HEP Specific Benchmarks of Virtual Machines on multi-core CPU Architectures

M. Alef
Forschungszentrum Karlsruhe, Karlsruhe, Germany
I. Gable
Department of Physics and Astronomy, University of Victoria, Victoria, Canada

Abstract

Virtualization technologies such as Xen can be used in order to satisfy the disparate and often incompatible system requirements of different user groups in shared-use computing facilities. This capability is particularly important for HEP applications, which often have restrictive requirements. The use of virtualization adds flexibility, however, it is essential that the virtualization technology place little overhead on the HEP application. We present an evaluation of the practicality of running HEP applications in multiple Virtual Machines (VMs) on a single multi-core Linux system. We use the benchmark suite used by the HEPiX CPU Benchmarking Working Group to give a quantitative evaluation relevant to the HEP community. Benchmarks are packaged inside VMs and then the VMs are booted onto a single multi-core system. Benchmarks are then simultaneously executed on each VM to simulate highly loaded VMs running HEP applications. These techniques are applied to a variety of multi-core CPU architectures and VM configurations.

Benchmark Technique

All measurements employed the HEP-SPEC06 bechmark because it has been shown to be linearly related to actual HEP application performance. HEP-SPEC06 is derived from SPEC CPU2006 from the Standard Performance Evaluation Corporation. HEP-SPEC06 came about from the 2007-2008 efforts of the HEPiX CPU Benchmarking working group to identify a suitable replacement for the now retired (February 2007) SPEC int 2000 benchmark that has been popular in the HEP community. We selected Scientific Linux (SL) as the operating system for all benchmarks because of its wide use in the HEP community. We selected the Xen Virtual Machine Monitor because of its ease of use with SL 5.2. The complete list of machines benchmarked is available in Table 1.

We are most interested in investigating the practicality of running n VMs where n is the number of cores per machine. In order to evaluate this efficiency of this method we compare benchmarks in three configurations: (1) 1 VM per core, (2) 1 VM using all the physical cores (i.e 1 VM per box), (3) the physical machine. Please refer to Figure 'Benchmarking Configuration'.

1 VM per core - n VMs are booted where n is the number of cores. Each VM is given a memory allocation as listed in Table 3. For example, an 8 core Intel Harpertown box with 16 GB of memory will have 8 VMs booted with 1900 MB of memory each, leaving 1184 MB of memory for the domain0 machine. Each virtual disk contains the SPEC CPU 2006 code. The benchmarks are then pre-compiled with HEP-SPEC appropriate flags on each VM. The HEP-SPEC06 benchmark is then executed simultaneously on all 8 VMs. This causes all VMs to compete for the resources of the physical CPUs in much the same way that the individual threads of the HEP-SPEC benchmark compete for resources on a multi-core box.

1 VM per box - In this case a virtual machine is created with VCPUs equal to the number of physical cores on the box. The benchmark is then executed in the same way as typical HEP-SPEC06 run with VCPUs standing in for physical cores.

Physical Machine - The machine was benchmarked in exactly the same way as any HEP-SPEC06 becnhmark run would be done; one benchmark process per physical core.

