

A COMPARISON OF POWERVM AND X86-BASED VIRTUALIZATION PERFORMANCE

Dr. Gaurav Kumar Jain^{#1}, Mukul Verma^{#2}

^{#1} Associate Professor, Department of Computer Science & Information Technology, Regional College For Education Research and Technology, Jaipur,

^{#2} M. Tech. Scholar Regional College For Education Research and Technology, Jaipur, Rajasthan

gaurav.rinkujain.jain@gmail.com
MUKUL.VERMA1986@GMAIL.COM

ABSTRACT

Virtualization has moved to the forefront of many organizations' IT agendas, because it reduces both the number of physical servers required and the associated costs. This independent study by the International Technology group found IBM PowerVM more cost-competitive over time, compared to VMware, in terms of IT costs and the costs of downtime.

KEYWORDS

Virtualization, it infrastructure, software partitioning, reduce costs, risk management, server virtualization

1. INTRODUCTION

Virtualization technologies allow IT organizations to consolidate workloads running on multiple operating systems and software stacks and allocate platform resources dynamically to meet specific business and application requirements. Leadership virtualization has become the key technology to efficiently deploy servers in enterprise data centers to drive down costs and become the foundation for server pools and cloud computing technology. Therefore, the performance of this foundation technology is critical for the success of server pools and cloud computing.

- Virtualization may be employed to:
- Consolidate multiple environments, including underutilized servers and

systems with varied and dynamic resource requirements

- Grow and shrink resources dynamically, derive energy efficiency, save space, and optimize resource utilization
- Deploy new workloads through provisioning virtual machines or new systems rapidly to meet changing business demands
- Develop and test applications in secure, independent domains while production can be isolated to its own domain on the same system
- Transfer live workloads to support server migrations, balancing system load, or to avoid planned downtime
- Control server sprawl and thereby reduce system management costs

2. POWERVM

With IBM Power™ Systems and IBM PowerVM virtualization technologies, an organization can consolidate applications and servers by using partitioning and virtualized system resources to provide a more flexible, dynamic IT infrastructure. PowerVM delivers industrial strength virtualization for AIX®, IBM i, and Linux® environments on IBM POWER processor-based systems; the Power Hypervisor™ supports multiple operating environments on a single system and is integrated as part of the system firmware. PowerVM offers the flexibility of

combining dedicated and shared resources in the same partition. IBM Power Systems servers and PowerVM technology are designed to deliver a dynamic infrastructure, reducing costs, managing risk and improving service levels.

2.1. Processor Virtualization

PowerVM's advanced dynamic logical partitioning (LPAR) capabilities allow a single partition to act as a completely separate AIX, IBM i, or Linux operating environment. Partitions can have dedicated or shared processor resources. With shared resources, PowerVM can automatically adjust pooled processor resources across multiple operating systems, borrowing processing power from idle partitions to handle high transaction volumes in other partitions.

PowerVM Micro-Partitioning™ supports up to 10 dynamic logical partitions per processor core. Depending upon the Power server, up to 254 independent virtualized servers can be run on a single physical Power server — each with its own processor, memory, and I/O resources. These partitions can be assigned at a granularity of 1/100th of a core. Consolidating systems with PowerVM can reduce operational costs, improve availability, ease management and improve service levels, while allowing businesses to quickly deploy applications.

Shared processor pools allow for the automatic non-disruptive balancing of processing power between partitions assigned to shared pools, resulting in increased throughput. It also provides the ability to cap the processor core resources used by a group of partitions to potentially reduce processor-based software licensing costs.

Shared dedicated capacity allows for the “donation” of spare CPU cycles from dedicated processor partitions to a shared processor pool. The dedicated partition maintains absolute priority for dedicated CPU cycles. Enabling this feature can help

to increase system utilization, without compromising the computing power for critical workloads in a dedicated processor. PowerVM logical partitioning for POWER6™ processor-based systems has received the Common Criteria Evaluation and Validation Scheme (CCEVS) EAL4+ certification for security capabilities.

