



# A Comparison of PowerVM and VMware Virtualization Performance

*April 14, 2010*

*Mala Anand, Yong Cai, William Bostic, Terry Thomas, Ian Robinson and Elisabeth Stahl*

## **POWER7 Virtualization Performance: Executive Summary**

With its 2010 launch of a new generation of Power Systems servers and blades based on the POWER7 processor architecture, IBM has extended its market lead and elevated system performance, throughput and energy efficiency to unprecedented levels that far outpace competitors. Most importantly, the new architecture also provides the foundation for the integrated PowerVM virtualization solution to deliver unrivaled scalability, flexibility and robustness. As a result, enterprise workloads deployed in PowerVM virtual machines (VMs) not only run faster on POWER7-based platforms, but they can also scale further and be optimized more efficiently.

This technical white paper demonstrates the extent of the performance lead that PowerVM enjoys over x86-based add-on virtualization products, by running identical virtualized workload benchmarks on comparable POWER7- and Intel-based systems. The benchmark results show the dramatic improvements that can be expected when deploying virtualized workloads on Power Systems servers, compared to the performance of those same workloads on x86-based platforms using a third-party virtualization product such as VMware vSphere 4. Key findings include the following:

- PowerVM on Power 750 performs up to 65% better than VMware
- PowerVM on Power 750 scales to four times more virtual CPUs than VMware in a virtual machine
- PowerVM on Power 750 scales linearly to use all CPUs, while VMware does not

It is clear from the published benchmark results that PowerVM on POWER7-based platforms not only offers vastly superior scalability than VMware vSphere on Intel x86-based servers, but it also makes more efficient use of system resources and imposes a negligible impact on performance. Many of these advantages relate to the fact that PowerVM is built directly into the firmware of all Power Systems servers, as opposed to x86-based virtualization products such as VMware vSphere, which are typically third-party software add-ons that are sold and installed separately.

In summary, the benchmark results published in this white paper prove that PowerVM on POWER7 platforms offers a far superior virtualization solution than VMware vSphere on Intel x86 platforms, with higher performance, broader scalability and increased flexibility.

# IBM POWER7 Virtualization with PowerVM

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, driving down costs and becoming the foundation for server pools and cloud computing technology. Therefore, the performance of virtualization 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

Today's POWER6 and POWER7 systems combine industry-leading performance, scalability and modularity to enable the most from an organization's investment and to build a flexible, responsive infrastructure that easily adapts and grows based on business needs. With a virtualization hypervisor built into every Power System, all performance benchmarks are achieved in a virtualized environment, unlike competitive systems that can be subject to lower performance when using third-party virtualization software.

POWER7 processor-based systems—the first generation of systems built for a smarter planet—offer balanced systems designs that automatically optimize workload performance and capacity at either a system or virtual machine level. Features include:

- TurboCore™ for maximum per core performance for databases
- MaxCore for industry-leading parallelization and maximum capacity throughput
- Intelligent threading technology to utilize more threads when workloads benefit
- Intelligent Cache technology to optimize cache utilization, flowing it from core to core
- Intelligent Energy that maximizes performance dynamically when thermal conditions allow
- Active Memory™ Expansion that dynamically provides more memory when needed

IBM® PowerVM™ in the POWER6 and POWER7 environment offers broader platform support, greater scalability, higher efficiency in resource utilization, and more flexibility and robust heterogeneous server management than ever before. IBM PowerVM has autonomic resource affinity resulting in higher workload performance in a virtualized environment. IBM POWER7 and PowerVM with its efficient virtualization are key to the success of cloud computing environments.

## 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 System 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.

### 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 1/100<sup>th</sup> 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.

### 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.

### I/O Virtualization

The Virtual I/O Server (VIOS) is an integral part of PowerVM, 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.

### 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.

### 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

### 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.

### Systems Management

IBM Systems Director Express, Standard and Enterprise Editions for Power Servers support the PowerVM environment and is the IBM tool for heterogeneous virtualization management of Power servers, as well as System x and System z. IBM Systems Director Editions support advanced management functions such as system discovery, workload lifecycle management, health monitoring, system updates, and topology mappings, as well as the ability to take action on defined event thresholds of monitored system components.

IBM Systems Director VMControl™ is a plug-in option that is included with the Systems Director Standard and Enterprise Editions, and 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 reducing the time and effort required to deploy workloads, increasing asset utilization, and enabling administrators to maintain high levels of availability with proactive monitoring and

collaborative troubleshooting.

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.

In the next section of the paper, the advantages of PowerVM will be discussed in greater detail.

## 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. POWER7™ high-end systems support up to 256 physical processors. and up to 1000 partitions.
- 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 and other third-party software must be installed on x86 hardware that leverages hardware-assist virtualization optimizations. In contrast, Power Systems servers implement 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 all while running with significantly less overhead.

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, Oracle VM 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 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. POWER7-based systems also have larger number of cores per socket, higher number of threads per core (SMT4), higher memory and more I/O bandwidth per core. This technology enables POWER7 systems to consolidate 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. IBM Power systems and PowerVM technologies allow organizations not only to consolidate but to do so at a higher consolidation ratio as IBM Power systems sustain performance at high utilization. To showcase PowerVM performance, IBM undertook a study to compare the performance of IBM PowerVM and VMware virtualization technologies that employs two industry-standard benchmarks. This study is highlighted in the next section of the paper.

## **Benchmark Comparison Study: PowerVM vs. VMware Virtualization Technologies**

Comparing virtualization technologies can be a complicated endeavor. As a starting point, this paper illustrates a simplified approach which can be built upon in the future. Two of the most common questions posed by those considering virtualization technologies are; 1) How efficient is the technology? and 2) How well does the technology scale?

To begin with, we propose perhaps the most simplified approach possible; an evaluation of how efficiently a single virtual machine can make use of a modern multi-core processor system. This is accomplished by configuring a single virtual machine, and then recording throughput and CPU utilization as a range of processing resources that are assigned to the VM. While a single VM isn't likely to be a common virtualization deployment choice, it does provide clear insight into the fundamentals of how a given virtualization technology performs, and it also gives insight into how that behavior might translate into more expansive deployments. Next, we look at what happens when the resources in a modern multi-core computer system are spread across multiple virtual machines, again tracking throughput and CPU utilization. As anyone familiar with these technologies would attest, a lot depends on the workload therefore multiple workloads are included in this study.

Two benchmarks AIM7 and DayTrader2.0 were selected for this study. AIM7, a well known open source benchmark, is used widely by UNIX computer system vendors to compare system performance. DayTrader2.0, which measures performance of Java™ Enterprise Edition (Java EE) technologies, is widely used to compare web application server performance by the industry. Note that IBM submitted the performance test plan to VMware prior to conducting the experiments.

AIM7 was recently run in IBM's performance lab comparing PowerVM and VMware performance. The AIM7 benchmark consists of a large set of atomic tests covering compute, memory and I/O intensive tests covering a wide range of operations such as numerical/matrix operations, memory allocate/move/copy, sort/search, and read/write. AIM7 has three predefined test suites each consisting of a mix of compute, memory and I/O intensive atomic tests with different weights namely compute, multi-user and database test suites. For this comparative study, the AIM7 compute server test suite was used.

## **AIM7 Single Virtual Machine Scaling (Scale-up Performance)**

This study used a POWER7-based IBM Power 750 Express system which has four sockets each with eight cores, and the HP DL370 with Intel® Xeon® 5570 which has two sockets, each with four cores. For the eight core comparison study between PowerVM and VMware, three sockets along with associated memory were disabled on the Power 750 system resulting in a eight core system (single chip system). In both cases the same guest OS SuSE 11 was used, x86\_64 version on VMware and Power Linux version on PowerVM virtual machine. The configuration details in each of the tests can be found in Appendix A.

The latest release of VMware vSphere 4.0 update1 was selected for the study. Even though the vSphere 4.0 update increased the limit of VCPUs per core from 20 to 25, it supported only up to eight virtual processors per virtual machine. Therefore a single VM scaling (scale-up) comparison between PowerVM and VMware was limited to eight virtual processors. The goal of this test was to measure the hypervisor efficiency in scaling virtual processors within a single virtual machine in isolation. However, it should be noted that PowerVM's virtualization technology is not restricted to only scaling eight virtual processors within a virtual machine (such as VMware's technology)

▮

To showcase POWER7 and PowerVM scale-up advantage over VMware, we also scaled up to 32 virtual processors within a single PowerVM virtual machine. Figure 1 depicts single Virtual Machine scaling test results of PowerVM and VMware vSphere 4.0 update 1. On PowerVM the scaling was linear from one virtual processor to 32 virtual processors and VMware results were also linear from one to eight virtual processors. The PowerVM absolute throughput (jobs/sec) surpassed the Intel Xeon 5570 and VMware results by up to 65% on the AIM7 benchmark on one, two, four and eight virtual processors. Using from eight to 32 virtual processors, PowerVM had linear scaling allowing a single virtual machine to use all the processor resources on the system.

The virtual machine configuration and the test results can be found in the tables following the graphs for each test in this paper, starting with Figure 1.

