as isolation, security, increased resource utilization and simplified management. Various implementations of system virtualization technology are available. Commercial implementations of these technologies are now available for enterprises to adopt and benefit.

Introduction

Enterprises, today, have over-provisioned hardware in their datacenters to accommodate surges in utilization and to ensure application isolation using dedicated servers. This has increased the Total Cost of Ownership (TCO) of applications and reduced ROI on IT infrastructure. Virtualization is an efficient means of addressing these problems. It is a technique through which hardware resources viz. processors, storage, I/O and networks on one or more machines can be transformed by hardware or software partitioning, time sharing and simulation / emulation of target machines into multiple execution environments, each of which can act as a complete system by itself [2]. Commercial implementations of virtualization technologies are now available.

Business Case

A typical datacenter hosts each business application on a separate server. It is required, in the view of the Enterprise IT Management, to isolate the application execution domain to reduce the IT infrastructure's vulnerability to software faults and security threats arising from a compromised application. More often than not, to achieve higher availability applications are deployed on a 2- way fail-over setup. This results in substantial over provisioning and poor utilization. From an economic standpoint, what is required is for systems to expand on demand and shrink when the demand is low to keep system utilization at an optimal level.

Server Consolidation

Enterprise Applications are isolated for security reasons into individual servers. If a fault or security vulnerability is known to exist in a particular application, it becomes easy to compromise the entire set of applications co-hosted on a physical server. To avoid this situation, each application is typically deployed on a separate server. This requires a dedicated server for each application hosted in the datacenter. As the number of servers increase in a data center, the operational complexities increase exponentially. Change management, patch updates, troubleshooting and other management activities become more complex with a large population of servers. High overhead costs are un-avoidable in such circumstances. Virtualization in all these cases can be a natural and automatic solution. Through virtualization it is possible to create multiple execution environments (Operating System instances) which are isolated from each other. This sand-boxed Virtual Machine (VM) can now host enterprise applications. It is possible to create several VMs on a single, target server and each of these VMs can host a single application, providing the necessary isolation. This improves utilization and drastically reduces the number of servers required, thus cutting the cost and effort of maintaining multiple servers.

On Demand Computing

Enterprises provide for excess capacity in terms of computing power, storage and bandwidth to accommodate peak usage with guaranteed Quality of Service (QoS). Virtualization promises to change this situation by allowing infrastructure to scale instantly to accommodate fluctuating demand. This is accomplished by bringing up several application instances on isolated VMs to handle surges in demand. This can happen seamlessly in matter of minutes, if not seconds. Virtualization provides a means through which a pool of servers can be shared among several application instances in their own execution environment, substantially improving utilization and reducing costs. An On Demand computing environment built around resource virtualization can help enterprises to automatically scale up or down application instances based on policies and demand. Virtualization can help in migrating applications from one physical server to another seamlessly. This can be quite helpful in the face of hardware faults. All these can help change the $2 * N$ fail over model for high availability to an $N+1$ adaptive model [17]. (N Servers need not be duplicated in an active/passive fail over configuration, instead N servers can have a single server on standby, which can take the place of the failed server inside a VM).

Impact of Virtualization in other areas

Other areas where virtualization has a significant role to play [2]

- Resource virtualization helps create a single system image in a distributed system of heterogeneous machines such as a grid computing platform.
- Legacy applications which need to run on older platforms/OS can be run on a virtual machine on a newer platform.
- Virtualization provides a means to run multiple operating systems side by side on a single machine.
- Job migration becomes simpler in the virtualized scenario.

- Virtualization enables the use of off-the-shelf operating systems and applications to run on newer hardware platforms.

There are, however, numerous challenges in providing a virtualization layer (Virtual Machine Monitor – VMM) over a hardware platform. Most servers, today, are based on IA-32 (Intel x86 32 Bit Architecture). IA-32, as it exists today, is not fully virtualizable.

Virtualization in Utility/Grid Computing

In the utility / grid computing scenario, virtualization can play a key role in provisioning and ensuring that resource allocation is fair and the grid job does not overwhelm the client/worker machine which is volunteering cycles/resources. Also the sand-box, the VM can create, protects the grid worker machine from being harmed by malicious code or faulty code in the grid job. Virtualization can help create a homogenous virtual layer on top of a heterogeneous grid platform. Running the grid job on the VM enables the grid middleware to account / meter the usage of resources by the grid job on a particular client machine. This accounting and metering aspect is very critical to the success and adoptability of the grid platform.