i386 VM - Physical Comparison

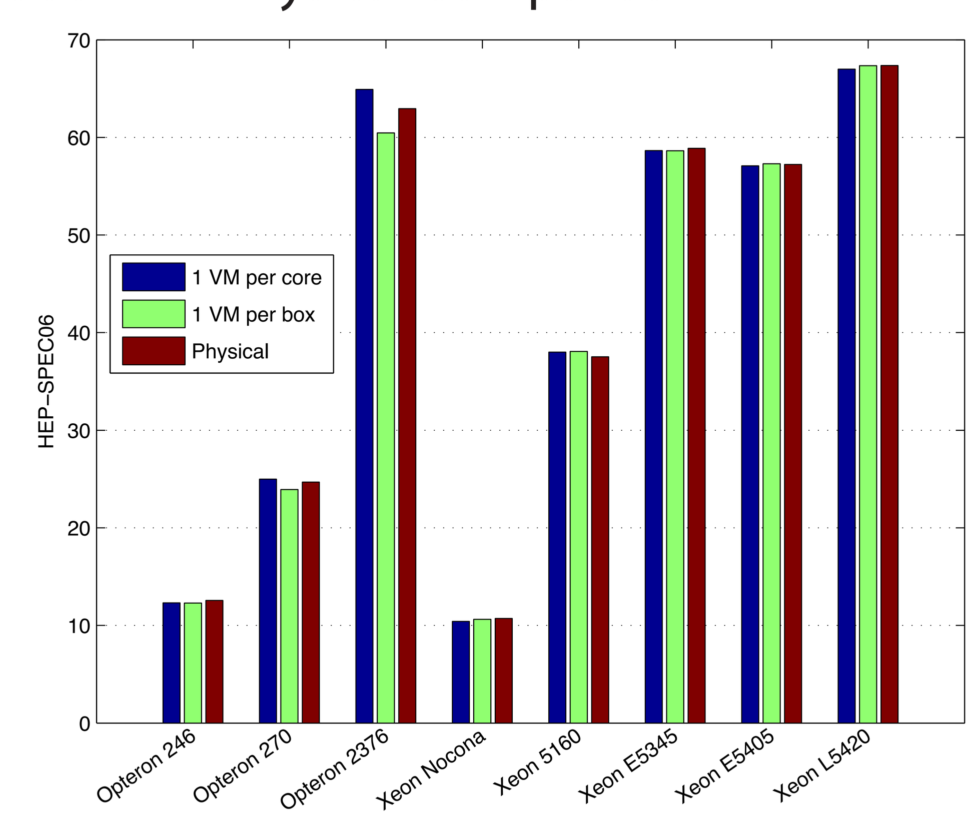

i386 VM % Difference from Physical

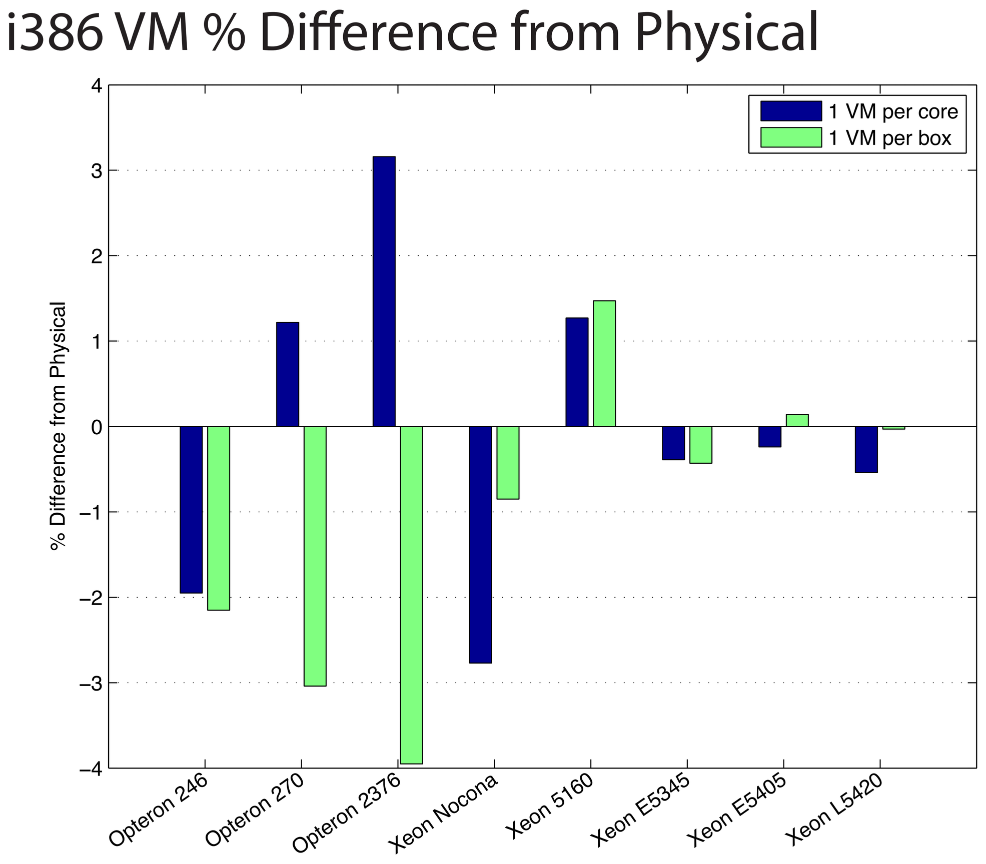

CPU Model	Mem (GB)	n Cores	Mainboard	Year
AMD Opteron				
246 (2.0 GHz SC)	4	2	MSI-9145	2003
270 (2.0 GHz DC)	8	4	MSI-9145	2005
2376 (2.3 GHz, QC, Shanghai)	16	8	Supermicro H8DMU+	2009
Intel Xeon				
Intel Xeon 3.00 GHz, SC (Nocona)	2	2	HP ProLiant DL360 G4, HT off	2004
5160 (3.0 GHz, DC, Woodcrest)	6	4	Intel Server Board S5000VCL	2006
E5345 (2.33 GHz, QC, Clovertown)	16	8	Supermicro CSE-812L-520CB	2007
E5405 (2.0 GHz, QC, Harpertown)	16	8	Dell PowerEdge 2950	2007
L5420 (2.5 GHz, QC, Harpertown)	16	8	Supermicro X7DCT	2008

Table 1. The benchmarking testbed machines. Note the SC, DC, and QC correspond to Single, Dual, and Quad Core.