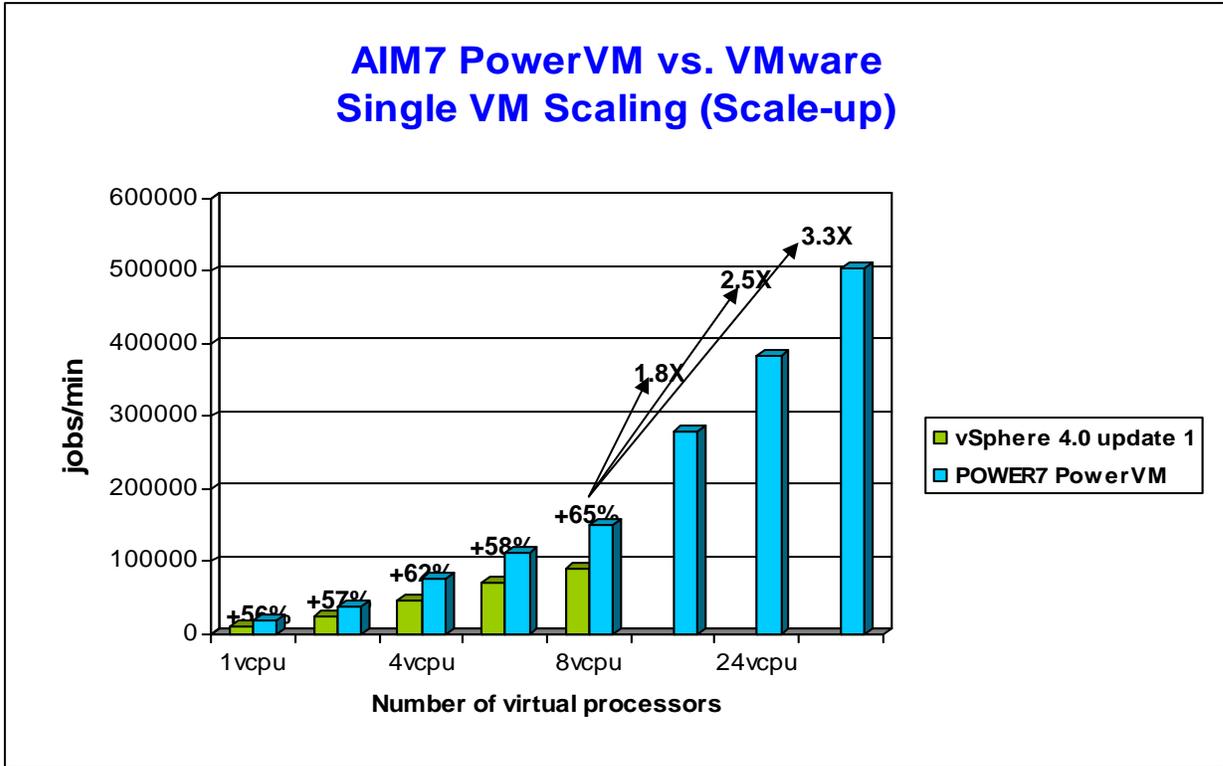


Figure1 AIM7 single VM scaling

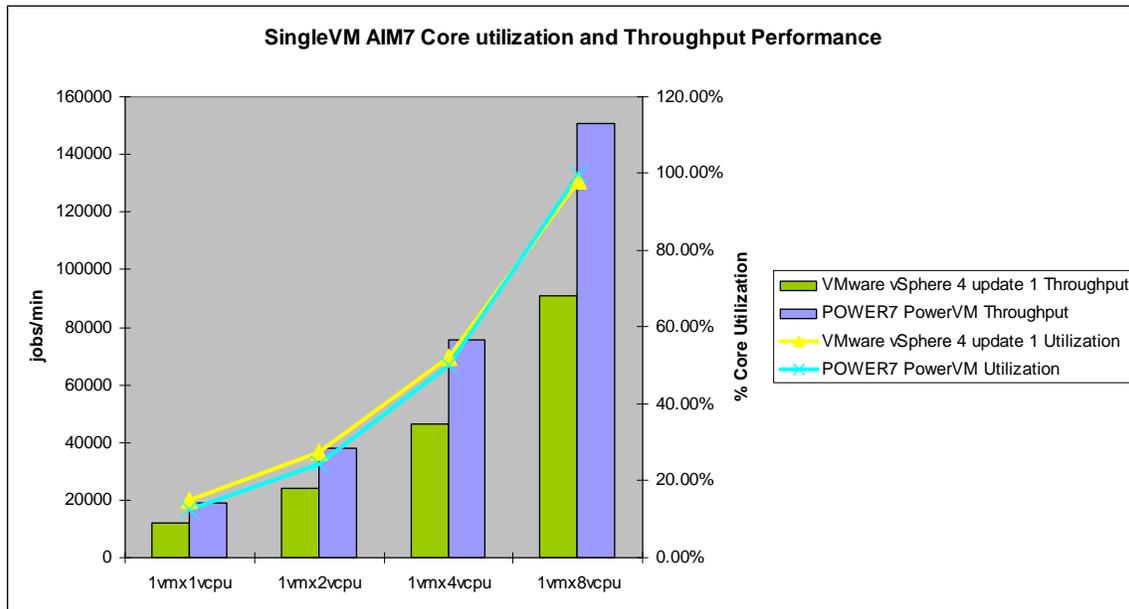
System Configuration for AIM7 Benchmark (Single VM)	Cores	Virtual CPUs	Jobs / min
IBM Power 750 3.5GHz DPSM mode, 4 sockets, 512 GB RAM, SMT4 enabled, PowerVM and SLES11 (Power Linux)			
Single Socket (8 cores), 128GB RAM	8	1	19084
Single Socket (8 cores), 128GB RAM	8	2	38007
Single Socket (8 cores), 128GB RAM	8	4	75587
Single Socket (8 cores), 128GB RAM	8	8	150558
4 sockets enabled with 512GB RAM	16	16	279817
4 sockets enabled with 512GB RAM	24	24	383632
4 sockets enabled with 512GB RAM	32	32	504020
HP DL370 Intel Xeon 5570 2.9 GHz 96GB RAM (HT and Turbo enabled in BIOS Intel VTx with EPT HW virtualization assist) VMware vSphere 4.0 update1 and SLES11 (GA, x86_64)			
2 Sockets (each with 4 cores), 96 GB	8	1	12170
2 Sockets (each with 4 cores), 96 GB	8	2	24164
2 Sockets (each with 4 cores), 96 GB	8	4	46595
2 Sockets (each with 4 cores), 96 GB	8	8	90910

Each AIM7 benchmark test was run to saturate the virtual processor utilization within the virtual machine in both PowerVM and VMware tests. The guest OS processor utilization metrics indicated close to 100% for each scale-up test. VMware esxtop (resource utilization collection tool) reported the core utilization of the virtual machine as shown in Figure 2 which parallels with PowerVM virtual machine utilization as plotted on y-axis. As virtual processors were scaled, CPU utilization was also scaled on both platforms. However, the throughput performance on PowerVM exceeds VMware at every single data point on the x-axis. Even though on both platforms the

virtual machines were configured with identical number of virtual processors, PowerVM leverages SMT4 even when the number of virtual processors (8) is lower than the number of logical processors (32) in the system. VMware on the other hand maps a logical processor to a virtual processor so at any point in time a VMware virtual machine configured with 8 virtual processors can only consume eight logical processors. As pointed out in VMware article [ref 16], “The best generalizations we can provide about HT on ESX are: 1. Until you have more vCPUs requesting processing power than there are physical cores, HT cannot hurt and provides no value. 2. Once you have more vCPUs requesting CPU than physical cores on the system, HT usually provides small Gains.” In HT enabled mode HP DL370 has 16 logical processors, since VMware restricts the number of virtual processors supported in a virtual machine, they need 2 virtual machines, each with 8 virtual processors to utilize all 16 logical processors in HT enabled mode. PowerVM allocates all the logical processors of a core to a virtual processor therefore it is able to utilize all the cores in a single virtual machine to help take advantage of all the available resources (leverages SMT4) for more efficient performance.

One could argue that the purpose of virtualization is not to run a single VM on a system, and when multiple VMs are consolidated VMware would consume all the cycles in the system. That is indeed a valid argument and such a scenario is covered later in this paper. However even in an environment where multiple VM's are consolidated, a single virtual machine could peak at times and could use most of the cycles in a system if other virtual machines are idle. So scaling-up virtual processors within a single virtual machine is also critical for workload consolidation and virtualization. ***By restricting the number of virtual processors in a single virtual machine, VMware is not leveraging hyper-threading technology of HP DL370 system.***

IBM POWER7 and PowerVM not only scaled linearly in a singleVM on this benchmark but could scale within a single VM leveraging all the processor cycles in the entire system if needed. The Power hypervisor efficiently dispatched processor cycles to the configured virtual processors that executed the workload resulting in higher throughput performance. The mapping of physical to virtual processors was done in whole core units by allocating all the SMT threads to a virtual machine. When there was a one-to-one ratio between physical processors and the configured virtual processors, PowerVM mapped a whole physical processor (including all the logical processors of a core) to a virtual processor. When there was a one-to-many ratio between physical processor and the configured virtual processors, the physical processor was shared among the multiple virtual processors. From the VMware results, it is obvious that VMware did not leverage all the logical processors of a core when there was a one-to-one ratio between available physical cores and configured virtual processors as shown in singleVM configuration. VMware would [need 2 virtual machines, each with 8 virtual processors to utilize all 16 logical processors available in HP DL370 system.](#)



**Figure 2: AIM7 Single VM Processor Utilization**

The results of two virtual machines each with 8 virtual processors indicated that PowerVM still beat VMware throughput performance in 2 virtual machine tests as shown in Figure 3 where both platforms virtual machines were consuming all system capacity (all available logical processors).

