

Performance analysis of selected hypervisors (Virtual Machine Monitors - VMMs)

Waldemar Graniszewski, Adam Arciszewski

Abstract—Virtualization of operating systems and network infrastructure plays an important role in current IT projects. With the number of services running on different hardware resources it is easy to provide availability, security and efficiency using virtualizers. All virtualization vendors claim that their hypervisor (virtual machine monitor - VMM) is better than their competitors. In this paper we evaluate performance of different solutions: proprietary software products (Hyper-V, ESXi, OVM, VirtualBox), and open source (Xen). We are using standard benchmark tools to compare efficiency of main hardware components, i.e. CPU (nbench), NIC (netperf), storage (Filebench), memory (ramspeed). Results of each tests are presented.

Keywords—virtualisation, virtualmachines, benchmark, performance, hypervisor, virtual machine monitor, vmm

I. INTRODUCTION

IN recent years the most popular IT projects have been based on cloud computing. With hardware resources, especially RAM, CPU power, storage (HDD), network interface cards (NICs) and other components, becoming cheaper, efficient access to those resources is crucial, and is conducted by specialized software - hypervisors. There are many of different hypervisors, e.g. VMWare, Xen, Hyper-V, Oracle VM, etc., which can be installed on almost all platforms. Some of them are better suited and tuned for available hardware. All providers claim that their solutions are the best, but results often depend on the used applications, i.e.: web servers, file servers, database applications, etc.

Server virtualization was a natural use of mainframes, where access to computing resources was done via terminals. There were also some formal requirements for computer architecture proposed [1]. With development of PCs, virtualization, as seen in 1980s, was not continued for x86 computers because of weak hardware and OS system resources. In the past 10-15 years, due to high performance and availability requirements for cloud computing, the role of virtualization solutions increased.

In this paper we would like to compare selected types of server virtualization and test their performance. These tests include CPU, RAM, HDD, and NIC performance and are conducted using standard benchmarking programs.

The remainder of this paper is structured as follows. Section II provides a brief description of all platforms described in the paper and a short review of related work. In Section III we present a description of the test environment and the methodology used to achieve performance comparison. Test

Authors are with the Faculty of Electrical Engineering, Warsaw University of Technology, Warsaw, Poland e-mail: waldemar.graniszewski@ee.pw.edu.pl, arciszea@ee.pw.edu.pl.

results for CPU, NIC, kernel compilation time and storage benchmarks' tests are presented in Section IV. Finally, in Section V, we draw some conclusions.

II. BACKGROUND AND RELATED WORK

In this section we present some general background for virtualisation technology (in Subsection II-A) and a short review of related work (in Subsection II-B).

A. Background

As mentioned earlier, in Section I, cloud computing and services provided by data centers require robust software for their operation. With data center server consolidation, the portability of each solution plays an important role. In the last decade both proprietary software like VMware ESXi, Microsoft Hyper-V, open source platforms, like Xen [2] and dual license programmes like VirtualBox [3] have been developed.

Looking for the virtualisation technology market share we consulted recent reports of one of the leading information technology research and advisory companies - Gartner, Inc. According Gartner's analysts about 80% of x86 server workloads are virtualised. Therefore firms compete in offering more lightweight software, supporting more workloads and agile development virtualisation solutions [4]. As of August 2016, Gartner's Magic Quadrant for x86 Server Virtualization Infrastructure specifies 8 companies (in alphabetical order): Citrix, Huawei, Microsoft, Oracle, Red Hat, Sangfor, Virtuozzo, and VMware. According to this report VMWare is seen as the market leader, followed by Microsoft. Earlier, in the years 2010 and 2011, also Citrix (the company that currently owns Xen) was also placed in this quadrant [5]. According to Hess, Gartner's report for the year 2015 is controversial [6]. The Author suggests that solution of the Virtuozzo company should be placed in the *Visionaries* instead of the *Niche Players*. He also suggests positions of *Challengers* for Citrix Xen Server and Red Hat's Enterprise Virtualization (RHEV) which also include containers. Another report prepared by Info-Tech Research Group placed Citrix Xen in the group of *Innovator* and Red Hat and Oracle as *Emerging Player*. Both VMware and Microsoft solutions remain in this report as *Champion* [7].

