

# Moodle Portal in Virtualized Environment – a Performance Analysis

Vladimír Mašín  
Charles University  
Faculty of Medicine  
in Hradec Kralove,  
Simkova 870,  
500 03 Hradec Kralove,  
Czech Republic  
Email: masin@lfhk.cuni.cz

Martin Kopeček  
Charles University  
Faculty of Medicine  
in Hradec Kralove,  
Simkova 870,  
500 03 Hradec Kralove,  
Czech Republic  
Email: kopecema@lfhk.cuni.cz

Josef Hanuš  
Charles University  
Faculty of Medicine  
in Hradec Kralove,  
Simkova 870,  
500 03 Hradec Kralove,  
Czech Republic  
Email: hanus@lfhk.cuni.cz

**Abstract**—The Moodle portal of our faculty is running in a virtualized environment together with other about 50 application servers and 60 virtualized desktops. Increasing traffic on the site (reaching over 400 000 views/posts monthly) forced us to assess its performance impact on the virtualization environment. The performance analysis identified processor cycles and disk operations as the bottlenecks of the system. We are planning to address these issues with increasing of the number of processor cores in our virtualization hosts and with a solid state disk upgrade of the disk array used in our virtualization environment in our next hardware upgrade cycle.

## I. INTRODUCTION

THE FACULTY of Medicine in Hradec Kralove is like most medical faculties massively overburdened with the combination of both research and educational tasks. Just for example, our Department of Medical Biophysics has only 8 full-time staff members, who are currently teaching 16 pre-graduate courses and 2 postgraduate courses, while at the same time they are actively involved in research projects in several different fields, ranging from mathematical statistics and applications of mathematics and statistics in medicine [1]–[3] through mathematical modeling of apheresis procedures [4]–[5], applications of shape memory materials in general medicine [6] and dentistry [7]–[8] up to inclusion of modern teaching method based on applications of information technology in medical education [9]–[11]. The conditions in the other departments are hardly any better.

The only solution of such situation is extensive application of the methods of unsupervised learning [12], which are in our faculty represented mainly by the e-learning courses running in the learning management system (LMS) Moodle.

Unfortunately there is an unavoidable downside of this approach – the ever-increasing demand for the computational and data storage resources. We are therefore currently planning a significant upgrade of our IT infrastructure, which we would like to base on a thorough analysis of the performance requirements and the potential bottlenecks of our current setup.

This work was supported by the program PROGRES Q40-09.

## II. CURRENT STATE OF OUR INFRASTRUCTURE

### A. Virtualization environment

Our faculty is currently using a virtualization environment consisting of a cluster of four identical hosts (DELL PowerEdge R810 servers equipped with two Xeon E7540 processors, each containing six cores running at 2.0 GHz, and 256 GB of RAM), connected through redundant 8 Gbps FibreChannel connections to the IBM DS3512 disk array containing 48x 600 GB and 12x 450 GB 15k rpm SAS drives organized in RAID 10 arrays. The hosts are running the VMware ESXi 6.0 hypervisors and are managed through the VMware vSphere 6 Standard.

Apart from the Moodle portal servers, there are about 50 various application servers (mostly web servers, but also MS Exchange mail servers and MS SQL database servers) and about 60 virtualized desktops running in the cluster.

### B. Moodle portal

The Moodle portal of our faculty is consisting of two virtual machines (VMs), both running a 64-bit version of the Ubuntu 14.04.1 operating system. The front-end server has 4 virtual CPUs, 6 GB RAM, and 150 GB of storage space allocated; this storage space is divided into three separate volumes, located at different physical arrays, and mounted as “root”, “moodledata”, and “logs”. It is currently running Moodle 3.2.2+. The back-end database server is running on a VM with 4 virtual CPUs, 5 GB of RAM, and 95 GB of storage space; it has just two storage volumes – a combined “root” + “logs” volume and a standalone “data” volume. The database server VM is currently running MySQL 5.5.55.

There are 381 active courses within our Moodle portal to the present day (May 2017); most of them are based on text and image information with only very limited use of video materials so far.

Our portal is used by about 1200 students and 400 staff members almost on a daily basis. Its usage pattern is showing strong year-to-year growth with peaks during the examination period of the winter term and periods of limited activity during the summer holidays (see Fig. 1).




Fig 1. Moodle portal usage statistics

### III. PERFORMANCE METRICS USED

Our analysis of the impact of our Moodle portal on the virtualization environment is mostly based on the detailed performance metrics recorded by the tools built into the vSphere management system (CPU, RAM, and disk usage – see Table I and Table II), supplemented with few statistical figures provided by Moodle itself (monthly usage statistics for different user roles – see Fig. 1).