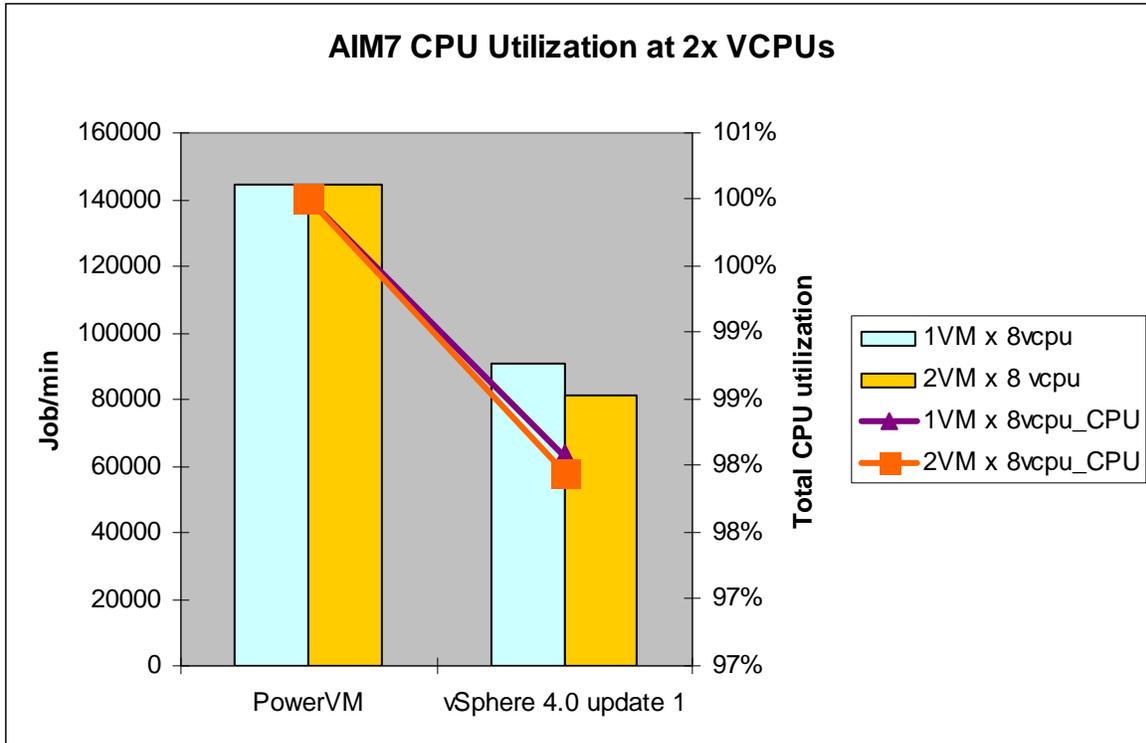
PowerVM guest OS tools report virtual processor utilization along with accurate physical processor utilization within a virtual machine. ***This accuracy is a result of tight integration of the guest OS, PowerVM and Power System hardware by IBM.*** In addition PowerVM also has the virtualized hardware event monitor which facilitates analysis of a virtual machine (LPAR) performance at CPI (cycles per instruction) level. Whereas in VMware virtualization environment, guest OS tools are not integrated with the underlying virtualization technology so the reported resource utilization does not show both the virtual processor and actual physical processor utilization within a virtual machine. The hardware monitoring tools will not produce valid results running in a guest OS on VMware as the hardware counters on the system are not exposed by VMware.

Note that the two tests with x86 hyper-threading enabled (Figure 3) were also repeated with hyper threading disabled on VMware. Without hyper-threading the throughput performance of 1 VM improved by 1% and 2 VM by 8% while the CPU utilization in both cases reached ~97% and ~99% respectively. The 2VM case where 16 virtual processors were used, the improved results in HT\_disabled mode led to the conclusion that ***VMware was not efficient in handling resources in Hyper Threading mode on this benchmark when one-to-one mapping between physical and virtual processors was configured.***

#### **AIM7 Multiple Virtual Machine scaling:**

Both the Power 750 and HP DL370 systems were configured to have equal number of cores (eight) for this test. That is, Model 750 was put in a single socket, 128 GB RAM configuration. The scenario in this section concentrated on the dimension of virtual machine scaling. By doing so, the study focused on the aggregate performance of the system as virtual machines were added from one to eight. That is, as virtual machines were added, the configuration of each virtual machine was adjusted so that the total virtual processors remained at 8 in each configuration.

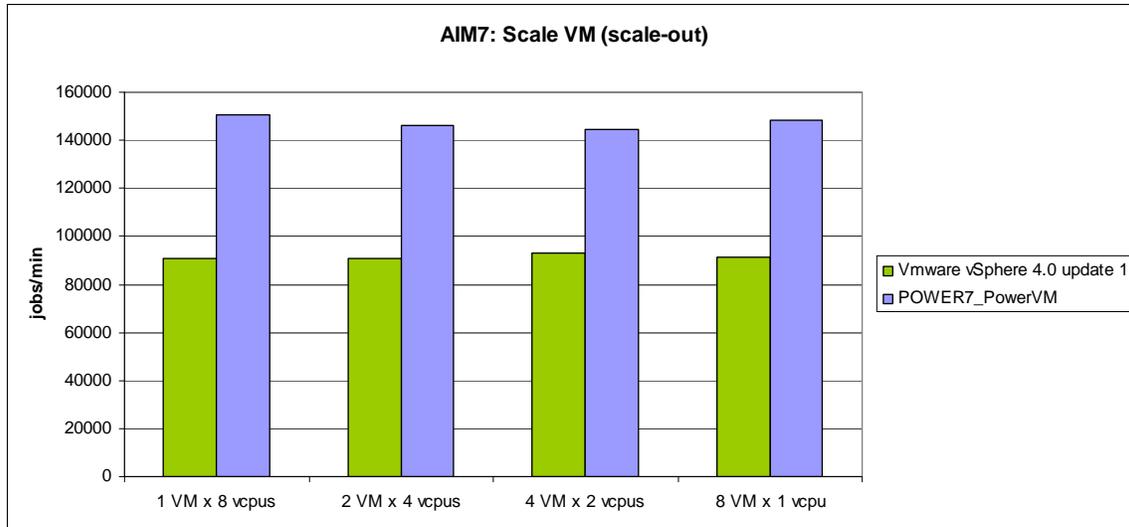
One virtual machine used 8 virtual processors, 2 virtual machines used 4 virtual processors each and so on. This scale-out effectively measured the throughput performance impact on different consolidation ratio. Both throughput performance and CPU utilization were measured for each test and the results primarily showed the difference between a scale-up (one VM) and scale-out (2 VM, 4 VM, 8 VM) performance on both virtualization technology and platforms.



**Figure 3: AIM7 CPU Utilization at 1x and 2x virtual processors**

PowerVM delivers the ability to allocate processor resources in dedicated or in shared mode. This flexibility affords customers options to best optimize the virtualization deployment for their unique situation. IBM offers this choice because there are different advantages between the shared and dedicated options and a 'one size fits all approach' can impose unnecessary limitations. As the virtualization configuration complexity was increased, with many virtual machines necessitating processor sharing and over commit, the shared mode was employed.

As more virtual machines were added, the shared mode processor management had a small impact on throughput performance compared to dedicated mode. Both PowerVM and VMware sustained throughput performance when virtual machines were scaled. However PowerVM absolute throughput performance was ~65% better than VMware at every scale-out experiment as shown in Figure 4. As discussed in the prior section, PowerVM used all the available compute capacity in SMT4 mode towards executing the workloads at every single scale-out test.



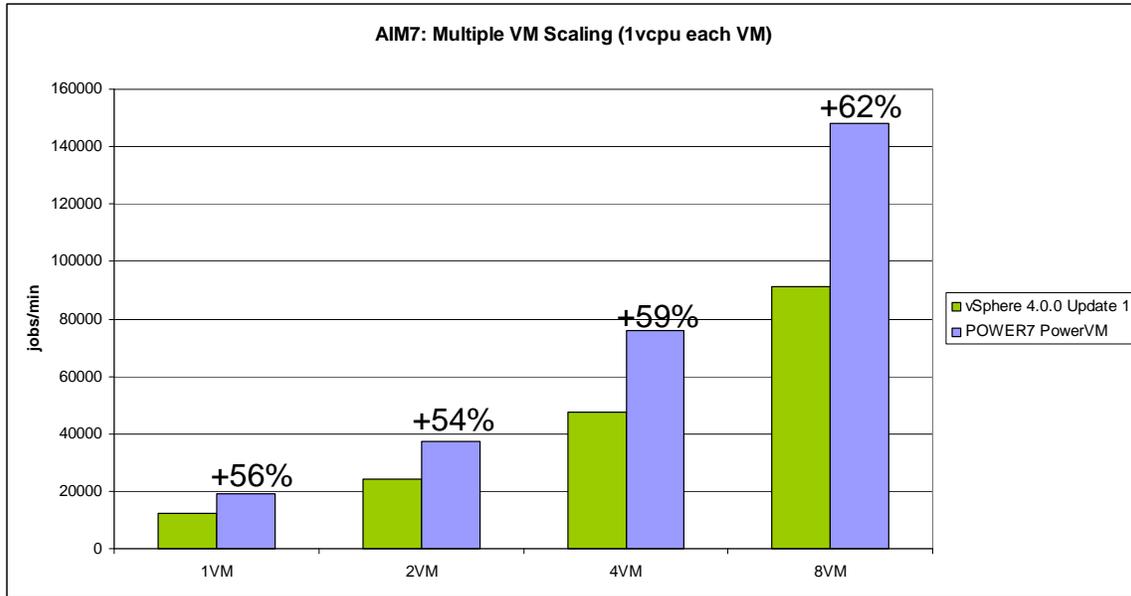
**Figure 4 AIM7 Virtual Machine Scale-out Performance**

#### **AIM7 Multiple Virtual Machine and Virtual Processors scaling:**

The next set of tests involved concurrently scaling both virtual machines and virtual processors. Eight identically configured virtual machines were created for this two dimensional scaling (scale-out performance) test. The virtual machines one to eight were booted one at a time to accommodate 1VM through 8VM tests. That is in the case of 1VM test, only one VM was booted and in the case of 8VM test, all eight VMs were brought on-line. There were three sets of tests conducted in this category by varying (1, 2 and 4) virtual processors in each set of tests. As virtual processors were scaled along with virtual machines, processor resources were over-committed. Processor over-committing in a consolidated environment is a common practice in a virtualized environment to maximize utilization while meeting on-demand peak load. In a processor over-commit environment, physical processor cycles are being time-sliced and used for different virtual processors.

The scaling test was started with one virtual machine followed by adding one VM at a time to scale up to 8VM. Identical load factor (minimum and maximum) for the benchmark was used to run on each set (1, 2, 4 and 8) of VMs concurrently. Throughput and processor utilization were measured in each test case. As the first set of scaling of one to eight VMs each with one virtual processor was completed, the virtual machines were reconfigured to increase the virtual processors to two and four respectively and repeated the scaling 1 to 8VM tests with two virtual processors and four virtual processors configuration. In this scenario, both virtual machines and virtual processors were scaled resulting in a mix of non-over commit and over commit of physical processors.

At one virtual processor (Figure 5) with the maximum of up to eight virtual machines, the CPU resources were not over committed (1VCPU x 8VM), but as virtual processors were added the over commit levels increased as the number of virtual machines were scaled to stress the system up to 4x over commit levels (4VCPU x 8VM). At both processor non over-commit and over-commit levels, IBM PowerVM surpassed the VMware AIM7 throughput performance at each test. Figure 5 demonstrates PowerVM superior performance when one virtual processor is mapped to one physical core (non over commit of processor resources).