Storage Virtualization

The central idea of virtualizing storage [16] is to combine several smaller storage devices with various attributes (performance, availability, capacity and cost/capacity) and present them as one or more virtual storage devices with better performance, availability, capacity and cost/capacity properties. This enhances the manageability of storage and better sharing of storage within an enterprise. Storage virtualization like system virtualization (discussed elsewhere in the document) can also greatly help businesses reduce cost.

Virtualization Challenges

As per Popek and Goldberg, the requirements for a virtualizable architecture [1] are as follows: For any conventional third generation computer, a virtual machine monitor can be constructed if the set of sensitive instructions is a subset of the set of privileged instructions. Formally, the virtual machine monitor (VMM) should exhibit the following three properties [1].

1. Efficiency Property: Provide the ability to execute innocuous instructions directly on hardware by-passing VMM
2. Resource Control Property: The VMM should be in complete control of the system. When the operating systems (running on top of VMM) try to access resources, the access should be routed through the VMM.
3. Equivalence Property: Any program running on top of VMM should perform in manner indistinguishable from the case when the VMM doesn't exist.

Unfortunately, IA-32 wasn't designed for virtualization and doesn't meet the criteria specified by Popek and Goldberg's rule. There are 17 privileged instructions in the Instruction Set Architecture (ISA) of IA-32 that don't cause a trap when executed by code executing in un-privileged mode. This prevents the VMM from gaining control, which is a pre-requisite for multiple operating system/execution environments to co-exist on a single machine. However, intelligent software hacks have been used to virtualize the IA-32 class machines/servers.

Full System Virtualization is a technique through which the target hardware is emulated in full by directly executing some instructions on the hardware and some through the VMM. The advantage of this technique is that the guest operating system (that runs on VMM) or the applications that are executed on the guest operating system need not be modified.

Para Virtualization is a technique through which the operating system is modified to avoid using the instructions that are not fully virtualizable and instead go through the VMM to execute those. Both the host OS and the guest OS (in some instances) must be modified and recompiled for para-virtualization. However, applications need not be modified.

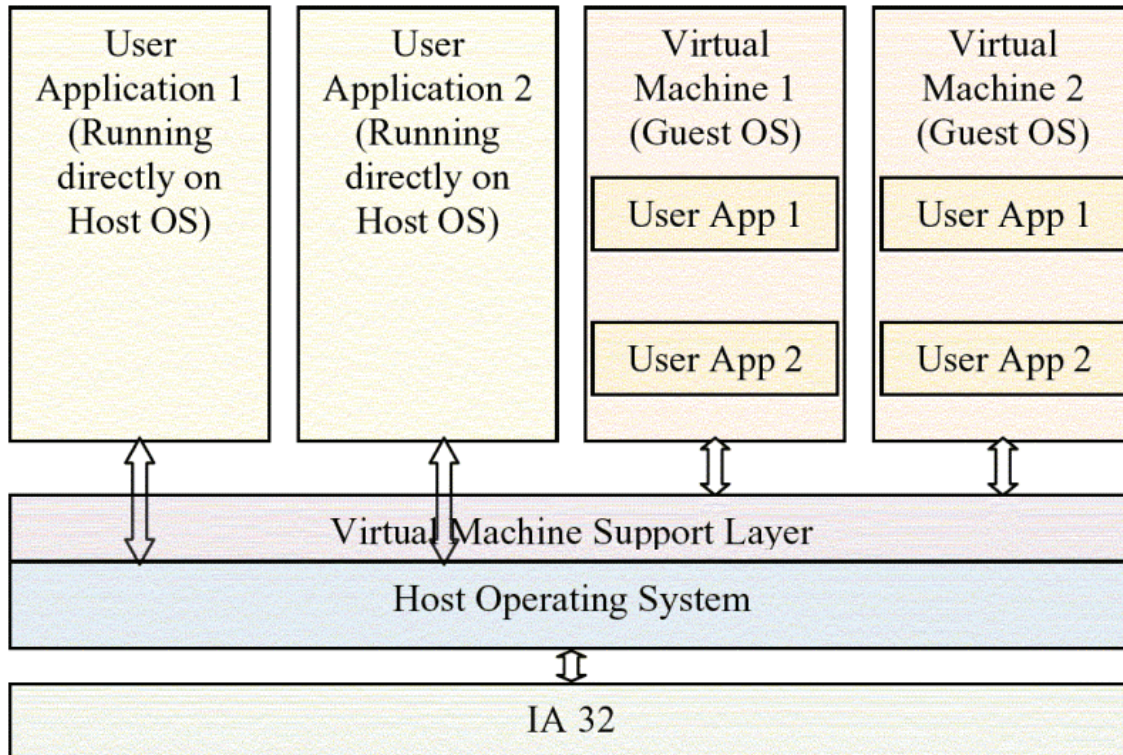
Hosted Virtual Machines is a method in which the VMM and guest OS run in the user space of the machine and the Host OS provides the services required by VMM. The applications running on the host OS and the guest OS share the same user space. This doesn't require any modification to the guest Operating System or the user applications; however the additional level of indirection normally causes heavy performance overheads. Unless optimized, the performance may degrade significantly.