CPU Type	n Cores	VM Type	Physical	1 VM per core	1 VM per box	% diff per core	% diff per box
Opteron 246	2	i386	12.56	12.32	12.29	-1.95	-2.15
Opteron 270	4	i386	24.68	24.98	23.93	1.22	-3.04
Opteron 2376	8	i386	62.94	64.92	60.45	3.16	-3.95
		x86_64	61.05	64.24	62.46	5.23	2.31
Xeon Nocona	2	i386	10.71	10.41	10.62	-2.77	-0.85
Xeon 5160	4	i386	37.52	38.00	38.07	1.27	1.47
Xeon E5345	8	i386	58.88	58.65	58.63	-0.39	-0.43
Xeon E5405	8	i386	57.22	57.08	57.30	-0.24	0.14
		x86_64	57.04	56.68	56.58	-0.63	-0.81
Xeon L5420	8	i386	67.36	66.99	67.34	-0.54	-0.03
		x86_64	67.31	68.22	67.23	1.35	-0.12

Table 2. The benchmark scores for all machines in the testbed. The scores are in HEP-SPEC06. Note that the number of cores is per machine and not per CPU.

VM Type	Hypervisor Version	Kernel Version	Disk Access
i386	Xen 3.0.3	2.6.18-92.1.13.el5xen i686	tap:aio
x86_64		2.6.18-92.1.13.el5xen x86_64	

Hardware Type	VM Memory Allocation
All AMD and Quad Core Intels	$n \times 1900$ MB
Intel Nocona	$n \times 870$ MB
Intel Woodcrest	$n \times 870$ MB

Table 3. Virtual Machine configuration; n is the number of physical cores.