2.2. Memory Virtualization

PowerVM features Active Memory™ Sharing, the technology that allows an organization to intelligently and dynamically reallocate memory from one partition to another for increased utilization, flexibility and performance. Active Memory Sharing enables the sharing of a pool of physical memory among logical partitions on a single server, helping to reduce the memory resource capacity in a consolidated environment by increasing memory utilization and driving down system costs. The memory is dynamically allocated amongst the partitions as needed, to optimize the overall physical memory usage in the pool. Along with shared memory, PowerVM also supports dedicated memory allocation, enabling partitions sharing memory and partitions with dedicated memory to coexist in the same system.

2.3. I/O Virtualization

Delivered with PowerVM is the Virtual I/O Server (VIOS), a special-purpose partition that can be used to virtualize I/O resources to client partitions. VIOS owns the resources that are shared with clients. A physical adapter assigned to the VIOS partition can be shared by one or more other partitions. VIOS is designed to reduce costs by eliminating the need for dedicated network adapters, disk adapters and disk drives, and tape adapters and tape drives in each client partition. With VIOS, client partitions can easily be created for test, development, or production purposes. PowerVM also supports dedicated I/O along with VIOS, on the same system. Therefore, a single system can have I/O hosted by VIOS for some partitions while

other partitions can have dedicated I/O devices.

2.4. Partition Mobility

Live Partition Mobility facilitates the migration of a running AIX or Linux partition from one physical server to another compatible server without application downtime for planned system maintenance, migrations, provisioning, and workload management.

2.5. PowerVM Lx86 Support for Linux Applications

PowerVM Lx86 is a cross-platform virtualization solution that enables the running of a wide range of x86 Linux applications on Power Systems platforms within a Linux on Power partition without modifications or recompilation of the workloads. This feature enables rapid consolidation of x86 applications onto Power Systems platforms to take advantage of the advanced performance, scalability, and RAS characteristics.

2.6. Workload Partitioning

PowerVM technology also supports another mode of virtualization capability called Workload Partitions (WPARs), a software partitioning technology that is provided by AIX. Introduced with AIX Version 6, WPAR does not have any dependencies on hardware features. WPAR enables consolidation of workloads on a single AIX operating system by providing isolation between workloads running in different WPARs. From an application perspective it is running in its own operating system environment. A key feature in WPAR is mobility; a running WPAR can be relocated from one VM to another VM irrespective of where the VMs are hosted. This feature enables applications to be migrated to another system during software upgrades and other planned maintenance, to balance workloads, to provision rapidly to meet growth dynamically and to improve energy efficiency by further consolidating on the fly during low load periods.

2.7. Systems Management

IBM Systems Director supports the PowerVM environment and is the IBM management tool for multiple, heterogeneous servers. IBM Systems Director supports advanced management functions such as workload lifecycle management, health check, and topology mappings, as well as the ability to take action on monitored events.

IBM Systems Director VMControl™ V2.2 is a plug-in option for IBM Systems Director that represents a transformation from managing virtualization to using virtualization to better manage an entire IT infrastructure. IBM Systems Director and VMControl are designed to help reduce the total cost of ownership in a virtual environment by decreasing management costs, increasing asset utilization, and linking infrastructure performance to business goals.

VMControl is available in three Editions, to suit the varying levels of virtualization deployment at client sites. VMControl Express Edition provides basic virtual machine lifecycle management; VMControl Standard Edition adds virtual appliance lifecycle management; and VMControl Enterprise Edition adds system pool lifecycle management. VMControl Standard Edition captures information from active systems and stores the information in a repository as reusable system images, also referred to as virtual appliances. VMControl Enterprise Edition allows users to create and manage system pools – or groups of virtual appliances deployed across multiple physical servers – as easily as managing a single entity. The advanced virtualization management capabilities of VMControl provide a pathway for organizations to build sophisticated cloud computing environments.

3. PowerVM Advantages

PowerVM offers a secure virtualization environment built on the advanced RAS features and leadership performance of the Power Systems platform. Numerous advantages of PowerVM exist including:

Higher resource utilization: PowerVM promotes high resource utilization by virtualizing resources including processors, memory and I/O across multiple virtual machines

Flexibility: PowerVM runs on all Power Systems servers, from blades to top of the line high-end servers. PowerVM provides the most flexibility in combining dedicated and shared resources in a partition supporting dynamic resource allocation.

Scalability: PowerVM supports partitions as small as 1/10 of a processor and as large as 64 physical processors. On a single Power system, up to 254 partitions can be accommodated.