Most commercial solutions on IA-32 are based on either para-virtualization systems or hosted virtual machine architectures. Some examples of hosted virtual machine are VMWare and Virtuozzo by SWSOft. Para-virtualization examples include Xen, Denali, Plex86 and UMLinux.

Virtualization Solutions

Hosted Virtual Machine Architectures

Though many hosted virtual machine solutions exist commercially, the most notable are VMWare from EMC, Virtual Server from Microsoft and Virtuozzo from SW Soft.



VMWare [9][12][14] has released virtualization solutions namely the VMWare Workstation, MWare ESX Server, VMWare V GSX Server and the VMWare Virtual Center. While the VMWare Workstation allows an individual workstation (PC) to be virtualized into multiple environments each executing different operating system (Windows & Linux), the VMWare ESX Server, GSX Server and Virtual Center concentrate on the server market. VMWare's server class solutions can be used for server consolidation and they can complement off-the-shelf blade servers quite well in further reducing the TCO. VMWare solutions are available on both Windows and Linux. The issue, however, with the VMWare server lies in the hosted model, because of which the performance of the applications hosted in the guest OS suffer. Some optimizations have been done on the server class solutions to boost performance and the results have been encouraging. The VMWare ESX Server runs directly on the hardware without a host OS providing low overhead environments for the guest OS.

Microsoft offers virtualization solutions similar to VMWare. Virtual PC is the solution targeting the workstations, while the MS Virtual Server 2005 [13] [15] is a product targeting server class machines running Windows Server 2003. The Virtual Server 2005 solution is available only for the Windows host operating system. It can support most x86 operating systems for the guest OS (unmodified) on the VMs. The Virtual Server 2005 is more optimized to run Windows based guest OS compared to other operating systems. The Virtual Server 2005 product is focused on solving issues like server consolidation, legacy application re-hosting and for simulating test environments for development and testing both stand alone and distributed software applications.

Virtuozzo [11] from SW Soft is another leading virtualization solution targeting the enterprise in the server space. Just as in the case of VMWare and MS Virtual Server, the Virtuozzo is also a hosted model. Virtuozzo is available for Windows as well as Linux. One key difference with Virtuozzo is that on a Virtuozzo for Windows the guest OS should also be Windows and similarly all the guest OS for Virtuozzo for Linux should be Linux. Virtuozzo offers two different kinds of solutions, one targeting enterprises and the other targeting hosting providers. For the hosting providers, it offers many models of shared and dedicated use of VMs across customers and applications. For the enterprise, the focus is primarily on server consolidation, process migration and QoS guarantees for applications on their Virtual Private Servers (VPS).

Para Virtualization Solutions

Solutions based on para-virtualization are not as commercially popular as the hosted virtualization model. The primary reason being that the para-virtualization solution requires modification of the kernel of both the host and the guest OS (in some cases). Several research projects taken up both by academia and startups/enterprises are in a fairly mature state and are ready for early-adoption. Notable projects/solutions in this spectrum are Xen from Cambridge University Research and XenSource, Denali from University of Washington, Plex86 and UMLinux from the open source community and Disco from Stanford. Since the para-virtualization solution requires modification of the OS for porting, there are difficulties with porting proprietary OS like Windows.

The Xen [4] project has made considerable in-roads into this technology and currently ports of Xen host OS are available for Linux. And work is in progress for porting NetBSD and Windows XP. Stable Xen ports are available for versions 2.4 & 2.6 of Linux. Performance benchmarking Xen against other virtualization solution (including the hosted model) have shown that Xen VMs performance is almost similar to that of bare Linux and the overhead is very minimal and the solution scales well over a large number of VM on a single server. If enterprises datacenters run Linux, then Xen may be a viable option. Xen supports unmodified Linux ports of Red Hat, SuSE and Mandrake for the guest OS.

