

# PowerVM, The Server Virtualization Without Limits

Mary Merline Rani<sup>1</sup>, Mrinmoyee<sup>1</sup>, M.N.Nachappa<sup>1</sup>  
<sup>1</sup>Department Of Computer Science, St. Joseph's College (Autonomous),  
Langford Road, Shanthinagar, Bangalore – 560027, India.  
mnnachappa@gmail.com

## ABSTRACT

*This paper is about PowerVM virtualization technologies. Maintaining Web performance during peak periods requires dynamic scalability, but achieving that kind of scalability often means overspending on an infrastructure that is underutilized most of the time. Virtual machine technology, or virtualization, is gaining momentum in the information technology community since it can satisfy peak scalability requirements without a costly, oversized infrastructure by replacing or complementing the physical servers with virtual ones. Virtual server services are designed to help reduce the cost and complexity of current Linux, AIX or OS/400 based infrastructure. It is the first multiplatform portfolio of virtual solutions for e-business. Recent advances in hardware and software technology have brought virtualization to the forefront of IT management. PowerVM allows business units to consolidate multiple workloads onto fewer systems, increasing server utilization, and reducing cost. PowerVM technology provides a secure and scalable virtualization environment for AIX, IBM i. Linux applications, hosted on the IBM Power series systems.*

**Keywords:** LPAR, Hardware Management Console(HMC), Virtual IO server (VIOS), PowerVM Hypervisor, Micro-Partitioning, Integrated Virtualization Manager, Virtual SCSI, Active Memory Sharing

## 1. INTRODUCTION

The server virtualization technology in majority of its implementations provides abstraction of the hardware resources of a physical server box to provide virtual resources to the multiple virtual server units built in that box. These virtual server units are sometimes referred to as emulations, partitions, guests, containers or instances based on the virtualization solution being used. These Virtual servers most often run different operating systems, in some cases limited in its choice of operating system based on the underlying proprietary physical server platform being used. They operate independently of each other and within the bounds of the virtual resources assigned to it.

The virtualized environment uses a special firmware/software to distribute and control the virtual resources to the virtual server units. This piece of firmware/software is named the **hypervisor**.

The key objective of server virtualization is to consolidate the services provided from independent physical servers onto a single physical server system, where these services are then provided from the virtual server units built using virtual resources made available by virtualization technology used in that physical server system.

The server consolidation helps business units cut down on their IT hardware investment as less data centre rack space is required due to server consolidation to host their IT hardware leading to lesser power consumption resulting in an optimized and environment friendly data centre hosting environment.

## 2. TYPES OF SERVER VIRTUALIZATION

All common techniques for Server Virtualization are based on the following models,

- VM (Virtual Machine) model
- Operating System level Virtualization and
- Para-Virtual Machine models.

### 2.1 Virtual Machine based Server Virtualization

The VM (Virtual Machine) based Server Virtualization uses VMM (Virtual Machine Monitor) or Hypervisor to coordinate instructions to CPU of the server and the host machine system resources in order to manage the running of guest operating systems without any modification. Since, VM (Virtual Machine)-based Server Virtualization provides masking of the hardware layer, any partitions or instances can be created within the server and can use different operating systems without the guest application systems needing to know about the host operating system environment or be aware of the hardware on which it is actually running.

Some of the well known products based on VM (Virtual Machine) based Server Virtualization are Microsoft Virtual Server, VMware Workstation, VMware Server (GSX), Virtual Iron, QEMU and VirtualBox.