Availability: Live Partition Mobility helps eliminate planned downtime by moving the partition while it is running to another server, upgrading or maintaining hardware without interrupting productive work.

PowerVM is commonly employed with enterprise class applications and workloads because of its level of sophistication and maturity. VMware vSphere is third-party software that must be installed on x86 hardware that leverages hardware-assist virtualization optimizations. In contrast, Power Systems servers implement a virtualization architecture with components embedded in the hardware, firmware and operating system software. The capabilities of this integrated virtualization architecture are thus significantly different and in many areas more advanced.

PowerVM, by enabling “firmware-based” partitions, provides greater partition isolation than software-based virtualization technologies. Firmware-based logical partitions (or VMs) reduce the potential for performance bottlenecks and contribute to higher levels of availability and security

than may be realized with software-based virtualization. They also contribute to increased linear scalability.

Power Systems servers and PowerVM capabilities are more granular and more closely integrated than is the case for VMware, Hyper-V and equivalent x86-based virtualization tools. The Power Systems platform also benefits from numerous industry-leading availability optimization features. These distinctive capabilities have caused widespread adoption of Power Systems servers to support transaction and database-intensive systems whose performance and uptime requirements are significantly more demanding than the norm.

The importance of workload management should be highlighted. Partitioning creates the potential for high levels of capacity utilization. The extent to which this potential will be realized in practice depends on the mechanisms that allocate system resources between and monitor and control workload execution processes across partitions. If these mechanisms are ineffective, a high proportion of system capacity may be idle at any given time. Close integration of partitioning and workload management capabilities is necessary to minimize risks that surges in workloads running in individual partitions will impact performance and availability. POWER6-based systems also have larger memory per core, higher memory and I/O bandwidth per core as well as larger cache per core compared to Intel Xeon 5500 series processors; therefore POWER6 systems can consolidated a higher number of partitions and can more effectively handle workload surges, leading to demonstrably higher performance.

PowerVM is ultimately well optimized to handle business-critical systems and complex multi-partition production environments. The next section of the paper will highlight a performance analysis of PowerVM and x86 technologies that employs industry-standard benchmarks.

4. Performance Comparison of PowerVM and x86 Virtualization Technologies

To compare the performance of virtualization technologies, it is useful to highlight industry standard and ISV benchmark results. *Every single Power Systems benchmark includes the hypervisor and is considered virtualized.* Virtualized and native performance should not be confused; benchmark results achieved without an active hypervisor should not be confused with results that did include a hypervisor.

This comparison will focus on three industry benchmarks: the online transaction processing benchmark TPC-C, the SAP Sales and Distribution two-tier benchmark and the Storage Performance Council SPC-1 benchmark.

AIM7 Multiuser and DayTrader 2.0 Benchmark performance is also highlighted, comparing Power Systems and

PowerVM with Intel Xeon 5570 processors and Windows 2008 Server Hyper-V, another x86-based virtualization technology.

4.1. Industry Standard Benchmark Performance

TPC-C, an OLTP system benchmark, simulates an environment in which a population of terminal operators execute transactions against a database. The benchmark is centered on the principal activities and transactions of an order-entry environment. These transactions include entering and delivering orders, recording payments, checking the status of orders, and monitoring the level of stock at the warehouses. Note that IBM is currently the only vendor publishing TPC-C results with virtualization technology such as PowerVM. Figure 1 shows the TPC-C benchmark results at 4, 8, 16, and 64 cores. Even with virtualization technologies inherently included as discussed above, the greater than linear scaling of the results is noteworthy.




Figure 1. IBM TPC-C Results

	System	Chips/cores/threads	Avail Date	\$/tpmC	tpmC
TPC-C	IBM Power 570	2/4/8	11/26/07	\$3.50	404,462
TPC-C	IBM Power 550	4/8/16	04/20/08	\$2.49	629,159
TPC-C	IBM Power 570	8/16/32	11/21/07	\$3.54	1,616,162
TPC-C	IBM Power 595	32/64/128	12/10/08	\$2.81	6,085,166

Table 1. IBM TPC-C Results

The SAP Sales and Distribution (SD) Benchmark covers a sell-from-stock scenario, which includes the creation of a customer order with five line items and the corresponding delivery with subsequent goods movement and invoicing.