The metrics recorded in the vSphere were chosen according to our preliminary observations of behavior of the whole virtualization system and the analysis of probable performance bottlenecks affecting our Moodle portal. CPU and RAM resources are often considered to be the most critical ones in the virtualized environments; on the other hand disk performance, especially disk I/O operations may be sometimes incorrectly given low priority. [13] We also considered recording of the network performance metrics, but we found them to be unnecessary in this particular application, as the current content of our Moodle portal was not very demanding on network resources (there were just few standard definition videos, no HD neither UHD videos stored). The over-

all utilization of network resources in our virtualized environment was also quite low.

The final set of the metrics therefore included:

- the actual CPU usage and the demand for CPU resources, both measured in MHz and summarized for all virtual cores in the VMs
- the amount of memory granted to the VMs and its actual usage
- the read and write disk operations summarized for all volumes in the VMs
- the highest value of disk latency observed in each time interval, aggregated for all volumes in the VMs

All these metrics were recorded separately for the front-end and the back-end VMs in 30 minute intervals for total duration of one week to accumulate a representative sample of the performance variation. The default aggregation methods provided by the vSphere were used in all metrics: The CPU, memory, and IOPS metrics were calculated as averages for each of the 30 minute intervals; the latency figures represented the highest values observed in each interval.

TABLE I.  
RECORDED PERFORMANCE METRICS – CPU AND MEMORY

	CPU Usage (MHz)	CPU Demand (MHz)	Memory Active (MBytes)	Memory Granted (MBytes)
<b>Moodle front-end (average value)</b>	102.7	156.4	469	6144
<b>Moodle front-end (maximum value)</b>	756	1807	2864	6144
<b>MySQL back-end (average value)</b>	37.8	47.4	162	5120
<b>MySQL back-end (maximum value)</b>	264	421	811	5120

TABLE II.  
RECORDED PERFORMANCE METRICS – DISKS

	Read operations per second	Write operations per second	Disk latency (ms)
Moodle front-end (average value)	5.7	0.2	2.2
Moodle front-end (maximum value)	240	10	101
MySQL back-end (average value)	0.1	7.4	1.0
MySQL back-end (maximum value)	10	144	41

#### IV. PERFORMANCE ANALYSIS

##### A. CPU resources

Even though the demand for CPU resources was relatively modest in both front-end and back-end VMs, it was consistently outstripping the available resources of the hosts where these VMs were running by a large margin. Both VMs were therefore CPU-limited not only in the peak load conditions but also during regular operation. The uneven character of the demand for the CPU resources can be nicely demonstrated in the chart plotting these metrics recorded for the Moodle front-end over time (see Fig. 2):




Fig 2. Moodle front-end – CPU usage and demand

##### B. Memory resources

Unlike the CPUs, the memory resources were not strained in any of the VMs even at the peak load, when the more demanding front-end VM consumed just close to 50% of its allocated memory, and the back-end VM managed to be even more prudent, consuming just up to about 16% of its allocated memory. This result should be attributed to the choice of the OS, which was installed without any memory intensive graphical user interface (GUI), as well as to the memory conserving features of the hypervisor.

##### C. Disk resources

Disk usage patterns were significantly different in each of the VMs, which is hardly surprising. The front-end (essentially a web server) was heavily leaning towards the read operations, while the back-end (database) performed mostly write operations. The absolute numbers of disk operations in both VMs were relatively low; but as the latency figures showed, these values were affected by the overall performance limits of the disk array used in our virtualization environment anyway. The coincidence of high IOPS and high latency figures can be illustrated by the chart depicting both of these metrics in the Moodle front-end (see Fig. 3):




Fig 3. Moodle front-end – aggregated disk performance

#### V. CONCLUSION

Our main conclusion should probably be that we were able to demonstrate that even a relatively large Moodle portal (381 courses, about 1600 users, up to 400,000 page views per month) could be running in our virtualized environment on very modest resources. The main reason of such low demands was without any doubt the character of e-learning materials presented in our portal – when we start using the video-based materials on a larger scale, the demands of the portal are undoubtedly going to increase significantly.

We were also able to identify CPU and disk performance as the two main bottlenecks affecting responsiveness of our Moodle portal. When we aligned these metrics to the activity logs in Moodle, we were able to identify the most resource intensive operations: The CPU activity in both front-end and back-end as well as the disk write operations were highly taxed by grading of simultaneously running tests and re-grading operations, while the disk read operations reached their high values in backup sessions (see Fig. 2 – 5):




Fig 4. MySQL back-end – CPU usage and demand

Finally, the results of this performance analysis provide us (and not only us) with the invaluable clues for a proper design of the virtualized environment capable of running of such demanding tasks without annoying lags.