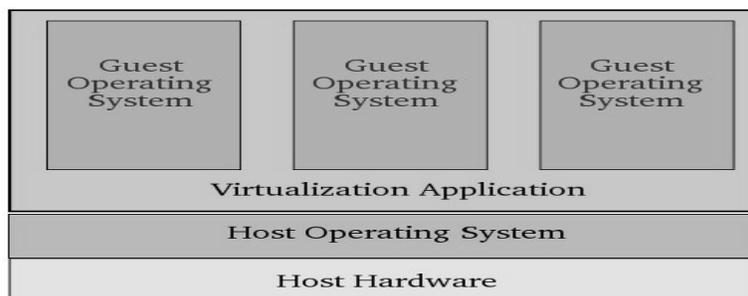


Figure-1 Virtual machine based Server Virtualization

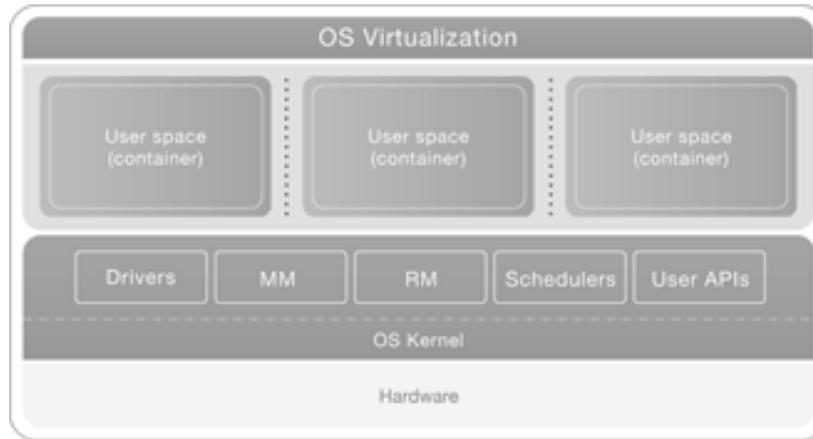
### 2.2 Operating System level Virtualization

Operating System level Virtualization is a server virtualization method where the kernel of the host Operating system allows for multiple isolated user-space instances, instead of just one. Such instances (often called containers, VEs, VPSs or jails) may look and feel like a real server, from the point of view of its owner. In other words, Operating System level Virtualization is based on the creation of isolated containers or partitions on a single physical server, and the use of O/S

instances for each guest application environment to work independently of the other partitions.

On Unix systems, this technology can be thought of as an advanced implementation of the standard chroot mechanism. In addition to isolation mechanisms, the kernel often provides resource management features to limit the impact of one container's activities on the other containers.

Some of the well known products based on Operating System level Virtualization are Parallels Virtuozzo, OpenVZ, FreeBSD Jail and Solaris Zones.



**Figure-2 Operating System level Virtualization**

### 2.3 Para-Virtual machine (PVM) model

The Para-Virtual machine (PVM) model also uses VMM (Virtual Machine Monitor) and enables running multiple operating systems in host server environments like VM (Virtual Machine)-based Server Virtualization. However, unlike the VM model, VMM (Virtual Machine Monitor) used in the para-Virtual machine (PVM) model, modifies the guest operating system's code to enable the running of guest

systems. This modification, also sometimes referred to as "porting," enables faster switching between operating system environments since, in its method of functioning, VMM makes use of privileged system calls less frequently.

Some of the well known products and technologies based on para-Virtual machine (PVM) is Trango Hypervisor, Xen VMM, LDOMs (Logical Domains technology from Sun) and IBM z/VM.



**Figure-3 Para Virtual Machine Server Virtualization**

## 3. SERVER VIRTUALIZATION WITH IBM POWERVM

PowerVM is a server virtualization solution based on IBM power technology.

It is a combination of advanced hardware features of Power technology supported by Value added software.

It allows Business units to consolidate multiple application workloads like Database, HTTP etc. onto fewer systems, increasing server utilization and reducing cost. PowerVM provides a secure and scalable server virtualization

environment for IBM AIX/ i-series and Linux based applications built upon the IBM Power "P" series server platforms optimized for virtualization.

**Key capabilities of PowerVM technology on IBM Power series servers to enable optimal hosting of Business applications in Data Centres:**

- This technology optimises server utilization and IO resource sharing to allow better use of IT assets so that the business applications have to incur lesser cost of ownership for its hosting imprint in the data center and its subsequent steady state support.

- It optimises the performance of the business application by improving its response time based on an increased operational speed brought about by dynamic reallocation of resources as needed to adjust to changes in business needs or a non-linear load on the application.
- It decouples the application workload like database / HTTP etc from the intricacies of the hardware resources in the power series server, thereby simplifying the IT infrastructure management for these applications. This further allows making business-driven policies to provide resources based on time/cost or agreed level of support.