For 2009, the SAP SD Benchmark was updated to the SAP enhancement package 4 for SAP ERP 6.0 to reflect the practice of SAP customers across all industries. The steps of the benchmark scenario remained unchanged but made the SD benchmark more resource-intensive. For this reason,

comparisons between SAP SD results on ERP Release 6 (Figures 2, 3 and 4) with SAP enhancement package 4 for SAP ERP 6.0 (Unicode) (Figure 5) have been avoided in this paper

Figure 2 compares a “Full Resource Partition” SAP SD 2-tier IBM result with a result that is noted as using “virtual CPUs.” Note that the users per core is approximately the same for each result, a conclusion that is expected as both configurations contain PowerVM and have implemented virtualization.

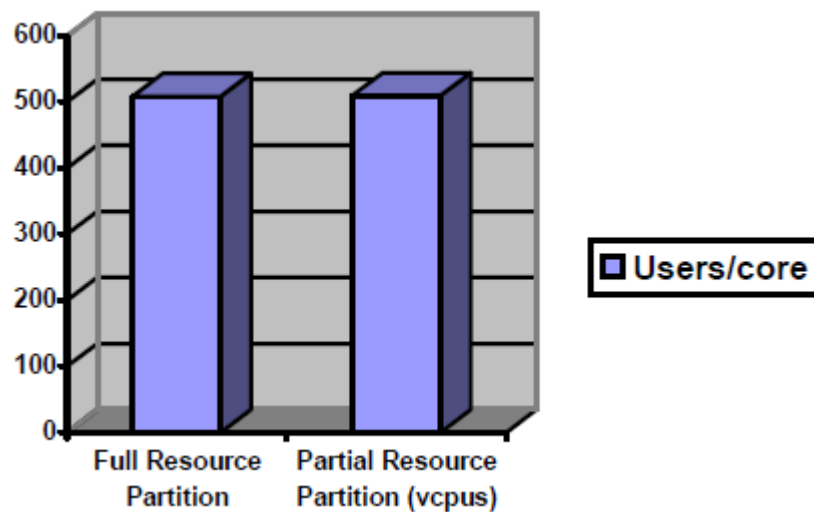


Figure 2. IBM SAP SD 2-tier “Full Resource Partition” result vs. “Partial Resource Partition (vcpus)” result

	System	Proc chips/cores/threads	Cert #	Users	Users/core	OS	ERP release
SAP SD 2-tier result	IBM Power 570	2/4/8	2007037	2035	508.75	AIX 6.1	6.0
SAP SD 2-tier “virtual”	IBM Power 570	2/4/8 using 2 virtual cpus	2008080	1020	510	AIX 6.1	6.0

Table 2. IBM SAP SD 2-tier “Full Resource Partition” result vs. “Partial Resource Partition (vcpus)” result

In contrast, Figure 3 depicts a comparison between “single system image in native mode without virtualization” results vs. results using VMware virtualization. Note that a major difference exists between these

two results; the “virtualized” results achieve fewer users per core. For these configurations, a price is paid in performance for the ability to participate in x86 virtualization technologies.




Figure 3. SAP SD 2-tier “Single system OS image in native mode” results vs. “VMware Virtualization” results

	System	Proc chips/cores/threads	Cert #	Users	Users/core	OS	ERP release
SAP SD 2-tier result	HP BL 460c blade x5355@ 2.66GHz	2/8/8	2007002	1841	230	Windows Server 2003	6.0
SAP SD 2-tier virtual	HP BL 460c blade x5355@ 2.66GHz	2/8/8	2007035	402	201	VMware 3.01/ Windows Server 2003	6.0

Table 3. SAP SD 2-tier “Single system OS image in native mode” results vs. “VMware Virtualization” results

Figure 4 shows users per core for virtualized results.