Denali [3], though similar to Xen in the underlying technology, doesn't truly emulate the underlying ISA (IA-32) in full after virtualization. It is therefore required for the guest OS to be ported to run over the virtualization layer. Denali has consciously taken the approach of not fully emulating the underlying ISA to achieve higher performance and scalability. No popular commodity OS (Linux, Windows) is known to have been ported for Denali.

User Mode Linux (UMLinux) [10] is another research project which primarily targets developers and kernel hackers to provide a platform for working with an isolated sand-box for experimentation, staging and debugging. This solution wasn't intended for commercial adoption for enterprise class servers as the performance of applications running in the VM is considerably lower than the bare Linux version. UMLinux ports are not available for other commodity OS.

Plex86 [6] and Disco [5] are other similar work in this area and have similar advantages and limitations to those listed in this category.

Other than the ones listed against the hosted VMs and para-virtualized VMs, there are other notable solutions. One such solution that warrants attention is the Intel Vanderpool technology [7] which aims at addressing the non-virtualizable aspect of the IA-32 architecture to make the system fully virtualizable while retaining full backward compatibility. AMD's Pacifica is similar work targeting AMD processors also based on IA-32. These technologies are yet to be made available commercially and these would extent revolutionize the virtualization solutions and products, by facilitating highly scalable and low overhead VMs for IA-32 architecture. The VM solution providers have taken cognizance of this fact and are working towards making products available on their new processor versions.

A recent introduction to the virtualization market is Virtual Iron VFe [8] from Virtual Iron. Virtual Iron VFe runs on Linux and promises to completely virtualize resources (processor, I/O and storage) across machine boundaries and create as many virtual machines from them as the enterprise may find suitable. It also gives an option to add and remove resource dynamically from a virtual machine.

Survey of Virtualization Technologies

Hosted Virtual Machine

		VMWare Servers (EMC)	Virtuozzo (SW Soft)	Virtual Server 2005 (Microsoft)
Support Environment	Targeting Host OS	Windows Server Linux Server	Windows Server Linux Server	Windows Server
	Guest OS Supported	Windows Server Family Linux Server (and also other OS)	Windows Server on Windows version Linux Server on Linux version	Windows Serverfamily (Optimized) Linux Server (and also other OS)
	Host OS shared among VMs	No	Yes (guest & host OS are identical, shared kernel)	No
Capability	Processors available to VM on multiprocessor hardware	1	16	1
	Max. memory available for VM	3.6GB	64GB	3.6GB
	Suspend & Resume	Available	Available	Available
	Number of VMs supported	< 64	> 100, > 1000 for Hosting providers	< 64
	SLA Guarantees (Provisioning)	Yes	Yes	Yes
Typical Usage	Server Consolidation	Yes	Yes	Yes
	Legacy Application Hosting	Yes	No	Yes
	Test/Staging Environment	Yes	No	Yes

Para Virtualization

		Xen	Denali	UM Linux
Support Environment	Targeting Host OS	Linux	No Commodity OS	Linux
	Guest OS	Linux (Win XP under development)	No Commodity OS	Linux
	Kernel recompile required	Yes	Yes	Yes
Capability	Limitations	Lower privilege level for guest OS	No Virtual Memory/ BIOS support. Interrupt semantics modified	Poor Performance
	Optimizations for speed	Access to hardware page tables. Supports Virtual Memory	Handling idle loops better for higher performance and scalability.	-
Typical Usage	Server Consolidation	Yes	Yes	-
	Kernel debugging, testing/Staging	No	No	Yes

Conclusion

The setup at enterprise datacenters, today, is sub-optimal. The need to isolate applications from one another and to accommodate peak loads has forced these data centers to be over provisioned. This has led to increased IT spending for procurement of hardware and software licenses and managing the servers. It is only natural for the enterprise to look for a solution that can consolidate these servers, improve utilization, while continuing to isolate application execution environments and handle peak loads.

Virtualization technology addresses these problems. Be it server consolidation or on demand computing needs, virtualization seems to be the natural choice. Virtualization solution can help enterprises in

- Consolidating servers and improving utilization
- Provide on demand computing by scaling up/down the number of Virtual Machines/ applications
- Reduce the number of servers required for fail over clustering, typically from 2*N servers to N+1 servers.
- Allow job/process migration smoothly across physical servers with near zero downtime. Allow smaller enterprises to share servers across production and development environments without compromising the security of the production setup.
- Continue to support legacy applications on older OS, while deploying the same along with newer applications on newer OS Versions on a single shared physical server.
- Facilitate utility computing through deployments on computational grids.