**Figure 5 AIM7 Multiple VM Scaling with 1 virtual processor**

System Configuration for AIM7 Benchmark (Multiple VM Scaling with 1 vcpu)	Cores	Total Virtual CPUs	Jobs / min
IBM Power 750 3.5GHz DPSM mode, Single socket (8 cores), 128 GB RAM, SMT4 enabled, PowerVM and SLES11 (Power Linux)	8		
1 VM		1	19027
2 VM		2	37330
4 VM		4	76004
8 VM		8	13149.8
HP DL370 Intel Xeon 5570 2.9 GHz 96GB RAM (HT and Turbo enabled in BIOS Intel VTx with EPT HW virtualization assist) VMware vSphere 4.0 update1 and SLES11 (GA, x86_64)	8		
1 VM		1	12170
2 VM		2	24125
4 VM		4	47638
8 VM		8	91162

In Figure 6, at two virtual processors per virtual machine configuration, PowerVM leveraged all the processing cycles of 8 cores in SMT4 mode (32 logical processors mapped to 8 virtual processors) at 4 virtual machines and sustained the throughput as virtual machines were scaled to eight resulting in mapping the same 32 logical processors onto 16 virtual processors (2X over commit). VMware on the other hand did not leverage hyper-threading efficiently at 4 virtual machines (using 8 out of 16 logical processors) and the throughput improved by 9% at 8 virtual machines (using 16 logical processors). So in this test set, VMware did not over commit processor resources. The disparity in when PowerVM and VMWare resources began to over-commit is attributed to the differences in their hypervisor technology. PowerVM allocates whole core (all 4 logical processors in SMT4 mode) to a virtual processor therefore the over commit starts earlier than VMware as its technology maps one logical processor (not whole core in HT mode) to a virtual processor.

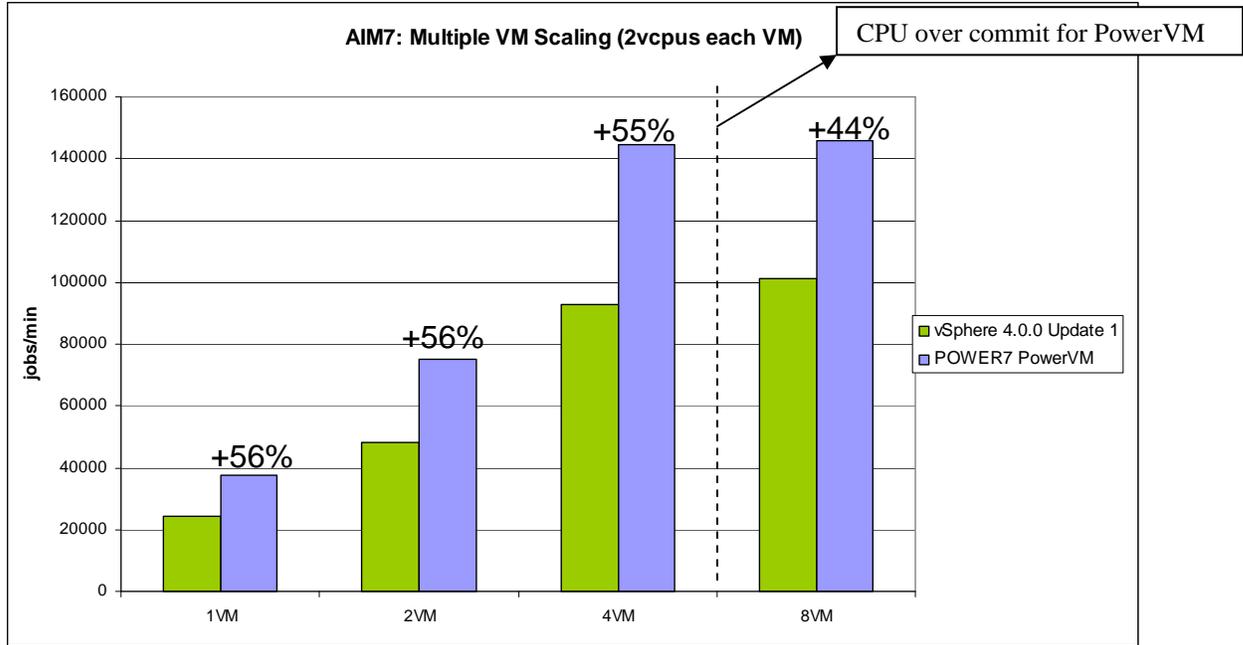


Figure 6 AIM7 Multiple VM Scaling with 2 virtual processors per VM

System Configuration for AIM7 Benchmark(Multiple VM Scaling with 2 vcpus)	Cores	Total Virtual CPUs	Jobs / min
IBM Power 750 3.5GHz DPSM mode, Single socket (8 cores), 128 GB RAM, SMT4 enabled, PowerVM and SLES11 (Power Linux)	8		
1 VM		2	37751
2 VM		4	75240.4
4 VM		8	144545.6
8 VM		16	145993
HP DL370 Intel Xeon 5570 2.9 GHz 96GB RAM (HT and Turbo enabled in BIOS Intel VTx with EPT HW virtualization assist) VMware vSphere 4.0 update1 and SLES11 (GA, x86_64)	8		
1 VM		2	24145
2 VM		4	48115
4 VM		8	92739
8 VM		16	101354

Figure 7 shows a four virtual processor configuration scenario, where PowerVM leveraged all the eight cores (32 logical processors) even at 2 virtual machines and then continued performance well at processor over commit levels (2x and 4x) at 4 virtual machines and 8 virtual machines respectively. In the 4 virtual machine test, the Power 750 system, using its “Intelligent Energy Optimization” mode, increased the system frequency as the system reached high utilization. The frequency reached close to ~3.8GHz improving the throughput by 9%. Overall the throughput at over commit levels improved and was sustained at 8 virtual machines as over commit levels doubled. Processor over-committing is frequently done in virtualized environments to help efficiently use and maximize resource utilization. Comparing PowerVM’s ability to overcommit resources vs. VMWare confirms VMWare is not efficient in overcommitting and PowerVM performs up to 61% better.

Figure 8 shows the actual processor utilization in this scenario across both platforms. Even though VMware esxtop core utilization matches with PowerVM core utilization curve, there is a difference in how the number of logical processors being dispatched onto virtual processors

across each test. At 4VM case VMware dispatched all 16 logical processors and went into over commit mode at 8VM. So depending on number of virtual machines active at a time, the underlying hardware hyper-threading capability may or may not be exploited. Nevertheless, VMware was still unable to match PowerVM's throughput performance even as VMware leveraged hyper-threading technology.

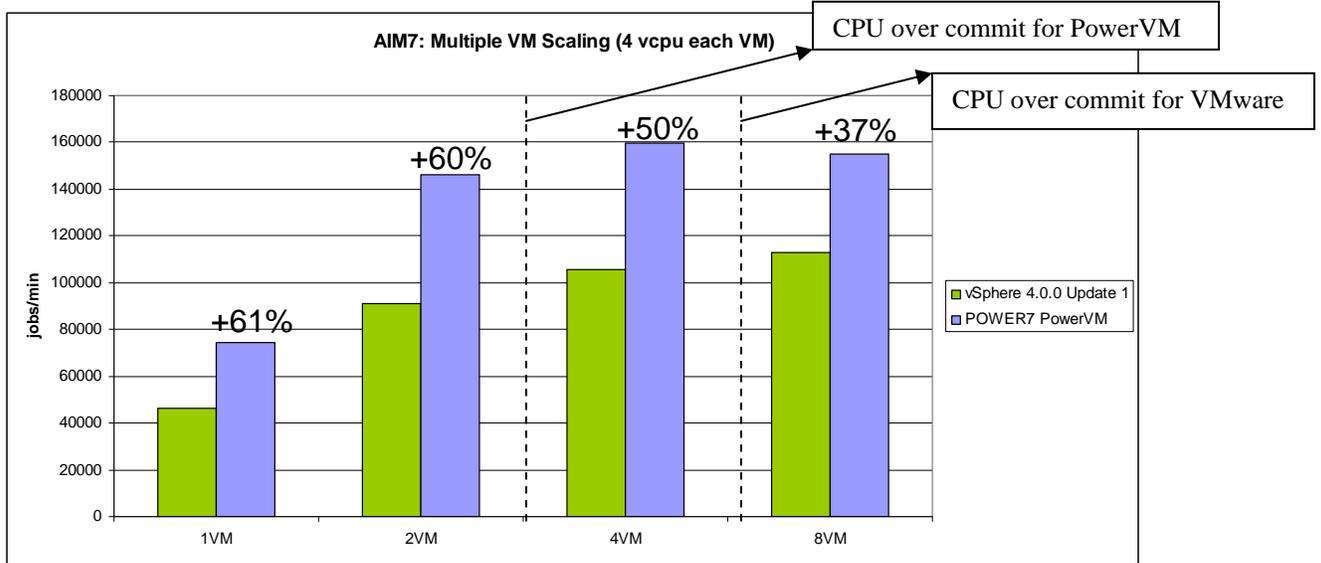
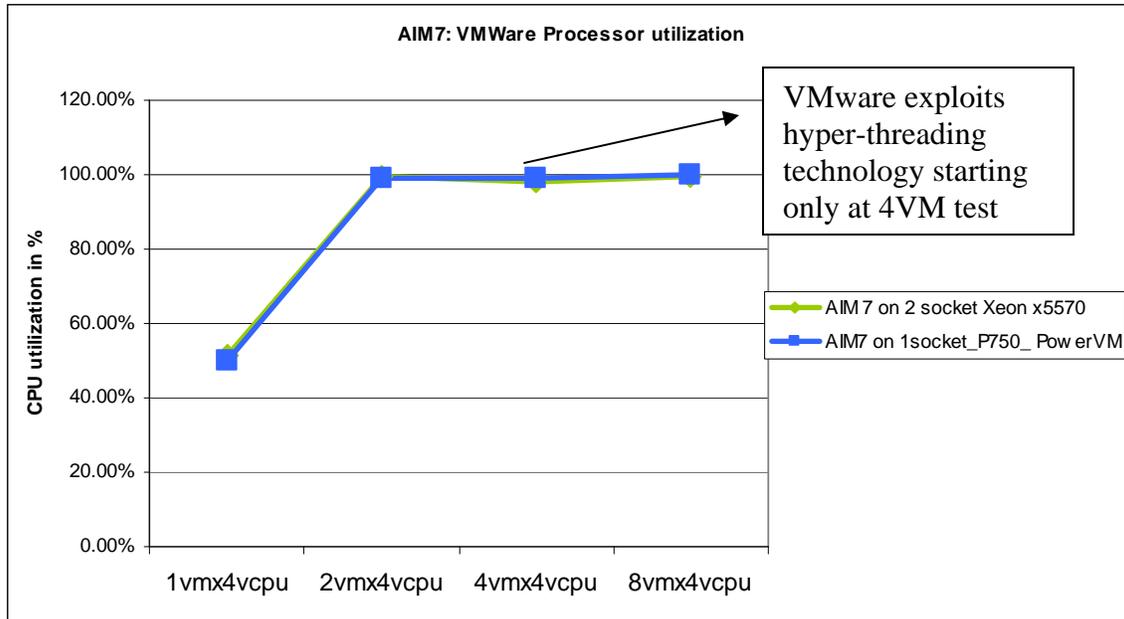