A diagrammatic representation of the use of PowerVM technology on Power series server platform:

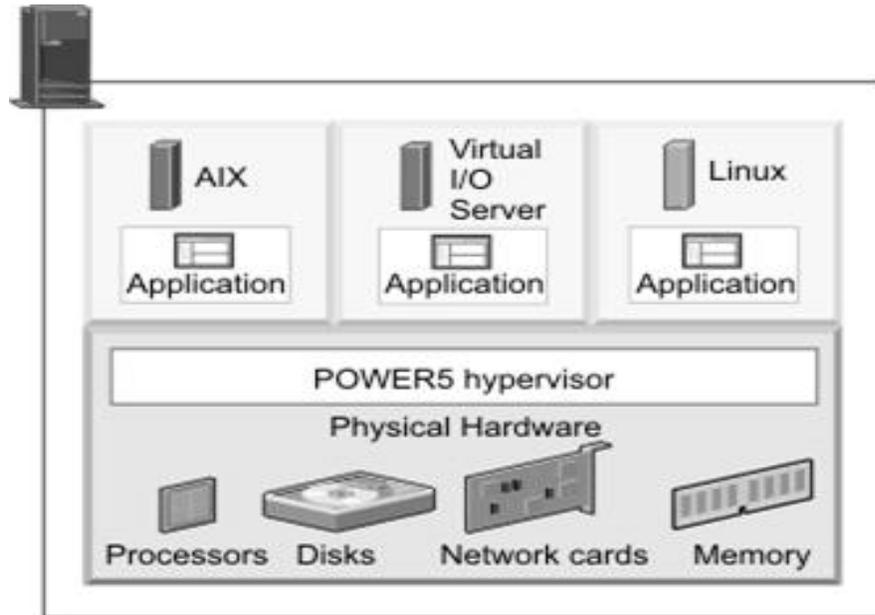


Figure-4 Power VM Server Virtualization

This shows a Power series server providing physical hardware resources like Processor, Disks, Network cards, Memory DIMMs for use by an application workload.

The Power5 Hypervisor, which is the firmware component of PowerVM, seats above the layer of hardware resources to act as a layer of abstraction to hide the physical hardware resources from the application workloads seating above, while controlling the distribution of the hardware resources to these workloads.

#### 4. KEY COMPONENTS OF POWERVM ENVIRONMENT

**LPAR (Logical partition)** : It is a subset of the servers's hardware resources, virtualized as a separate server. In effect, a physical machine can be partitioned into multiple logical partitions, each hosting a separate operating system

**HMC (Hardware Management Console)**: It is a technology invented by IBM for the purpose of providing a stand alone appliance used to create, change or delete logical partitions. Its task is to assign both physical and virtual hardware resources to a logical partition (LPAR) and provide terminal access to all the Lpars.

**VIOS (Virtual I/O Server)**: The Virtual I/O Server is the software component of the Powervm feature, it runs a modified AIX operating system shell, in a separate logical partition (LPAR) and interacts with the Hypervisor. This software facilitates the sharing of physical I/O resources

between client logical partitions built in the physical server. The VIOS owns the physical hardware resources, such as Ethernet and storage interface cards which are connected to the external switch fabric facilitating connection to external network and storage subsystems. The network or storage resources reachable through the VIOS are then shared with the client logical partitions by the hypervisor. This enables the administrators to minimize the number of physical servers deployed in their network.

With VIOS, client partitions can communicate between each other on the same physical server without access to physical Ethernet adapters.

Access to Physical Ethernet adapters are provided through a Virtual Ethernet adapter at the Client partition which is mapped by the hypervisor to the physical Ethernet adapter configured in the Virtual IO server to allow the applications running in the client partitions to reach the external networks.

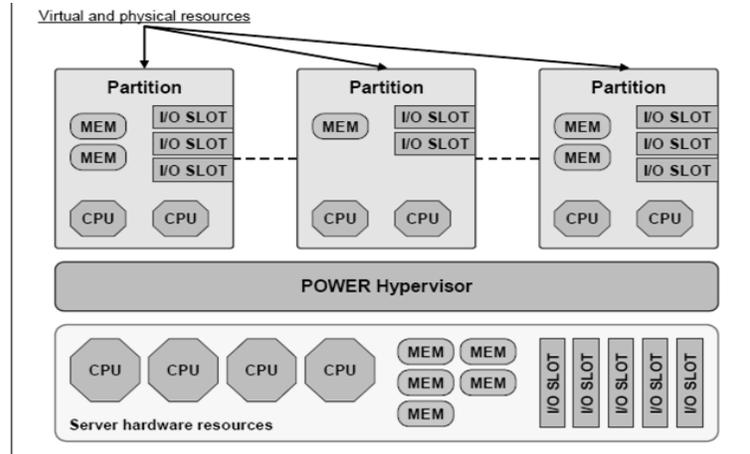
Similarly, access to external storage subsystem is provided by the mapping of a virtual SCSI adapter in the client partition to a Physical Fiber adapter in the Virtual IO server by the hypervisor, the physical Fibre adapter is in turn connected to external storage subsystems, allowing applications on the client partitions to read/write data on the external storage.