Taking into account market share of the virtualisation technology, in this paper, we evaluate performance of different solutions: proprietary software products (ESXi, Hyper-V, OVM, VirtualBox), and open source (Xen).

One of the most popular [8] classifications of hypervisors, also known as virtual machine monitors (VMM) is:

- Type 1: native or bare metal hypervisors,

- Type 2: hosted hypervisors.

In our selection of hypervisors for tests purposes only, VirtualBox is a Type 2 VMM. The remaining hypervisors are Type 1, i.e. native or bare metal. However, Hwang et al. claims that Xen, possess characteristics of both types [9].

Another often used [10] classification is:

- para-virtualization (PV),
- full virtualization (FV),
- hardware-assisted virtualization (HVM).

Owing to strong competition between all of the market players current hypervisors implement and can use at least 2 of mentioned techniques, i.e. para-virtualization (PV), full virtualization (FV) or hardware-assisted virtualization (HVM). Also in the case of VirtualBox, software detects if processor supports hardware virtualisation or not and switch to appropriate mode [3].

Due to high computation cost for virtual machines, so called overhead for guest operating system, in the last decade, there were also introduced container based operating systems. One of the proprietary solutions is Oracle Zone (former Solaris Zone) [11]. Another open source, container-based OS virtualization platform example is Linux-VServer [12].

B. Related work

In recent years, many benchmark comparisons has been published. VMware has published a performance comparison of their VMware ESX server 3.0.1 and Xen 3.0.3-0 [13]. The company used SPEC, Passmark and Netperf tools for its benchmark.

Danti compared KVM versus VirtualBox 4.0 [14]. He tested Windows 2008 R2 install time, Windows 2008 R2 installer load, Debian base system install time, host resource utilization and guests system performances.

Li et al. [15] tested three hypervisors: a commercially available one (the exact name of which was not explicitly provided in their paper) and open source software, i.e. KVM, Xen. They ran several MapReduce benchmarks, such as Word Count, TestDSFIO, and TeraSort and further validated their observed hypothesis using microbenchmarks.

Elsayed and Abdelbaki [16] tested and quantitatively and qualitatively compared the performance of the VMware ESXi5, Citrix Xen Server 6.0.2 and Hyper-V2008R2 Hypervisors. They used a customized SQL instance as workload simulating real life situations.

Varette et al. [17] evaluated performance and energy-efficiency of Xen 4.0, KVM 0.12 and VMware ESXi 5.1 in High Performance Computing (HPC) implementation. They used the Grid'5000 platform [18] to test hypervisors in a flexible and very close to real HPC environment.

Hwang et al. [9] not only compared four different virtualization platforms, but also tried to find and understand the strengths and weakness of each tested solution. They investigated Hyper-V implemented with Microsoft Windows Server 2008R2, KVM ver. 2.6.32-279, vSphere 5.0, and Xen ver. 4.1.2.

Morabito et al. [19] presented in their work a detailed performance comparison of traditional hypervisor based virtualization and lightway solutions (containers). They compared Kernel-based Virtual Machine (KVM), Linux Xen Containers (LXC), cloud solution (OSv) and Linux container (Docker).

Due to our research interest, infrastructure needs, and according to the market share of the virtualisation technology (see Subsection II-A), we evaluate performance of ESXi, Hyper-V, OVM, VirtualBox and Xen in this paper.

III. METHODOLOGY

As mentioned in Subsection II-B, there is a number of test approaches and tools used for OS and hardware performance analysis, e.g. [9], [11], [13]–[17], [19]. The methodology used in this study is similar to that used by researchers from George Washington University [9], but in contrary to their and other benchmarks, we tested newer versions of both native (bare metal hypervisors) and hosted hypervisors.