Figure 7 AIM7 Multiple VM Scaling with 4 virtual processors

System Configuration for AIM7 Benchmark (Multiple VM Scaling with 4 vcpus)	Cores	Total # of VCPUs	Jobs / min
IBM Power 750 3.5GHz DPSM mode, Single socket (8 cores) 128 GB RAM, SMT4 enabled, PowerVM and SLES11 (Power Linux)	8		
1 VM		4	74624
2 VM		8	146007.3
4 VM		16	159668
8 VM		32	155061.8
HP DL370 Intel Xeon 5570 2.9 GHz 96GB RAM (HT and Turbo enabled in BIOS Intel VTx with EPT HW virtualization assist) VMware vSphere 4.0 update1 and SLES11 (GA, x86_64)	8		
1 VM		4	46136
2 VM		8	90832
4 VM		16	105834
8 VM		32	112661



**Figure 8: AIM7 Benchmark Processor Utilization**

### DayTrader 2.0 Benchmark

DayTrader2.0, an open source benchmark, was used to compare POWER7 and PowerVM performance with Intel Xeon 5570 processors and VMware vSphere 4.0 update 1. DayTrader 2.0 is a stock trading benchmark built around the paradigm of an online stock trading system.

The benchmark was configured in three-tier mode where the back-end database was installed on a separate system and the system under test was the middle-tier hosting the web application server. The client workload was generated using JIBE drivers on a separate IBM System x<sup>®</sup> system. For both platform tests the same back-end database server and the client (JIBE drivers) systems were used. The middle tier Web Application Server was run on a POWER7-based IBM Power 750 Express server for the PowerVM study and an HP DL370 Intel Xeon 5570 server for the VMware study. The guest OS used in virtual machine for PowerVM was AIX 6.1 TL 04 and for VMware the guest OS was RHEL 5.4 version.

IBM DB2<sup>®</sup> v9.7 was installed on the back-end database server and IBM WebSphere<sup>®</sup> V7.0.07 was installed on the Web Application Tier system (SUT). Further details on the hardware, software, and the test plan details can be found in Appendix A.

In the case of a single virtual machine (scale-up) experiment, as virtual processors were scaled, additional Web Application Server instances along with the Daytrader2.0 applications were added to get the best performance in both the platforms. Where multiple Web Application Server instances were used they were independent instances, administered separately, and executing independently with no load balancing performed across the Daytrader2.0 applications across multiple Web Application Server instances within a virtual machine. There were a set of tests performed in scale-up configuration: 1) One Web Application Server instance per virtual processor and 2) One Web Application Server instance per two virtual processors. The best performing throughput was selected for both platforms for comparison.

Daytrader2.0 was built on a core set of Java EE technologies that includes Java Servlets and JavaServer Pages (JSPs) for the presentation layer and Java database connectivity (JDBC), Java Message Service (JMS), Enterprise JavaBeans (EJBs) and Message-Driven Beans for the back-end business logic. There are two modes that could be used to run the benchmark, namely 1) EJB mode and 2) JDBC mode. The Daytrader2.0 architecture (Figure 9) shows the transaction processing flow for each of these two modes. In JDBC mode, the transaction bypasses EJB containers and accesses the database directly.

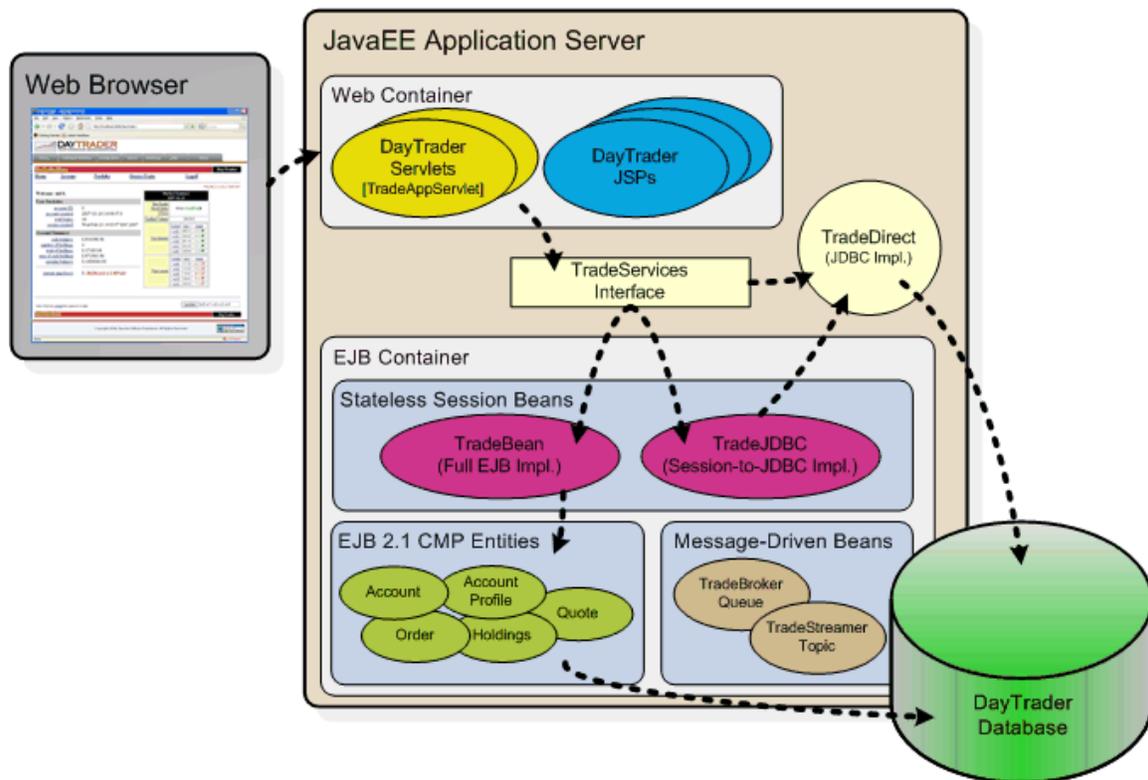


Figure 9 DayTrader2.0 Architecture

### DayTrader2.0 Single Virtual Machine Scaling:

Single VM scaling tests were performed in both EJB and JDBC mode of the Daytrader2.0 benchmark. Each DayTrader2.0 test was conducted in four phases; each phase was run for three minutes. The first and second phases were warm up phases and the benchmark results were collected over the third and fourth phases. Results were determined by measuring the transaction throughput on each platform for comparison. The optimal transaction throughput was created by using the number of clients that resulted in consistent response times with a high CPU utilization (~90 to ~100% based on guest OS CPU metrics).

The EJB mode results as shown in Figure 10 clearly indicate that PowerVM throughput performance on this benchmark is far better than VMware at each 1, 2, 4 and 8 virtual processors configuration. The throughput scaling performance for 2, 4 and 8 virtual processors was 1.8x, 3.4x, 6.7x on PowerVM while the throughput scaling performance on VMware was 1.9x, 3.2x and 4.8x. Both platforms had equal “relative scaling” at two and four virtual processors; however only PowerVM had linear scaling from four to eight virtual processors. Since this workload was network intensive, four gigabit Ethernet ports were used on the SUT (both platforms) to ensure no network bottleneck occurred during the benchmark run. The VMware results shown in Figure 10 and the relative performance shown in the table following Figure 10 were better than what had been published by HP on a VMware/Intel Xeon 5570 system [9].

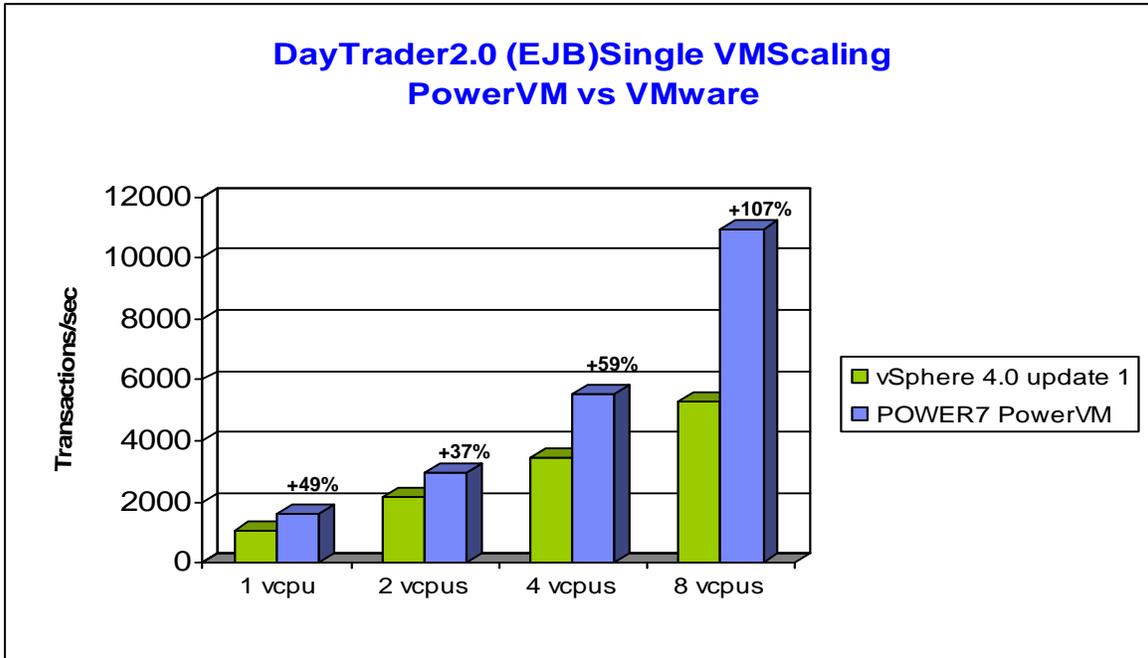


Figure 10 DayTrader2.0 Single VM Scaling (EJB Mode)