**NIM (Network Installation Manager)**: It is a separate AIX Server which is used to install IOS in the virtual I/O server partition over the network. IOS is the modified AIX shell which is run in the virtual I/O server partition .It represents the

software part of PowerVM technology. NIM is also used to install AIX operating system onto client logical partitions.

**PowerVM Hypervisor:** The POWER Hypervisor is a firmware layer sitting between the hosted operating systems and the server hardware, as shown in the figure-5 below. It shows a physical system running three logical partitions. Each logical partition is running an operating environment and

applications which are independent of the other logical partitions, although they are sharing the same physical resources. The POWER Hypervisor has no specific or dedicated processor resources assigned to it and it is automatically installed and activated, regardless of system configuration.



**Figure-5 Power hypervisor**

The POWER Hypervisor performs the following tasks:

- Enforces partition integrity by providing a security layer between logical partitions.
- Provides an abstraction layer between the physical hardware resources and the logical partitions using them. It controls the dispatch of virtual processors to physical processors, and saves and restores all processor state information during virtual processor context switch.
- Controls hardware I/O interrupts and management facilities for partitions.

It supports multiple operating environments on a single system. It is the foundation for virtualization on IBM PowerVM . The hypervisor manages system virtualization, including creating logical partitions and dynamically moving resources across multiple operating environments. The hypervisor provides the ability to divide physical system resources into isolated logical partitions. Each logical partition operates like an independent system running its own operating environment: AIX, IBM i, Linux, and the Virtual I/O Server (VIOS). The hypervisor can assign dedicated processors, I/O, and memory, which can be dynamically reconfigure as needed, to each logical partition. The hypervisor can also assign shared processors to each logical partition using its **Micro-Partitioning** feature. Unknown to the logical partitions, the hypervisor creates a shared processor pool from which it allocates virtual processors to the logical partitions as needed. In other words, the hypervisor creates virtual processors so that logical partitions can share the physical processors while running independent operating environments.

**Micro-Partitioning:** Micro-Partitioning is the technology that allows distributing the processing capacity or allocating fraction of physical processors among one or more logical partitions. A logical partition using fractions of processors is also known as a Shared Processor Partition or Micro-Partition. Micro-Partitions run over a set of processors called a Shared Processor Pool. The benefit of Micro-Partitioning is that it

allows significantly increased overall utilization of processor resources within the system.

Virtual processors are used to let the operating system manage the fractions of processing power assigned to the logical partition. From an operating system perspective, a virtual processor cannot be distinguished from a physical processor, unless the operating system has been enhanced to be made aware of the difference. Physical processors are abstracted into virtual processors that are available to partitions. The meaning of the term *physical processor* in this section is a *processor core* .When you create a logical partition, you can assign entire processors for dedicated use, or you can assign partial processor units from a shared processor pool

**Dynamic Logical Partitioning:** Processor, memory, and I/O resources can be dynamically moved between logical partitions (LPARs) using Dynamic logical partitioning (DLPAR).

It increases the flexibility, allowing selected system resources, such as processors, memory, and I/O components to be added and deleted from logical partitions while they are operational. The ability to reconfigure dynamic LPARs encourages system administrators to dynamically reallocate all available system resources to reach the optimum capacity for each defined dynamic LPAR.

## 5. INTEGRATED VIRTUALIZATION MANAGER

Integrated Virtualization Manager (IVM) is a management tool that combines partition management and Virtual I/O Server (VIOS) functionality into a single partition running on the system. The IVM features an easy-to-use point-and-click interface and is supported on blades and entry-level to mid-range servers. Using the IVM helps lower the cost of entry to PowerVM virtualization because it does not require a Hardware Management Console.