x86_64 VM Comparison

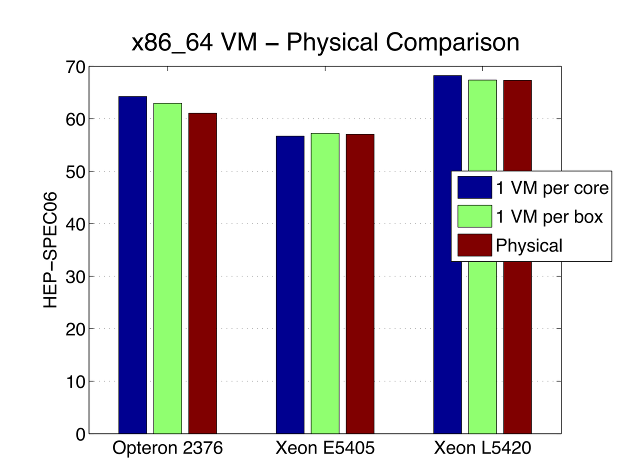

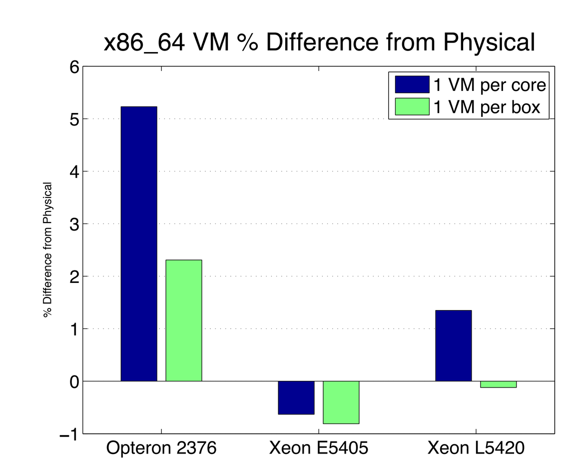

Benchmarking Configuration

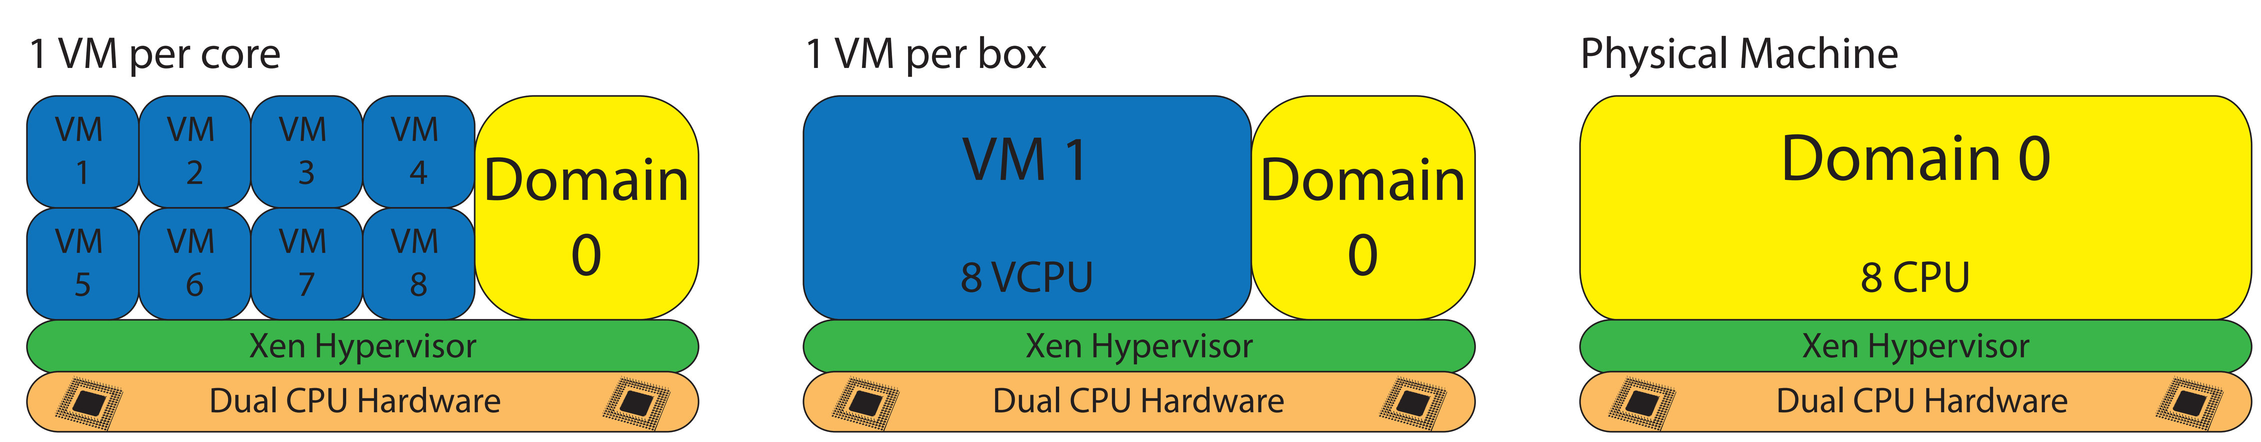

Conclusion

Our results have shown that, in terms of CPU performance, it is indeed very feasible to run highly loaded virtual machines on current generation multi-core CPUs. We established this using the HEP-SPEC06 benchmark which has been proven to map to real HEP application performance. No VM benchmark suffered more then a 4% decrease in performance relative to the physical box.

The 2008 generation Opteron 2376 (Shanghai) showed striking performance gains of 3.16% (i386 VM) and 5.23% (x86_64 VM) when running 1 VM per core. However, the 1 VM per box case suffered the largest VM performance decrease (-3.95%) of any CPU benchmarked. The dual core Opteron 270 behaved similarly with smaller scale differences.

All quad core Intel CPUs benchmarked exhibit negligible performance differences between the physical box in all VM configurations. For example the Xeon L5420 (Harpertown) showed performance loss of -0.51% and -0.03% with i386 VMs, in the 1 VM per core and 1 VM per box cases respectively. The largest difference among Intel CPUs is the 2003 generation Xeon 3.0 GHz(Nocona) where a -2.77% decrease is seen in the 1 VM per core case.

There now appears to be no CPU performance barriers for the wide spread adoption of virtualization technologies such as Xen in the HEP community. This could enable HEP applications to take advantage of the flexibility offered by completely virtualized software stacks for running on computing resources which were not compatible with HEP software requirements.

Note: Full references and further details available from the corresponding paper.

Karlsruhe Institute of Technology
GridKa
University of Victoria