It has to be mentioned that the hypervisor performance tests were conducted on a single physical machine with an Intel Core 2 Duo E8400 dual-core (clock speed of 3GHz). The test computer was also equipped with 4GB of DDR2 RAM, as well as a 5400rpm 60 gigabyte hard drive. Each hypervisor was tested in an isolated manner, with a complete hard drive format between installations. An Ubuntu 12.04 LTS virtual machine was created on each hypervisor.

It should be noted that Oracle VirtualBox, as the only hosted hypervisor, was run on Windows Server 2012. The remaining part of the testing methodology remains the same, using the same guest system. OVM is based on the Xen hypervisor [20] and thus, as with Citrix XenServer, is run natively on the hardware [21] and, together with ESXi and Hyper-V, they are bare metal solutions.

Performance tests were undertaken on the same operating system - Ubuntu 12.04 LTS. Each major component of the virtual machine was tested separately. These components were the CPU, memory, hard drive, and network interface.

When a virtual machine is created, a certain number of virtual CPUs (vCPU) is allocated to it. The number of vCPUs limits how many physical cores the machine can use, but does not guarantee CPU time. Processing power is allocated to machines according to load and other parameters used by the hypervisor, eg. VM priority. This study has tested the performance of the VM with one and two vCPUs.

The hypervisor serves as a layer between the virtual machine's operating system and the host's physical memory, in order to provide data integrity and isolation of VMs. Thanks to hardware-assisted virtualization, this is accomplished via EPTs (Extended Page Tables, for Intel chipsets) or RVI (Rapid Virtualisation Indexing, for AMD). Both of these technologies allow for a large increase of speed compared to software memory virtualization [10]. All tested hypervisors make use of these technologies, albeit in varied ways, which results in different performance. Each test VM was equipped with 2GB of virtual RAM.

Hard drive IO is often the main factor causing slowdown of virtual machine operation, because the hypervisor must

emulate operation of a physical hard drive. To determine how efficient vHD implementations are across the tested hypervisors, IO speed tests of various files were conducted. Disk usage similar to observed in an email, file, and WWW server has also been simulated.

The network interface card (NIC) is one of the most important elements of a virtual machine. A vast majority of virtualization is used to make network services such as fileservers or WWW servers available. The two main factors in scoring a NIC are throughput and latency. Technology allowing for hardware-assisted virtualisation of NICs is being developed [22], but it is not widely used yet. This study conducted a throughput test on each machine.

After testing each component, a general test was run, which allowed for a more general view on each machine's performance. This test was Linux kernel compilation, which was done twice on each machine and the time required for it was averaged.

IV. RESULTS

In this section we provide results of the tests for each platform, i.e. Hyper-V, ESXi, OVM, VirtualBox, Xen. In Subsection IV-A we present test of the CPU with nbench. Next Subsection IV-B shows results of the NIC benchmark (netperf). In Subsection IV-C time of Linux kernel compilation is measured. Results of storage tests (filebench) are presented in Subsection IV-D. Finally, in Subsection IV-E, results of memory tests (ramspeed) are shown.

A. CPU Test (nbench)

CPU tests have been conducted with the use of nbench [23], a program which runs ten different benchmarks:

- numeric sort - sorting of an array of 32-bit integers,
- string sort - sorting of an array of strings of randomised length,
- bitfield - various bitwise operations,
- emulated floating-point - various floating-point operations,
- Fourier - calculating a Fourier transform,
- assignment - an algorithm for task allocation,
- IDEA - an encryption algorithm,
- Huffman - a lossless compression algorithm,
- neural net - a neural net simulation,
- LU decomposition - a method of matrix factorization

The results for CPU test are shown on the Fig. 1, respectively for 1vCPU (Fig. 1a) and 2vCPU (Fig. 1b).

As can be clearly seen in the results, the use of a Type 2 hypervisor (VirtualBox) means a great loss of CPU performance because of host OS overhead. Other tested VMMS are similar in performance to each other, with the results not going under 90% of the best result in each test.