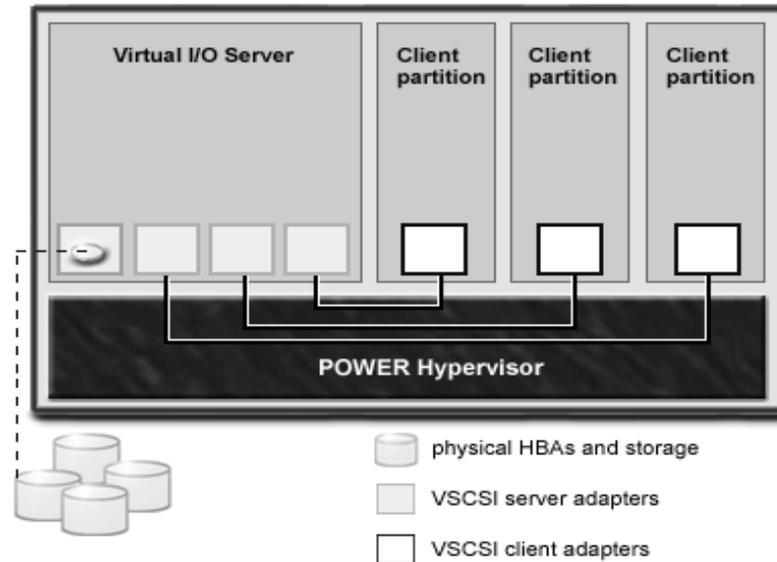


Figure-6 Virtual SCSI configuration

Virtual SCSI allows client logical partitions to share disk storage and tape or optical devices that are assigned to the Virtual I/O Server logical partition. There are virtual SCSI adapters in the server partition as well as client partition which are mapped together with the help of the hypervisor.

## 6. ACTIVE MEMORY SHARING

Active Memory Sharing is an IBM PowerVM advanced memory virtualization technology that provides system memory virtualization capabilities to IBM Power Systems, allowing multiple logical partitions to share a common pool of physical memory. The physical memory of an IBM Power System can be assigned to multiple logical partitions in either a dedicated or shared mode.

The system administrator has the capability to assign some physical memory to a logical partition and some physical memory to a pool that is shared by other logical partitions. A single partition can have either dedicated or shared memory.

## 7. CONCLUSION

PowerVM provides the industrial-strength virtualization solution for IBM Power Systems servers and blades. Based on more than a decade of evolution and innovation, PowerVM represents the state of the art in enterprise virtualization and is broadly deployed in production environments worldwide by most Power Systems owners.

The IBM Power Systems family of servers includes proven workload consolidation platforms that help clients control costs while improving overall performance, availability and energy efficiency. With these servers and IBM PowerVM virtualization solutions, an organization can consolidate large numbers of applications and servers, fully

virtualize its system resources, and provide a more flexible, dynamic IT infrastructure.

Security on PowerVM is built directly into the hardware, running checks on all system files at bootup looking for malicious changes. It has superior configuration flexibility.

PowerVM virtual servers can be configured to use as much or as little of the overall system as they need. Virtual servers can be set up to use as little as 10% of a single processor core, or to take over the entire systems' resources.

Thus we can see that PowerVM is a very useful and powerful technology for server Virtualization.

## REFERENCES

1. Power Systems Enterprise Servers with PowerVM Virtualization and RAS, By: *Dino Quintero; JinHoon Baek; Guillermo Diez; Hassan Elsetohy; Debra Francis; Bing He; Rajesh Jeyapaul; Anil Kalavakolanu; Tejaswini Kaujalgi; David Kgabo; Ricardo Puig; Vani Ramagiri* Publisher: *IBM Redbooks*
2. Virtualization For Dummies By *Bernard Golden*
3. IBM PowerVM Virtualization Managing and Monitoring, *An IBM Redbooks publication*
4. Virtualization: From the Desktop to the Enterprise (Books for Professionals by Professionals) By *Chris Wolf and Erick M. Halter*

### Online Resources

- [1] <http://en.wikipedia.org/wiki/PowerVM>
- [2] <http://www.ibm.com/>
- [3] <http://www.redbooks.ibm.com/>
- [4] <http://my.safaribooksonline.com/book/>