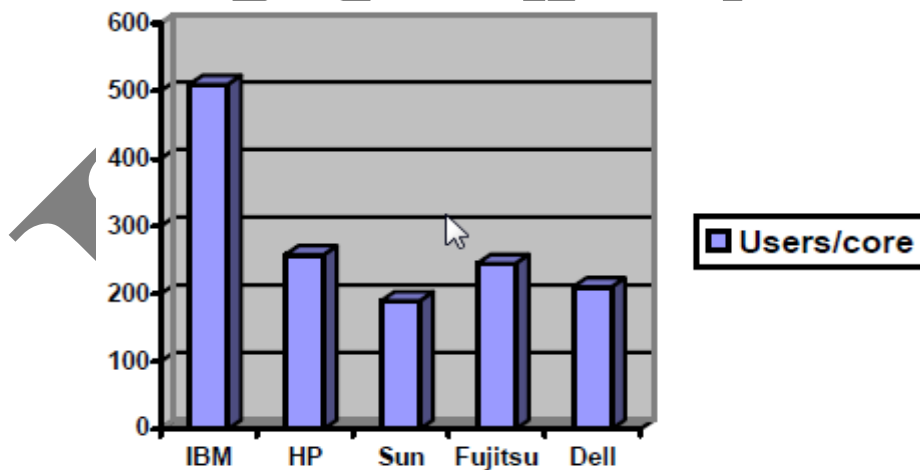


Figure 4. SAP SD 2-tier “virtual” results

	System	Proc chips/cores/threads	Cert #	Users	Users/core	Virtualization SW/OS	ERP release
SAP SD 2-tier virtual	IBM Power 570	2/4/8 using 2 virtual cpus	2008080	1020	510	PowerVM/ AIX 6.1	6.0
SAP SD 2-tier virtual	Dell PowerEdge 2900	2/8/8 using 2 virtual cpus	2007048	421	210.5	VMware ESX 3.0.1/Windows Server 2003	6.0
SAP SD 2-tier virtual	Fujitsu BX630	4/8/8 using 2 virtual cpus	2008008	490	245	VMware ESX 3.5/Windows Server 2003	6.0
SAP SD 2-tier virtual	HP DL580	4/16/16 using 2 virtual cpus	2008007	516	258	VMware ESX 3.5/Windows Server 2003	6.0
SAP SD 2-tier virtual	Sun Fire X2200	2/4/4 using 2 virtual cpus	2007067	380	190	VMware ESX 3.0.1 / Windows Server 2003	6.0

Table 4. SAP SD 2-tier “virtual” results

Finally, Figure 5 compares an IBM result with another virtualized result using the enhancement package 4 for SAP ERP 6.0 (Unicode) kit. As noted above, the IBM result used in these comparisons can be a “Full resource partition” result or one that

states “virtual cpus” (partition resource partition) since every Power Systems result uses PowerVM. This comparison confirms that the SAP enhancement package 4 for SAP ERP 6.0 (Unicode) kit does not change the basic performance and conclusions above.

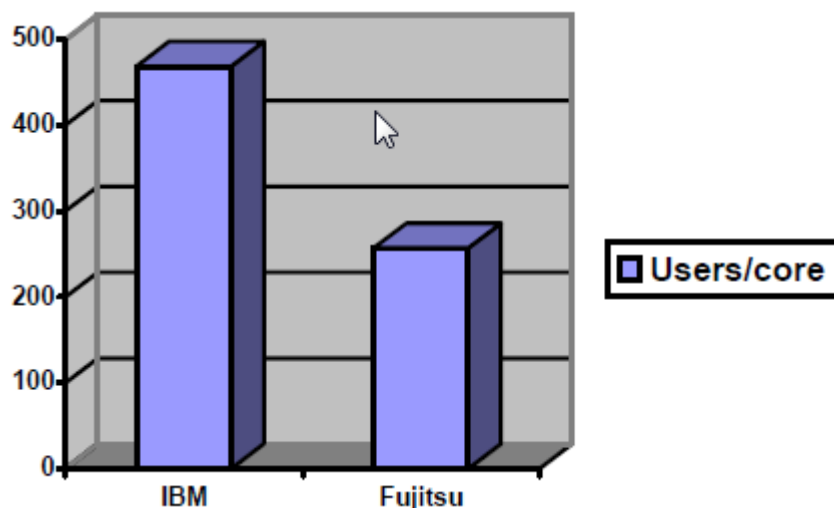


Figure 5. SAP Enhancement package 4 for SAP ERP 6.0 (Unicode) Kit Results

	System	Proc chips/cores/threads	Cert #	Users	Users/core	Virtualization SW/OS	ERP release
SAP SD 2-tier result	IBM Power 550	4/8/16	2009023	3752	469	PowerVM/AIX 6.1	6.0
SAP SD 2-tier result	Fujitsu	2/8/16 using 8 virtual cpus	2009029	2056	257	VMware ESX 4.0/SuSE Linux Enterprise Server 10	6.0