B. NIC Test (Netperf)

Netperf is a NIC throughput benchmark, utilising a TCP_STREAM test [24]. The measured result is maximum




Fig. 2. Results of NIC throughput test.




Fig. 3. Results of kernel compilation time.

interface throughput, in megabits per second. The results of this test are shown on the Fig. 2

As evident, ESXi has achieved the highest result, having a throughput of approximately 24Gbit per second. It must be noted, however, that the results of Hyper-V, OVM and XenServer, approx. 16Gbit/s, are also enough for a vast majority of cases.

C. Kernel compilation

Linux kernel compilation was the next benchmark. It was run twice, as a test allowing for an overview of the machine's performance. The results are shown on the Fig. 3.

As evident from the above graph, ESXi has achieved the best result again, finishing compilation after 82 minutes on average. The difference between Type 1 hypervisors and VirtualBox is very clear here as well. The compilation time on that system was much longer, because it has to access resources via the computer's operating system.

D. HDD Test (Filebench)

Filebench is a disk load simulator, allowing for simulation of activity similar to activity observed while using the system as an email server, fileserver or webserver [25]. The main benchmark in this case was operations per second. Results for HDD tests with 1vCPU and 2vCPU are shown on the Fig. 4b and Fig. 4b respectively.




Fig. 1. Results of CPU test with nbench for 1vCPU (1a) and 2vCPU (1b)

ESXi and OVM Server seem to have an advantage over the other systems in most cases. Hyper-V has achieved a higher score as a fileserver than otherwise - this simulation uses large file I/O. VirtualBox has, again, achieved surprisingly high scores in fileserver tests, despite scoring very low in the other categories.

E. Memory test (ramspeed)

Ramspeed (ramsmg, for multi-cpu systems) is a tool which allows for memory throughput benchmarking [26]. It conducts four test for integers, and four tests for floating-point numbers:

- copy - copying of data from position to position
- scale - multiplication while copying
- add - addition of data from two positions and insertion in a third
- triad - a combination of add and scale

Results for the memory tests with 1vCPU and with 2vCPU are shown on the Fig. 3 and Fig. 4 respectively.


For multiprocessor VMs on Type 1 hypervisors, the results are similar, varying only by up to 10%. However, there is a noticeable discrepancy in the results for a single CPU VM in Hyper-V. This may suggest that this hypervisor is optimised mainly for multicore VMs. Hyper-V is the only hypervisor, which shows a marked improvement in memory throughput after increasing the vCPU count - by approximately 38%. The leader in both cases is ESX. What may be surprising is the

very good result achieved by VirtualBox in Integer Add and Integer Triad tests. This may suggest that in some cases it is capable of fully utilising the memory bus, which usually is not allowed by the operating system.


V. CONCLUSION

The clear leader in VM performance is ESXi. The surprisingly low scores of Hyper-V 2012R2 (as compared to 2008R2) may be caused by higher system requirements, and the low performance of the testbench system. Tests also clearly indicate that in a majority of cases Type 1 hypervisors have great advantages over Type 2 solutions, thanks to direct access to the system's resources. A large difference between these results and the results of J. Hwang, S. Zeng, F. y Wu, and T. Wood from George Washington University and IBM is noticeable. For example, XenServer has much higher scores, in many cases reaching ESXi results. Additionally, even though Hyper-V 2008 results are similar to the previous ones, 2012 shows a surprising loss in performance.

As evident from the conducted benchmarks and analysis, the main candidates for use in enterprise environments are the solutions created by Microsoft and VMWare. Most of the remaining systems do not have a majority of features necessary in large scale virtualization. It must be, however, noted that XenServer has achieved good scores, and is a very affordable solution, which makes it a good choice for smaller scale operations, eg. a small office. Because of complex configuration



