Performance Evaluation of Cloud Computing

Rodrigo B. Vianna*, Luiz F. Bittencourt.

Abstract

This research's goal is to model a performance relationship between the resources of a private cloud and of a public cloud and determine their capabilities to meet application requirements.

Key words:

Cloud computing, Openstack, Performance.

Introduction

Cloud computing refers to the use of storage capacity and computing on shared servers, connected to the user through the Internet^[1]. This form of computing has been consolidated as a standard of offering computational capacity as a service.

Decisions about the management of these virtualized resources can be challenging^[2], being usual the challenge of allocating virtual machines.

This study's hypothesis is: private clouds when overloaded show a performance drop that can be avoided through the use of public clouds virtual machines. Therefore, better management of the local resources of a private cloud can bring benefits to running applications.

Results and Discussion

Making use of Openstack4j – a ^[3]Java OpenStack client SDK which allows provisioning and control of an OpenStack deployment – private cloud virtual machines (VMs) were benchmarked. This private cloud is located at the ^[4]Computer Networks Laboratory at UNICAMP, being composed of a single node ^[5]Openstack installation.

The chosen benchmarks for comparing the performance degradation were ^[6]Stream (for memory bandwidth) and ^[7]Sysbench (for CPU).

The used methodology consisted in, starting with 0 VMs, instantiating a new VM, and benchmarking all the existent VMs concurrently, until the limit of 9 VMs. This limit was due to the RAM limitation in the server - 16 GB with each new instance using 2 GB.

For the Stream benchmark, 10 tests were done each time, so an overlap was possible, assuring concurrence.

Figure 1. Figure 1. Median of the memory access rate for the Stream benchmark (higher is better).

Figure 2. Average running time for the Sysbench CPU benchmark (lower is better)

Conclusions

The results suggest that for memory intensive applications, having multiple VMs on the same physical machine is going to heavily affect their performance. On the other hand, the CPU-intensive application in Sysbench had a lower degradation, allowing multiple instances in the same physical machine, to a certain extent.

Acknowledgement

The authors would like to thank CNPq/PIBIC/UNICAMP for the financial support, and thank Marcelo C. S. Araújo for his help with OpenStack.

^[1]M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "A view of cloud computing," Commun. ACM, vol. 53, no. 4, pp. 50-58, April 2010.

^[2]L. F. Bittencourt, E. R. M. Madeira and N. L. S. Da Fonseca, "Scheduling in hybrid clouds," in IEEE Communications Magazine, vol. 50, no. 9, pp. 42-47, September 2012.

^[3]www.openstack4j.com/ [Accessed 04 July 2017]

^[4]<http://www.lrc.ic.unicamp.br/lrc/> [Accessed 04 July 2017]

^[5]<http://www.openstack.org/> [Accessed 04 July 2017]

^[6]<http://manpages.ubuntu.com/manpages/xenial/man1/sysbench.1.html> [Accessed 04 July 2017]

^[7]<https://www.cs.virginia.edu/stream/> [Accessed 04 July 2017]