Fig 5. MySQL back-end – aggregated disk performance

Based on these findings we decided to increase number of processor cores in the planned new virtualization hosts, include solid-state disks, able to provide at least  $10^5$  IOPS, in the planned new disk array, and separate virtualized servers and VDI machines to different clusters, if possible.

#### ACKNOWLEDGMENT

We would like to thank to our IT support staff from the Computer Technology Center of our faculty for their invaluable help with collecting of the performance metrics used in this article.

#### REFERENCES

- [1] D. Kordek, "Statistical Analysis of Subconscious Human Behaviour," in *Aplimat 2009: 8th International Conference. Proceedings*, Bratislava, Slovakia, 2009, pp. 783-789.
- [2] D. Jezbera, D. Kordek, J. Kriz et al., "Walkers on the circle," *J. Stat. Mech.-Theory Exp.*, vol. 2010, no. 01, Jan. 2010, <http://dx.doi.org/10.1088/1742-5468/2010/01/L01001>
- [3] D. Kordek, "The definition of optical systems aberrations to secondary school students regarding their knowledge of mathematics," in *AIP Conference Proceedings*, vol. 1804, 2017, pp. 030004-1 - 030004-6.
- [4] M. Blaha, V. Masin, P. Stransky et al., "Optimization of the therapeutic procedure during LDL-apheresis – a computerized model," *Transfus. Apher. Sci.*, vol. 32, no. suppl. 2, pp. 149-156, Apr. 2005, <http://dx.doi.org/10.1016/j.transci.2004.10.022>
- [5] V. Masin, M. Blaha, P. Stransky et al., "Optimization Of Therapeutic Procedure During LDL Apheresis – Verification Of The Computerized Model In The Clinical Practice," *Transfus. Apher. Sci.*, vol. 36, no. 1, pp. 39-45, Feb. 2007, <http://dx.doi.org/10.1016/j.transci.2006.10.004>
- [6] J. Zahora, A. Bezrouk, J. Hanus, "Models of stents – Comparison and applications," *Physiol. Res.*, vol. 56, pp. 115-121, 2007.
- [7] A. Bezrouk, L. Balsky, I. Selke Krulichova et al., "Nickel-titanium closed-coil springs: evaluation of the clinical plateau," *Rev. Chim.*, vol. 68, no.5, pp. 1137-1142, May 2017.
- [8] A. Bezrouk, L. Balsky, M. Smutny et al., "Thermomechanical properties of nickel-titanium closed-coil springs and their implications for clinical practice," *Am. J. Orthod. Dentofac. Orthop.*, vol. 146, no. 3, pp. 319-327, Sep. 2014, <http://dx.doi.org/10.1016/j.ajodo.2014.05.025>
- [9] J. Hanus, T. Nosek, J. Zahora et al., "On-line integration of computer controlled diagnostic devices and medical information systems in undergraduate medical physics education for physicians," *Phys. Medica*, vol. 29, no. 1, pp. 83-90, Jan. 2013, <http://dx.doi.org/10.1016/j.ejmp.2011.12.002>
- [10] J. Zahora, J. Hanus, D. Jezbera et al., "Remotely Controlled Laboratory and Virtual Experiments in Teaching Medical Biophysics," in *6th International Conference of Education, Research and Innovation (iceri 2013). Proceedings*, Seville, Spain, 2013, pp. 900-906.
- [11] J. Hanus, J. Zahora, V. Masin et al., "On-Line Incorporation of Study and Medical Information System in Undergraduate Medical Education," in *6th International Conference of Education, Research and Innovation (iceri 2013). Proceedings*, Seville, Spain, 2013, pp. 1500-1507.
- [12] J. Feberova, T. Dostalova, M. Hladikova et al., "Evaluation of 5-year Experience with E-learning Techniques at Charles University in Prague. Impact on Quality of Teaching and Students' Achievements," *New Educ. Rev.*, vol. 21, no. 2, pp. 110-120, 2010.
- [13] J. Baker, (2015, Dec. 4), Top 10 VMware Metrics to help pinpoint bottlenecks [Online]. Available: <http://capacitymanagement-metron.blogspot.cz/2015/12/top-10-vmware-metrics-to-help-pinpoint.html> [Accessed: 2017, May 9].