(a) 1vCPU



(b) 2vCPU

Fig. 5. Results of memory test with memtest for 1vCPU (5a) and 2vCPU (5b)

and its relatively small feature set, Oracle's product may be recommended mainly for institutions that already use many solutions created by the company, eg. Oracle Database. The drawbacks have less impact then, thanks to the integration features OVM offers, as well as good results in benchmarks. Hyper-V and VMWare offer similar, wide capabilities. In previous years, the license cost for VMWare's software was much higher, however recently, with the release of Windows Server 2012R2, Microsoft has raised the price of their solution, causing the price points to become similar. Considering the total unit price for Windows Server 2012 R2 Datacenter and System Center 2012 R2 Datacenter, approx. 7000 USD (after a 19.2% discount due to purchasing over 25 licenses) [27], and the license cost for vSphere, approx. 5000 USD, VMWare's product appears to be the better choice in price/quality. However, Microsoft also offers gratis operating system licenses, which allows for saving 800 US Dollars per virtual machine, which is not an option for VMWare, assuming, as is common in enterprise, the use of Microsoft solutions. Hyper-V also allows for easier implementation, thanks to quick integration in a pre-existing Windows infrastructure. However, ESXi has shown better results in benchmarks, trumping Hyper-V in

nearly every category. VirtualBox has scored surprisingly well in benchmarks in some categories. However, this doesn't remove the large loss of performance observed in other tests. Ultimately, the choice must be made with consideration for benchmark results, but also features and cost. Each company must select a solution based on its own needs and budget. Due to the benchmark results in some cases showing large discrepancies from the expected results, indubitably a second round of tests on a more powerful testbench would be useful. This would allow to eliminate the factor of low hardware performance, which may improve the results of eg. Hyper-V 2012R2, as well as allowing for installation of a newer version of ESXi.

ACKNOWLEDGMENT

We thank the anonymous reviewers for their careful reading of our manuscript and their many insightful comments and suggestions.

We would also like to thank Danuta Ojrzeńska-Wójter for her comments and suggestions which helped with manuscript preparation and Marcin Iwanowski and Krzysztof Szczypiorski for their valuable remarks.




Fig. 4. Results of hard disk test with filebench for 1vCPU (1a) and 2vCPU (1b)