Hosted Virtual Machine Architecture seems to be more mature compared to other solutions. Deploying products like VMWare Server, Virtuozzo and Virtual Server in datacenters can greatly cut down the TCO. Open source products like Xen and Denali have the potential to challenge the hosted model, considering the fact that these solutions are much more scalable and have very low virtualization overhead. Also the changes to the hardware / processor by Intel's Vanderpool and AMD's Pacifica can change the situation, making it even more favorable for rapid adoption of virtualization technologies. The caveat is that, though virtualization can certainly contribute to economic benefits in terms of hardware and personnel cost reduction, in some cases the software cost may increase due to licensing model of server applications and also the cost of virtualization software. A prudent approach would be to evaluate the overall cost benefit considering all the three aspects (hardware, personnel and software costs) before the virtualization exercise is taken up.

References

1. G. J. Popek and R. P. Goldberg, "Formal requirements for virtualizable third generation architectures," *Comm. ACM*, vol. 17, no. 7, pp. 412–421, 1974.
2. Amit Singh, "An Introduction to Virtualization", <http://www.kernelthread.com/publications/virtualization/>
3. A. Whitaker, M. Shaw, and S. D. Gribble. "Denali: Lightweight virtual machines for distributed and networked applications." In Proceedings of the USENIX Annual Technical Conference, Monterey, CA, June 2002. 10
4. Paul Barham , Boris Dragovic , Keir Fraser , Steven Hand , Tim Harris , Alex Ho , Rolf Neugebauer , Ian Pratt , Andrew Warfield, "Xen and the art of virtualization", Proceedings of the nineteenth ACM symposium on Operating systems principles, October 19-22, 2003, Bolton Landing, NY, USA
5. Edouard Bugnion, Scott Devine, Kinshuk Govil, and Mendel Rosenblum. "Disco: Running Commodity Operating Systems on Scalable Multiprocessors." *ACM Transaction on Computer Systems (TOCS)*, Vol. 15, No. 4 (Nov. 1997).
6. "The Plex86 Project," 2003. [Online]. Available: <http://plex86.sourceforge.net/>
7. Rich Uhlig, Gil Neiger, Dion Rodgers, Amy L. Santoni, Fernando C.M. Martins, Andrew V. Anderson, Steven M. Bennett, Alain Kagi, Felix H. Leung, Larry Smith. "Intel Virtualization Technology," *Computer*, vol. 38, no. 5, pp. 48-56, May 2005.
8. Virtual Iron [Online]. Available: <http://www.virtualiron.com>
9. Mendel Rosenblum, Tal Garfinkel. "Virtual Machine Monitors: Current Technology and Future Trends," *Computer*, vol. 38, no. 5, pp. 39-47, May 2005.
10. User-mode Linux [Online]. Available: <http://user-mode-linux.sourceforge.net/>.
11. Virtuozzo: "A Complete Server Virtualization and Automation Solution {Product White Paper & Data Sheet}". Source: <http://www.sw-soft.com/>
12. Carl A. Waldspurger, "Memory resource management in VMware ESX server", *ACM SIGOPS Operating Systems Review*, v.36 n.SI, Winter 2002
13. Microsoft Virtual Server 2005 Technical Overview, [Online]. Available: <http://www.microsoft.com/windowserversystem/virtualserver/overview/vs2005tech.mspx>
14. Jeremy Sugerman, Ganesh Venkitachalam, Beng-Hong Lim, "Virtualizing I/O Devices on VMware Workstation's Hosted Virtual Machine Monitor", Proceedings of the General Track: 2002 USENIX Annual Technical Conference, p.1-14, June 25-30, 2001
15. Microsoft Virtual Server 2005 Product Overview, [Online]. Available: <http://www.microsoft.com/windowserversystem/virtualserver/overview/vs2005prod.mspx>
16. Paul Massiglia, Frank Bunn, "Introduction to Virtualization", *Virtual Storage Redefined: Technologies and Applications for Storage Virtualization*, Ch. 1
17. White Paper on "Server Consolidation and Virtualization", Corosoft Inc



For more information, contact askus@infosys.com

About Infosys

Many of the world's most successful organizations rely on Infosys to deliver measurable business value. Infosys provides business consulting, technology, engineering and outsourcing services to help clients in over 30 countries build tomorrow's enterprise.

For more information about Infosys (NASDAQ:INFY), visit www.infosys.com.