Table 5. SAP Enhancement package 4 for SAP ERP 6.0 (Unicode) Kit Results

5. Integrated Storage Virtualization and Solid State Technology Performance

As shown in Figure 6, IBM recently published a new Storage Performance Council SPC-1 benchmark result using the IBM Power 595 with PowerVM and Solid State Drives (SSD). SPC-1 consists of a single workload designed to demonstrate the performance of a storage subsystem while performing the typical functions of business critical applications. Those applications are characterized by

predominately random I/O operations and require both queries as well as update operations.. Examples of those types of applications include OLTP, database operations, and mail server implementations. It also includes some sequential IO to represent a level of batch component as might occur in a typical production environment. As detailed in the SPC-1 report, this configuration embraced virtualization to achieve a leadership result. The configuration entailed five virtual machines, one for AIX to drive the workload, and four instances of PowerVM VIOS for redundancy and mirroring. This result is the current top SPC-1 result.

	System	SPC-1 IOPS	\$/IOP	Data Protection	Date
SPC-1	IBM Power 595 with PowerVM and SSD	300,993.85	10.71	Mirroring	October 2009

Figure 6.SPC-1 Result with PowerVM

6. Application Benchmark Performance: POWER6 and PowerVM vs. Intel Xeon 5570 and Hyper-V

AIM7 Multiuser Benchmark, a well known open source benchmark, was recently run in IBM's performance lab. The test suites stress the CPU, memory, and I/O, covering a wide range of operations. Figure 7 depicts the Virtual Machine scaling test results between PowerVM and Microsoft Windows Hyper-V, an x86 virtualization

technology, using the AIM7 Compute Benchmark. The systems under test include an IBM Power 550 with POWER6 processors and an HP DL370 with Intel Xeon 5570 processors, running the same level RHEL 5.3 operating system on both. The results highlight superior Power systems and PowerVM performance on this benchmark surpassing Intel Xeon 5570 processors and Hyper-V by greater than 70% across 1 to 8 VMs.

7. AIM7 VM Scaling




Figure 7. AIM7 PowerVM vs. Windows Hyper-V

System Configuration	Cores	Virt SW	Virtual Machines	Jobs / min
IBM Power 550 5GHz 64GB RAM, RHEL 5.3 SMT enabled	8	PowerVM	1	17111
			4	66753
			6	101214
			8	135590
HP DL370 Intel Xeon 5570 2.9 GHz 96GB RAM (HT and Turbo enabled in BIOS Intel VTx with EPT HW virtualization assist) RHEL 5.3 (GA, x86_64)	8	Hyper-V	1	9692
			4	38114
			6	54371
			8	71813

Table 7. AIM7 PowerVM vs Windows Hyper-V

8. DayTrader 2.0 Benchmark

DayTrader2.0, an open source benchmark, was used to compare POWER6 and PowerVM performance with Intel Xeon 5570 processors and Microsoft Windows Hyper-V. DayTrader 2.0 is a stock trading system benchmark build around JAVA EE technologies. DayTrader2.0 was configured in three-tier mode where the back-end database was installed on a separate system. The same back-end database server and the

client (JIBE drivers) systems were used in both PowerVM and Hyper-V cases. The only difference between the two configurations compared in this test was the system under test for the application server tier. The study compared the Power 550 with PowerVM and HP DL370 G6 Intel Xeon 5570 system with Hyper-V technology for the application server tier running AIX and Windows 2008 Server Data Center version respectively. The back-end database used IBM DB2® v9.7 and the application server tier used IBM

WebSphere® V7.0.05 middleware software stacks.

The study reveals that when running an EE JAVA application on PowerVM and POWER6 systems surpass Hyper-V and

Intel Xeon 5570 processors system by 61% with 8 vcpus. (Note that Hyper-V does not support 8vcpus in a single VM; therefore on Hyper-V, 2 VMs each with 4 vcpus were used to compare with 8 vcpus results on PowerVM.)