REFERENCES

- [1] G. J. Popek and R. P. Goldberg, "Formal requirements for virtualizable third generation architectures," *Commun. ACM*, vol. 17, no. 7, pp. 412–421, Jul. 1974. [Online]. Available: <http://doi.acm.org/10.1145/361011.361073>
- [2] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield, "Xen and the art of virtualization," *SIGOPS Oper. Syst. Rev.*, vol. 37, no. 5, pp. 164–177, Oct. 2003. [Online]. Available: <http://doi.acm.org/10.1145/1165389.945462>
- [3] Oracle. Virtualbox. [Online]. Available: <https://www.virtualbox.org/>
- [4] T. J. Bittman, P. Dawson, and M. Warrilow. (2016) Magic quadrant for x86 server virtualization infrastructure. [Online]. Available: <https://www.gartner.com/doc/reprints?id=1-3E2WESI&ct=160804&st=sb>
- [5] M. Kędziora. (2014) Co to jest magiczny kwadrat gartnera (gartner magic quadrant)? [Online]. Available: <https://blogs.technet.microsoft.com/mkedziora/2014/07/16/co-to-jest-magiczny-kwadrat-gartnera-gartner-magic-quadrant/>
- [6] K. Hess. (2015) Gartner's magic quadrant for x86 server virtualization infrastructure is a head scratcher. [Online]. Available: <http://www.zdnet.com/article/gartners-magic-quadrant-for-x86-server-virtualization-infrastructure-is-a-head-scratcher/>
- [7] Info-Tech Research Group Inc. (2016) Vendor landscape: Server virtualization. [Online]. Available: <https://www.infotech.com/research/ss/it-vendor-landscape-server-virtualization>
- [8] Oracle. (2011) Oracle vm user's guide 3.0, e18549-03. [Online]. Available: https://docs.oracle.com/cd/E20065_01/doc.30/e18549.pdf
- [9] J. Hwang, S. Zeng, F. y. Wu, and T. Wood, "A component-based performance comparison of four hypervisors," in *2013 IFIP/IEEE International Symposium on Integrated Network Management (IM 2013)*, May 2013, pp. 269–276.
- [10] VMware. (2008) Understanding full virtualization, paravirtualization, and hardware assist. [Online]. Available: https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/VMware_paravirtualization.pdf
- [11] D. Price and A. Tucker, "Solaris zones: Operating system support for consolidating commercial workloads," in *Proceedings of the 18th USENIX Conference on System Administration*, ser. LISA '04. Berkeley, CA, USA: USENIX Association, 2004, pp. 241–254. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1052676.1052707>
- [12] S. Soltész, H. Pötzl, M. E. Fiuczynski, A. Bavier, and L. Peterson, "Container-based operating system virtualization: A scalable, high-performance alternative to hypervisors," in *Proceedings of the 2nd ACM SIGOPS/EuroSys European Conference on Computer Systems 2007*, ser. EuroSys '07. New York, NY, USA: ACM, 2007, pp. 275–287. [Online]. Available: <http://doi.acm.org/10.1145/1272996.1273025>
- [13] VMware. (2007) A performance comparison of hypervisors. [Online]. Available: https://www.vmware.com/pdf/hypervisor_performance.pdf
- [14] G. Danti. (2011) Kvm vs virtualbox 4.0 performance comparison on rhel 6. [Online]. Available: <http://www.ilsistemista.net/index.php/virtualization/12-kvm-vs-virtualbox-40-on-rhel-6.html?limitstart=0>
- [15] J. Li, Q. Wang, D. Jayasinghe, J. Park, T. Zhu, and C. Pu, "Performance overhead among three hypervisors: An experimental study using hadoop benchmarks," in *2013 IEEE International Congress on Big Data*, June 2013, pp. 9–16.
- [16] A. Elsayed and N. Abdelbaki, "Performance evaluation and comparison of the top market virtualization hypervisors," in *Computer Engineering Systems (ICCES), 2013 8th International Conference on*, Nov 2013, pp. 45–50.
- [17] S. Varrette, M. Guzek, V. Plugaru, X. Besson, and P. Bouvry, "Hpc performance and energy-efficiency of xen, kvm and vmware hypervisors," in *2013 25th International Symposium on Computer Architecture and High Performance Computing*, Oct 2013, pp. 89–96.
- [18] Grid'5000. [Online]. Available: <http://grid5000.fr>
- [19] R. Morabito, J. Kjällman, and M. Komu, "Hypervisors vs. lightweight virtualization: A performance comparison," in *Cloud Engineering (IC2E), 2015 IEEE International Conference on*, March 2015, pp. 386–393.
- [20] Oracle. (2013) Oracle vm server for x86 virtualization and management. [Online]. Available: <http://www.oracle.com/us/technologies/virtualization/oraclevm/026996.pdf>
- [21] Citrix. (2016) Xenserver tech info. [Online]. Available: <https://www.citrix.com/products/xenserver/tech-info.html>
- [22] S. Tripathi, N. Droux, T. Srinivasan, and K. Belgaied, "Crossbow: From hardware virtualized nics to virtualized networks," in *Proceedings of the 1st ACM Workshop on Virtualized Infrastructure Systems and Architectures*, ser. VISA '09. New York, NY, USA: ACM, 2009, pp. 53–62. [Online]. Available: <http://doi.acm.org/10.1145/1592648.1592658>
- [23] BYTE. nbench. [Online]. Available: <http://www.tux.org/~mayer/linux/bmark.html>
- [24] Hewlett-Packard. Netperf. [Online]. Available: <http://www.netperf.org/netperf/>
- [25] Filebench. [Online]. Available: <https://github.com/filebench/filebench/wiki>
- [26] R. M. Hollander and P. V. Bolotoff. Ramspeed. [Online]. Available: <http://alasisir.com/software/ramspeed/>
- [27] Microsoft. (2014) Datacenter tco tool. [Online]. Available: <http://www.datacentertool.com>