System Configuration for DayTrader 2.0 EJB Mode	Cores	WebSphere Instances	# of vcpus	Transactions/sec	Relative Scaling	Response Time
<b>Single VM scaling</b>						
IBM Power 750 3.5 GHz, Single socket (8 cores) , 128GB RAM, AIX v6.1 TL04 SMT4 enabled, IBM WebSphere v7.0.0.7, PowerVM	8	1	1	1630	100%	0.021
IBM Power 570 4.7GHz 64GB RAM, AIX v6.1 TL03 SMT2 enabled, DB2 v9.7, JDBC		1	2	2990	183%	0.011
4 JIBE Drivers		4	4	5536	339%	0.013
		8	8	10951	671%	0.010
HP DL370 Intel Xeon 5570 2.9 GHz 96GB RAM (HT and Turbo enabled in BIOS Intel VTx with EPT HW virtualization assist) RHEL 5.4, IBM WebSphere v7.0.0.7, VMware vSphere 4.0 update 1	8	1	1	1088	100%	0.032
IBM Power 570 4.7GHz 128GB RAM, AIX v6.1 TL03 SMT enabled, DB2 v9.7, JDBC		1	2	2175	199%	0.016
4 JIBE Drivers		2*	4	3467	318%	0.017
		4*	8	5286	485%	0.027

\*Results with 4 and 8 Web Application Server instances are lower than 2 and 4 Web Application Server instances

Similar results were seen when the benchmark was run in JDBC mode (Figure 11). PowerVM had significantly higher throughput performance compared to VMware at one through eight virtual processors. As was the case for EJB mode, PowerVM had superior scale-up performance within a single SMT virtual machine for JDBC modes on the DayTrader2.0 benchmark.

The guest OS CPU utilization metrics showed high (~90 to ~95%) CPU utilization, however the VMware esxtop which reported reliable resource utilization showed ~98% core utilization at 8 virtual processor configuration as shown in Figure 12. Again similar to AIM7 single virtual machine tests covered in earlier section, VMware was not leveraging hyper-threading efficiently in these tests (only 8 out of 16 logical processors are dispatched at any moment in its execution time).

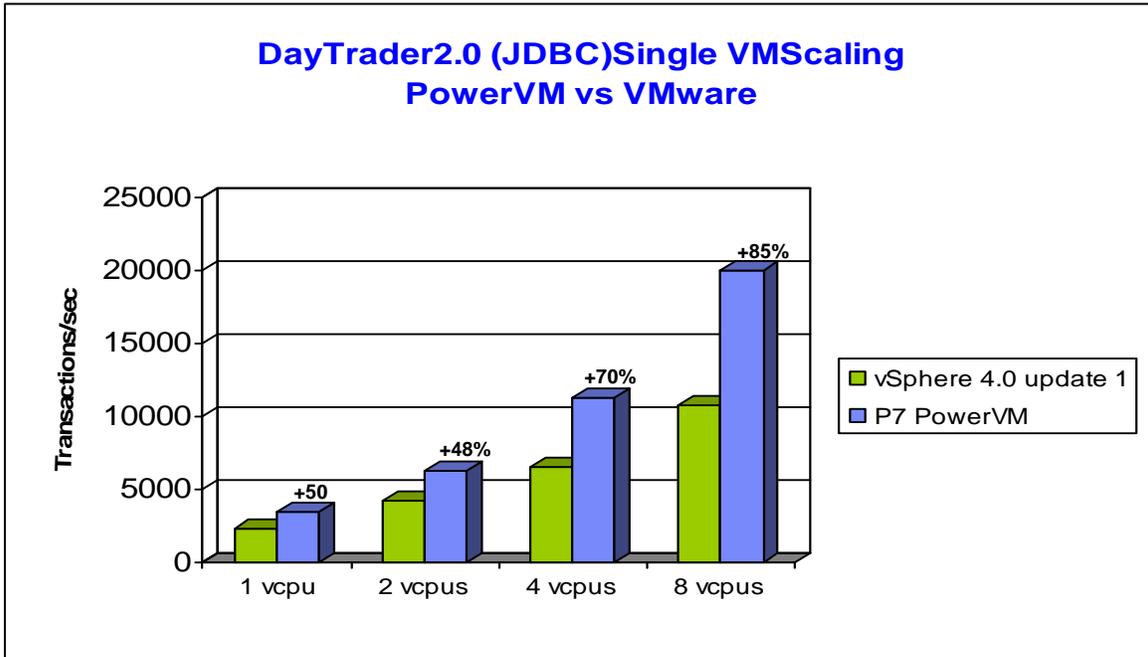
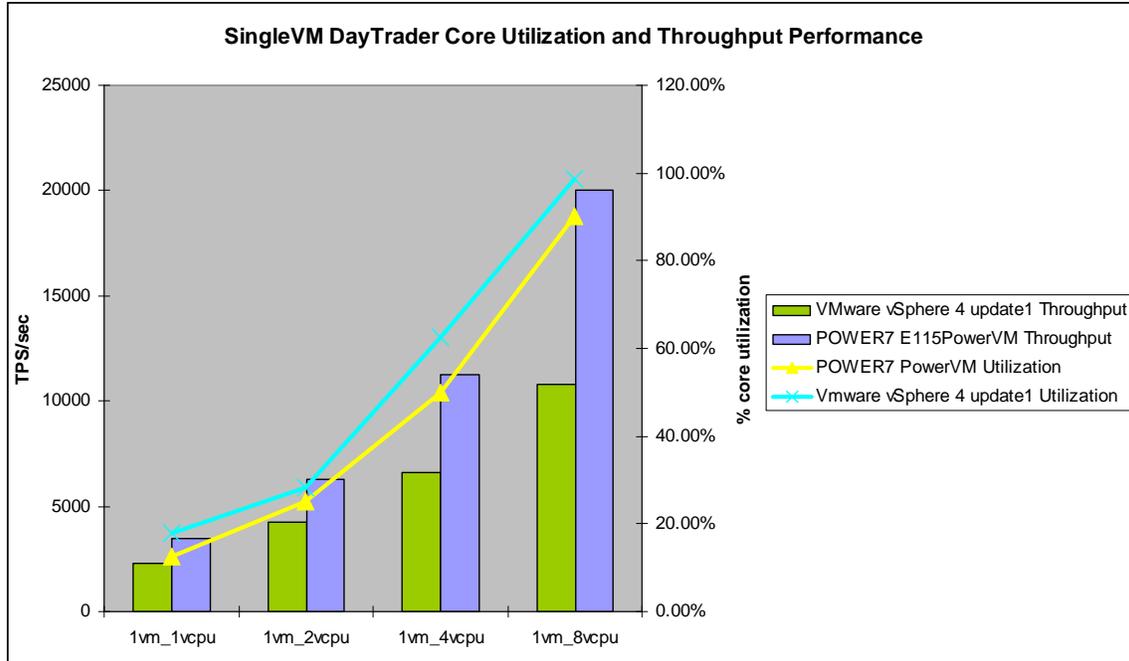


Figure 11: DayTrader2.0 Single VM Scaling (JDBC Mode)

System Configuration for DayTrader 2.0 JDBC Mode	Cores	Websphere Instances	# of vcpus	Transactions/sec	Relative Scaling	Response Time
<b>Single VM scaling</b>						
IBM Power 750 3.5 GHz, single socket (8 cores), 128GB RAM, AIX v6.1 TL 04 SMT4 enabled, IBM WebSphere v7.0.0.7, PowerVM	8	1	1	3467	100%	0.010
IBM Power 570 4.7Ghz 64GB RAM, AIX v6.1 TL03 SMT2 enabled, DB2 v9.7, JDBC		1	2	6257	180%	0.005
4 JIBE Drivers		2	4	11273	325%	0.006
		8	8	20024	577%	0.007
HP DL370 Intel Xeon 5570 2.9 GHz 96GB RAM (HT and Turbo enabled in BIOS Intel VTx with EPT HW virtualization assist) RHEL 5.4, IBM WebSphere v7.0.0.7, VMware vSphere 4.0 update 1	8	1	1	2306	100%	0.015
IBM Power 570 4.7Ghz 128GB RAM, AIX v6.1 TL03 SMT enabled, DB2 v9.7, JDBC		1*	2	4226	183%	0.008
4 JIBE Drivers		4	4	6616	286%	0.012
		4*	8	10775	467%	0.013



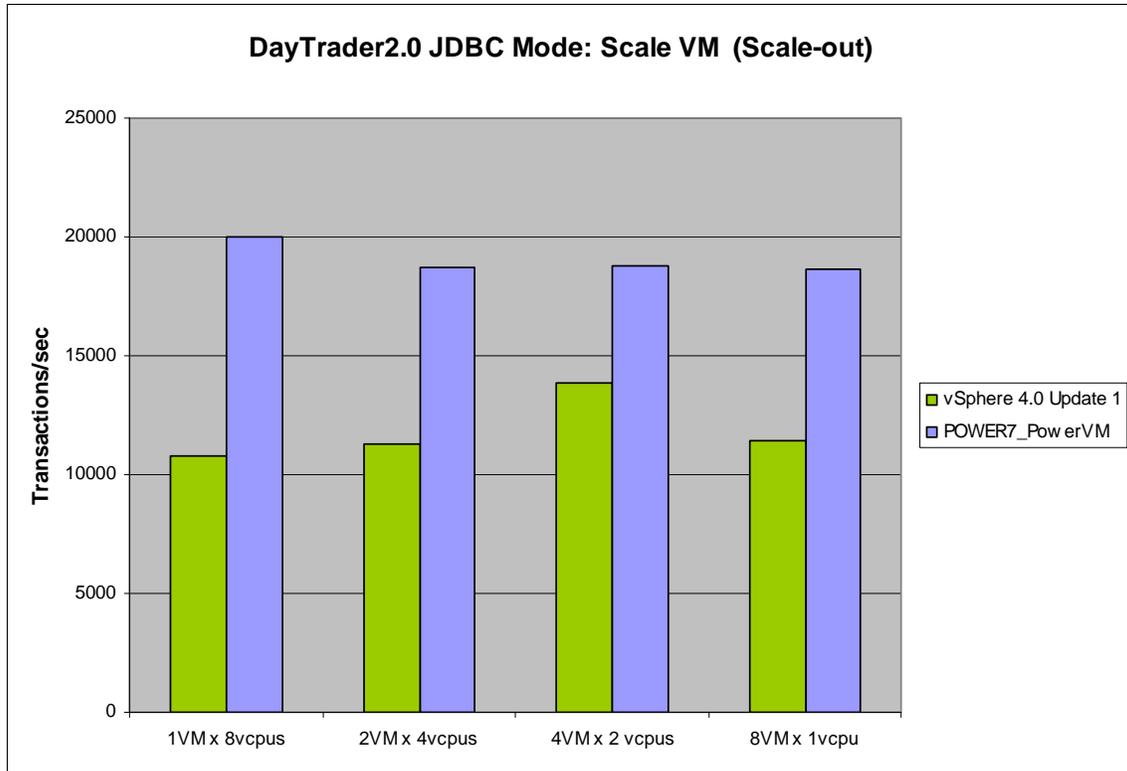
**Figure 12: DayTrader2.0 Processor Utilization**