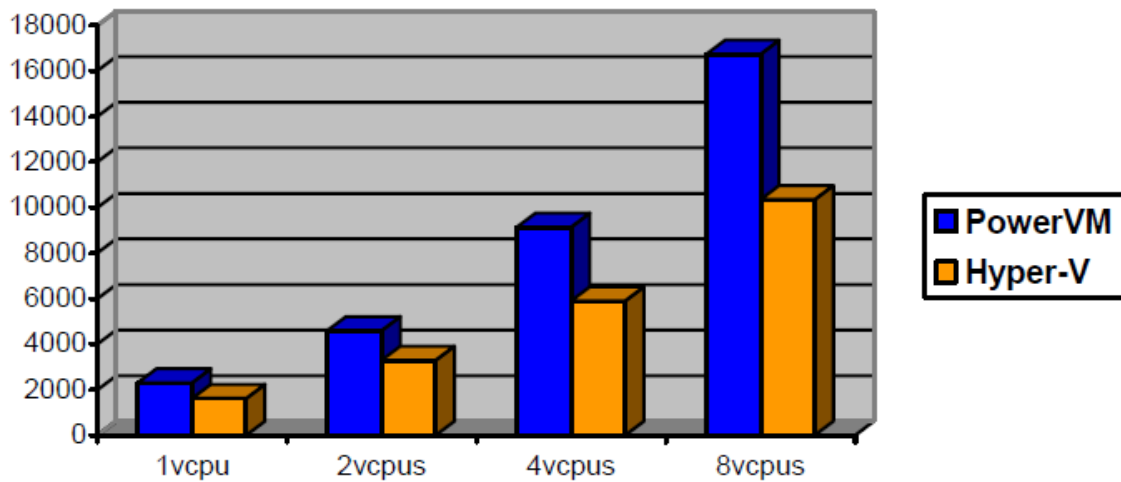


Figure 8. DayTrader 2.0 PowerVM vs. Windows Hyper-V

System Configuration	Cores	Virt SW	Virtual Processors	Page Transactions/sec
IBM Power 550 4.2 GHz 128GB RAM, AIX v6.1 TL03 SMT enabled, IBM WebSphere v7.0.0.5	8	PowerVM	1	2320
IBM Power 570 4.7GHz 64GB RAM, AIX v6.1 TL03 SMT enabled, DB2 v9.7, JDBC			2	4609
4 JIBE Drivers			4	9132
			8	16732
HP DL370 Intel Xeon 5570 2.9 GHz 96GB RAM (HT and Turbo enabled in BIOS Intel VTx with EPT HW virtualization assist) Windows 2008 Server Data Center, IBM WebSphere v7.0.0.5	8	Hyper-V	1	1618
IBM Power 570 4.7GHz 128GB RAM, AIX v6.1 TL03 SMT enabled, DB2 v9.7, JDBC			2	3300
4 JIBE Drivers			4	5899
			8	10346

Table 8. DayTrader 2.0 PowerVM vs. Windows Hyper-V

9. CONCLUSION

Many organizations around the globe are looking to take advantage of virtualization technologies. With IBM Power Systems and PowerVM, clients can achieve virtualization with outstanding performance.

As can be seen from the virtualization performance results, the scalability and performance of the PowerVM benchmarks are unsurpassed. Whether running an OLTP, SAP, storage or application workload, the PowerVM virtualization

technology produces leadership results. In fact, a Power Systems benchmark containing virtualization technologies may not even state that it uses PowerVM; as seen in the study, organizations can count on outstanding performance with no increased cost for every type of workload.

ACKNOWLEDGEMENTS

We would like to thank everyone, especially developers, early adopters, I.T. enthusiast and family members who provided support and followed up with us.

REFERENCES

- [1.] <http://www.tpc.org>
- [2.] <http://www.sap.com/benchmark>
- [3.] <http://www.storageperformance.org/results>
- [4.] <http://www.ibm.com/systems/power/software/virtualization/index.html>
- [5.] <ftp://ftp.software.ibm.com/common/ssi/pm/sp/n/pod03015usen/POD03015USEN.PDF>
- [6.] http://www.vmware.com/files/pdf/vsphere_performance_wp.pdf
- [7.] <http://www.sap.com/solutions/benchmark/index.epx>
- [8.] <http://www.storageperformance.org/results>
- [9.] <http://aimbench.sourceforge.net/>
- [10.] <http://cwiki.apache.org/GMOxDOC20/daytrader.html>