### DayTrader2.0 Multiple Virtual Machine Scaling (Scale-out Performance)

In this scenario scale-out performance was measured running the DayTrader2.0 benchmark. This section covers two types of scale-out performance; the first category measures how well a system performed as more virtual machines were added splitting the processor resources resulting in more granular virtual machines. This test measures the optimum number of virtual machines in a given system. In the second category where virtual machines were scaled again one to eight, all eight virtual machines were configured equally but only the required virtual machines (1, 2, 4 and 8) were booted for performance evaluation.

In the first category of scale-out performance where multiple virtual processors were configured, multiple WebSphere instances were created along with the DayTrader2.0 application. In the second category all virtual machines were created equal and each was installed at the same level of guest OS, one instance of WebSphere and the DayTrader2.0 application.

The first category of scale-out tests results are shown in Figure 13 where each test set (1VM, 2VM, 4VM and 8VM) was configured to use eight virtual processors. As discussed earlier in the paper, PowerVM dedicated processor mode was used in 1VM\_8vcpu while multiple virtual machines tests (2VM, 4VM and 8VM) used shared processor mode. As the virtual machine scaled there was a small drop in throughput; however the throughput sustained consistently as more virtual machines were added. PowerVM surpassed VMware in every test in this category in throughput performance. VMware seemed to perform better in the 4VM case compared to other tests. There was no scaling pattern found in this test.



**Figure 13: DayTrader 2.0 Scale-out Performance**

System Configuration for DayTrader 2.0 JDBC Mode	Cores	# of VMs	# of vcpus/VM	Transactions/sec	Response Time
<b>Multiple VM Scaling (1vcpu per VM)</b>					
IBM Power 750 3.5 GHz single socket 8 cores, 128GB RAM, AIX v6.1 TL04 SMT4 enabled, IBM WebSphere v7.0.0.7, PowerVM	8	1	8	20024	0.007
IBM Power 570 4.7Ghz 64GB RAM, AIX v6.1 TL03 SMT2 enabled, DB2 v9.7, JDBC		2	4	18696	
4 JIBE Drivers		4	2	18787	
		8	1	18639	
HP DL370 Intel Xeon 5570 2.9 GHz 96GB RAM (HT and Turbo enabled in BIOS Intel VTx with EPT HW virtualization assist) RHEL 5.4, IBM WebSphere v7.0.0.7, VMware vSphere 4.0 update1	8	1	8	10775	0.013
IBM Power 570 4.7Ghz 128GB RAM, AIX v6.1 TL03 SMT enabled, DB2 v9.7, JDBC		2	4	11262	
4 JIBE Drivers		4	2	13851	
		8	1	11436	

The second category scaling virtual machines one at a time showed linear scaling on POWER7 and PowerVM surpassing VMware in absolute throughput performance by 36% to 63% as seen in Figure 14. VMware had relative scaling (2x, 3.6x) up to 4VM and fell short (4.9x) at 8VM scaling.

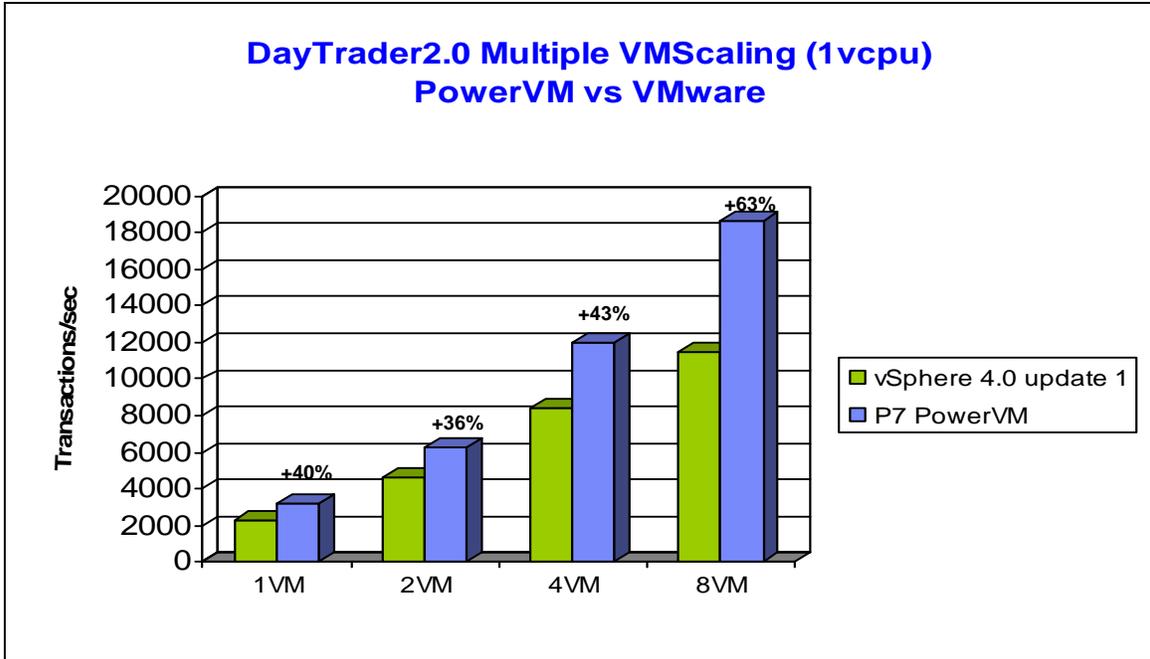


Figure 14: Multiple VM Scaling

System Configuration for DayTrader 2.0 JDBC Mode Multiple VM Scaling (1vcpu per VM) with 1 WAS instance per VM	Cores	# of VMs	# of vcpus	Transactions/ sec	Relative Performance	Response Time
IBM Power 750 3.5 GHz single socket 8 cores, 128GB RAM, AIX v6.1 TL 04 SMT4 enabled, IBM WebSphere v7.0.0.7, PowerVM	8	1	1	3227	100%	0.010
IBM Power 570 4.7Ghz 64GB RAM, AIX v6.1 TL03 SMT2 enabled, DB2 v9.7, JDBC		2	2	6306	195%	
4 JIBE Drivers		4	4	11945	370%	
		8	8	18639	577%	
HP DL370 Intel Xeon 5570 2.9 GHz 96GB RAM (HT and Turbo enabled in BIOS Intel VTx with EPT HW virtualization assist) RHEL 5.4, IBM WebSphere v7.0.0.7, VMware vSphere 4.0 update1	8	1	1	2306	100%	0.015
IBM Power 570 4.7Ghz 128GB RAM, AIX v6.1 TL03 SMT enabled, DB2 v9.7, JDBC		2	2	4612	200%	
4 JIBE Drivers		4	4	8324	360%	
		8	8	11436	495%	

## 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. For every scenario covered in this paper, IBM POWER7 and PowerVM had superior performance, higher efficiency in using system capacity, and better linear scaling within a single virtual machine and across multiple virtual machines compared to VMware virtualization technology on Intel Xeon 5570 processors. IBM Power Systems with PowerVM superior performance and features such as reliability, security, high availability, and resiliency, are well positioned for Cloud computing and for Smarter Planet solutions today and in the future.

In summary this competitive study has shown that IBM PowerVM has

- Linear scaling within a single virtual machine
- Leveraged maximum configured processor capacity efficiently
- Higher throughput performance for both AIM7 and DayTrader2.0
- Superior virtual machine and virtual processor scaling
- Sustained performance in processor over-commit mode
- Provided accurate accounting of resource usage within a virtual machine
- Supported Guest OS tools as-is without compromising its results
- Tighter integration across system, hypervisor, guest OS
- Overall better performance than VMware vSphere 4.0 update 1

## Appendix A

### Test Methodology used for comparing PowerVM to VMware application benchmark Performance

The performance evaluation would characterize efficiency of virtualization manager (hypervisor) and virtual machine scalability.

The following experiments were conducted at the IBM facilities, using HP DL370 G6 and POWER7™ Systems:

1) Examined the effect that adding virtual processors incrementally had on throughput performance in a single Virtual Machine.

**Note:** This test methodology shows how virtual (and physical) processors scale within a virtual machine.

2) Examined the effect that adding virtual machines had on throughput performance. Throughput was monitored as the number of virtual machines was scaled from 1 to 8. Throughput in each virtual machine was also evaluated with varying numbers (1, 2 and 4) of virtual processors.

**Note:** This test methodology shows the effect of multiple virtual machines running on a system in a non-over commit and an over commit resource environment.

Two different workloads, one compute-intensive and another covering Java, were used for the study as these two benchmarks run on multiple operating systems and platforms. Each of the above (1 and 2) tests were run using the same workload (homogeneous) on each virtual machine concurrently.

To ensure fair comparison across platforms, and remove variability across each set of tests, the following actions were taken:

1. Deployment of similar VM configurations in terms of virtual processors and memory allocated per VM.
2. The same set of “benchmark parameters” used across platforms
3. Except for the virtualized workloads under investigation, the remainder of the hardware and software are common across the two platforms (DB server, JIBE clients etc.)
4. Tuning was performed based on best practices of respective platforms
  - VMware vSphere 4.0 update 1, RHEL 5.4, AIX 6.1, TCP/IP, WebSphere, and JAVA tuning

## Hardware (SUT) and Virtual Machine Configuration

IBM POWER7™ Server:

IBM POWER7™ Power 750 Express system was used for this study to showcase the capabilities of IBM's PowerVM virtualization technology. The system was configured with 4 sockets, 3.5 GHz, 32 cores supporting up to 4 threads (SMT4) per core as well as with 512GB of RAM. The system was used in two different configurations, 1) all 4 sockets and its associated memory (512GB) were enabled and 2) only one socket and its associated memory (128GB) was enabled for 8 core comparison study.

HP ProLiant DL 370 G6 Intel Xeon 5570 Server:

HP ProLiant DL370 G6, a large form factor rack high performance Intel Xeon 5570 based server, system was used for this study to compare the capabilities of latest VMware virtualization technologies. The system was configured with 2 sockets, 8 cores supporting up to 2 threads per core (HT mode). The system was also enabled for Turbo Mode, Intel VTx with EPT HW Virtualization assist.

## Systems used

System Configuration	Internal and external Storage	Network
IBM POWER7™ Model 750 3.5 Ghz, 8 cores per socket POWER7™ Processors, 128GB RAM per socket	8 Internal SAS 146.8GB disks	4 Host Ethernet Gbit ports
HP ProLiant DL 370 G6, 2.9 Ghz, 8 cores (2 sockets) Intel Xeon 5570 Processors, 96GB system RAM	Embedded SAS array controller with five 400GB disk and 1 72GB SAS internal disk	2 x HP NC375i quad-port Gigabit Network adapter
IBM Power 550 4.2 Ghz, 8 cores POWER6 Processors (Application Server)	1 Internal SAS 146GB disk	2 x FC 5767 dual port Gig-E
IBM Power 570 4.7 Ghz, 8 cores POWER6 Processors (DB2 Server)	2 x FC 5759 dual port FC adapter  IBM DS5300 with 2 EXP5000 drawers  1 Logical Volume 16 drives for logging  1 Logical Volume 16 drives for Database	2 x FC5767 dual port Gig –E
4 IBM x366 Intel Xeon 3.2Ghz 2 core, 2GB RAM	Internal disks	Gigabit Ethernet card in each system

## Software used

Category	PowerVM	VMware
Hypervisor	Power Hypervisor (Model 750 in-built hypervisor)	VMware ESX 4 Update 1 Version 4.0 Update 1 Build Number 208167 Release Date 2009/11/19

Guest OS	SuSE 11 (Power Linux)  AIX 6.1H TL 04	SuSE 11 GM x86_64  RHEL 5.4 x86_64
Middleware	IBM WebSphere® V7.0.07,  IBM DB2® v9.7	IBM WebSphere® V7.0.07,  IBM DB2® v9.7

### VMware Virtual Machine Configuration details:

1. VMware Virtual Machine was created using Virtual machine version 7 which was compatible with ESX 4.0 and greater hosts. It also provided greater virtual machine functionality.
2. Virtual Disk LSI Logic Parallel was used. It was noted that the LSI Logic Parallel adapter and the LSI Logic SAS adapter offer equivalent performance. (vSphere help)
  - a. VMware Best Practices recommend pvscsi interface for disk I/O intensive workloads since the two benchmarks covered in the study were not disk I/O intensive, pvscsi was not opted.
3. VMXNET3 was used for private network for DayTrader2.0 benchmark run.
4. Updated to the latest VMware Tools
5. Memory affinity was enabled
6. Collected esxtop –ab and vmstat from the virtual machine.

### System Tuning:

HP DL370 System:

- System (BIOS) default settings were used (HT enabled, Turbo mode enabled)

IBM Power 750 Express System:

- Intelligent Energy Optimization Mode was enabled

### AIX (PowerVM Guest OS) and Linux (VMware Guest OS) Tunings

- Large pages used for WAS instance

-Xss128k -Xgcpolicy:gencon -Xnloa -Xmo512m -Xmn2560m -Xgcthreads4 -  
Xcodecache16m -Xgc:tenureAge=2,noAdaptiveTenure -Xdisableexplicitgc -  
Djava.net.preferIPv4Stack=true -Dsun.net.inetaddr.ttl=0 -Xtrace:none -Xlp

Bound WAS instance to a core in case of multiple WAS instances

### AIX Tunings\*:

raso -r -o mtrc_enabled=0	(disables lightweight memory trace)
ctctrl -P memtraceoff	(disables component trace)
errctrl -P errcheckoff	(disables run-time error checking)
skeyctl -k off	(disables storage keys)

\*Disclaimer: These AIX tunings were done for benchmarking reasons customers do not change any of these tunables in a productive environment unless they are asked by AIX support team in the process of resolving an issue.

### DayTrader2.0 Test Methodology:

- Single VM scaling 1, 2, 4 and 8 cores for application server
- Same Database system used for both platform tests
- Same JIBE clients used for both platform tests
- Similar WebSphere tunings for both platform tests
- Tests were repeated with one WebSphere instance per 1 core and one WebSphere instance per 2 cores on both platforms. Best results were selected for each platform for comparison.
- Each test was run to drive CPU utilization to maximum.

## References

- 1) Comparing Costs of IBM PowerVM and x86 VMware for Enterprise Server Virtualization  
<http://ftp.software.ibm.com/common/ssi/sa/wh/n/pol03023usen/POL03023USEN.PDF>
- 2) IBM PowerVM  
<http://www.ibm.com/systems/power/software/virtualization/index.html>
- 3) IBM PowerVM Virtualization  
<http://ftp.software.ibm.com/common/ssi/pm/sp/n/pod03015usen/POD03015USEN.PDF>
- 4) What's New in VMware vSphere™ 4  
[http://www.vmware.com/files/pdf/vsphere\\_performance\\_wp.pdf](http://www.vmware.com/files/pdf/vsphere_performance_wp.pdf)
- 5) AIM7 Benchmarks  
<http://aimbench.sourceforge.net/>
- 6) DayTrader 2.0 Benchmarks  
<http://cwiki.apache.org/GMOxDOC20/daytrader.html>
- 7) VMware vSphere 4.0 update release notes  
[http://www.vmware.com/support/vsphere4/doc/vsp\\_esx40\\_u1\\_rel\\_notes.html](http://www.vmware.com/support/vsphere4/doc/vsp_esx40_u1_rel_notes.html)
- 8) VMware vSphere 4.0 Performance Best Practices Guide  
[http://www.vmware.com/pdf/Perf\\_Best\\_Practices\\_vSphere4.0.pdf](http://www.vmware.com/pdf/Perf_Best_Practices_vSphere4.0.pdf)
- 9) HP White Paper  
[http://www.vmware.com/files/pdf/Alliances-WebSphere-vSphere4\\_Nehalem-Performance-HPProliant-Servers.pdf](http://www.vmware.com/files/pdf/Alliances-WebSphere-vSphere4_Nehalem-Performance-HPProliant-Servers.pdf)
- 10) IBM Power Systems Advantages  
<http://www.ibm.com/systems/power/advantages/>
- 11) Optimizing Virtual Infrastructure with PowerVM and the IBM Systems Director VMControl Performance Summary  
<http://ftp.software.ibm.com/common/ssi/sa/wh/n/pow03033usen/POW03033USEN.PDF>
- 12) Red Hat Enterprise Linux  
<http://www.redhat.com/rhel>
- 13) SUSE Linux Enterprise  
<http://www.novell.com/linux>

14) VMware HyperThreading

[http://www.vmware.com/pdf/esx21\\_hypertreading.pdf](http://www.vmware.com/pdf/esx21_hypertreading.pdf)

15) VMWare Scheduler

[http://www.vmware.com/files/pdf/perf-vsphere-cpu\\_scheduler.pdf](http://www.vmware.com/files/pdf/perf-vsphere-cpu_scheduler.pdf)

16) VMware Hyper threading

<http://communities.vmware.com/docs/DOC-5101.jsessionid=6AA9D4DDD708B00ED77577039080B36B>

17) Interpreting esxtop statistics

<http://communities.vmware.com/docs/DOC-9279>

We appreciate the feedback received from Charles Gitamore, Susan Schreitmueller, Charles Farrell, Andrew Theurer, Dirk Michel, Herman Dirk, and Bruce Mealey which helped to improve the quality of the paper.



© IBM Corporation 2010  
IBM Corporation  
Systems and Technology Group  
Route 100  
Somers, New York 10589

Produced in the United States of America  
March 2010  
All Rights Reserved

This document was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this document in other countries.

The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

IBM, the IBM logo, ibm.com, POWER, POWER5, POWER6, Power Systems, and PowerVM are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml)

Other company, product, and service names may be trademarks or service marks of others.

IBM hardware products are manufactured from new parts, or new and used parts. In some cases, the hardware product may not be new and may have been previously installed. Regardless, our warranty terms apply.

Copying or downloading the images contained in this document is expressly prohibited without the written consent of IBM.

This equipment is subject to FCC rules. It will comply with the appropriate FCC rules before final delivery to the buyer.

Information concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of the non-IBM products should be addressed with those suppliers.

All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, to evaluate the performance of a system they are considering buying.

When referring to storage capacity, 1 TB equals total GB divided by 1000; accessible capacity may be less.

The IBM home page on the Internet can be found at: <http://www.ibm.com>.

The IBM Power Systems home page on the Internet can be found at: <http://www.ibm.com/systems/power/>

The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.

UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Linux is a trademark of Linus Torvalds in the United States, other countries or both.

Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of the Microsoft Corporation.

Intel, Itanium and Xeon are registered trademarks and MMX and Pentium are trademarks of Intel Corporation in the United States and/or other countries.

AMD Opteron is a trademark of Advanced Micro Devices, Inc.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. In the United States and/or other countries.

TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).

SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECcapc, SPECchpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corporation (SPEC).

SAP, mySAP and other SAP product and service names mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world.

SPC Benchmark-1 and SPC Benchmark-2 are trademarks of the Storage Performance Council.

[POW03041USEN-01](#)