Implementing a Cloud Computing Service in an Academic Institution

Emmanuell D. Carreño[1,2], Eduardo Roloff[1], Manuel G. Florez[2], and Philippe O. A. Navaux[1]

[1]Universidade Federal do Rio Grande do Sul
Instituto de Informática
Programa de Pós-Graduação em Computação
{edcarreno,eroloff,navaux}@inf.ufrgs.br

[2]Universidad Industrial de Santander
Facultad de Ingenierías Físico-Mecánicas
Escuela de Ingeniería de Sistemas e Informática
mgflorez@uis.edu.co

Abstract

In academic environments, the infrastructure present and the limitations on the amount of resources which they have, made it necessary to find ways for using the idle capacity of the hardware for the benefit of their members, it is there where cloud computing could allow efficient utilization of this idle capacity. Cloud computing represents a paradigm shift in computing, from a model focused on providing single services to a model providing multiple services on shared hardware, this model represents a big opportunity for datacenters to make a better usage of their resources. This research work presents the implementation of a user cloud computing infrastructure as a service model, in order to support projects by members and research groups from an academic environment.

1. Introduction

Cloud computing [9] [12] has become a very important subject in the last years [6]. It represents a change in the way datacenters are used and also a developing paradigm shift, not only about computing and the way it is used, but also in the way it is accessed [8].

The service models that rely on cloud computing created an industry moving billions of dollars per year [4]. Despite being a relatively new paradigm in the market, it presents so many features, such a high degree of usability and usage under the on-demand model, that it has become a promising technology in a few years [10].

Open Source cloud software present a great opportunity for organizations interested in experimenting with cloud computing to do it, without the burden of big starting investments in capital or in infrastructure, because they could use their current infrastructure and figure out their real needs before investing massively on this new paradigm. However, there is a lack of studies regarding to the implementation of real environments.

This paper proposes an implementation of an actual setting in an academic environment, seeking to assess the ability of commodity machines to meet the needs of an academic community, and determine some of the costs and the computing power of the implementation.

2. Related Work

Hosseini, Greenwood and Sommerville [7] studied a migration process of a commercial application to the Amazon EC2 Cloud. The migration was conducted by a company in the UK, which has done the migration of their IT system, from an in-house data center to Amazon EC2. They show that the system infrastructure in this case study would have cost 37% less over 5 years on EC2, and that using cloud computing could have potentially eliminated 21% of the support calls received by the studied system.

Steinbauer et al. [11] proposed the use of cloud computing for academic environments, to be used as support for the classroom. The authors proposed a configuration using the OpenNebula cloud toolkit. The applications considered in this study were Hadoop and Storm, that were useful in the academic context of the paper.

Waßman et al. [13] provide a comprehensive study about energy consumption of Virtual Machines. They proposed a system to help the user to decide where to deploy his application to achieve more cost efficiency. The analysis is conducted in terms of software behavior, and they evaluate the proposal using several machines running a set of virtual machines (VMs).

3. Motivation

The Systems Engineering and Computer Science School (EISI) at the Universidad Industrial de Santander did not had a service to perform application development and testing for its academic community. Furthermore, the computing infrastructure was inefficient, because it was underused and was difficult to be accessed by the community members. This work aimed to change these two factors, by implementing an infrastructure as a service cloud computing platform.

Initially, we performed an analysis of the needs of the undergraduate projects developed in the previous two years [2]. In the results, we noticed that the projects was mostly regarding software development, and a high percentage of them were in the area of web development using some common frameworks. The results allowed us to identify the type of cloud service that was needed by the EISI to help these projects; an IaaS model would make it easier for the EISI community to test those frameworks, develop directly on them and, if they reach a production stage with their projects, they could be kept in the cloud service in production usage.

4. Project Design

During the design of the project, the desired architecture was defined. It had some prerequisites that need to be followed.

- All tools used must be free.
- There should be no incompatibility of licenses or limitations on use of the different tools.
- They must operate on the Debian operating system.

The free of charge aspect of the tools was required by this project, because it is for an academic environment and it is intended to be replicated without incurring upfront costs, that could be caused by operating systems or applications licenses. Also, the applications required by the platform should not be covered in licenses restraining their usage by more than one user or limited to specific hardware or software.

The implemented prototype consists of three essential parts: a hypervisor, that manages the entire life cycle of the virtual instances; a management tool, for orchestrating provided services and resources and a main server or frontend, who controls the security and is the entry point of the services.

In this section, we describe the hardware available to the project. Also the hypervisor selection process is explained. After the hypervisor selection, the cloud manager was chosen.

4.1. Hardware & Operating System

The physical hardware used in this project consists of two nodes, they are Dell PowerEdge R720 servers, each one with a Intel Xeon E5640 processor, with eight cores and 24 GB of RAM. The servers are interconnected using a Gigabit Ethernet link, they are connected to the campus LAN through a Fast Ethernet connection. We used the Debian Squeeze operating system. One virtual machine in one of the nodes was used as Frontend.

4.2. Hypervisor Selection

Virtualization [5] is one of the key features of a cloud computing environment, therefore, the selection of a good hypervisor is critical to the success of a real implementation. In this section, we discuss our decision of which hypervisor will be used to control all the VMs in our environment. Our choice was to use VirtualBox as hypervisor.

We used VirtualBox because of their characteristics, its license is GPL v2, the amount of RAM, vCPUs and VMs that can used is almost unlimited, limited only by the amount of hardware resources availables, and supports guest virtual machines with arquitectures different to the host machine.

4.3. Cloud Manager

The selected cloud manager was OpenNebula[1]. One big advantage of OpenNebula is the possibility of VM migration of Virtualbox instances. This feature is very useful in an academic environment because, in this scenario, we can provide a better resource usage.

Another important feature of OpenNebula is the modularity of his components. This means that we can change the hypervisor from VirtualBox to another one in future new nodes, without creating big changes in the environment. OpenNebula also has the capability to expand its functionalities by using scripting, without needing to change the source code or to upgrade the environment.

1 OpenNebula: http://opennebula.org

5. Evaluation

5.1. Maximum Load

We define virtual hardware configurations to allow several users to share the same physical hardware without compromising the stability of the server. The specifications of every configuration were defined based on the expected number of users of every instance. To calculate the maximum number of concurrent instances on every node, we performed an evaluation. For a web programming course, a small template with 256 MB of RAM, 10GB of HDD and 0.25 of a CPU core (Aprox 666 Mhz) was enough to serve 50 concurrent students uploading and downloading assets, while they were developing their course assignments. In this scenario, the peak load of the virtual CPU was 5% and the RAM usage was 95 MB, only 37.10% of the memory assigned to the VM.

We expected to be able to run at most 30 VMs per node using this template. We calculate this number by dividing the 8 cores of the processor by the 0.25 used by each VM, minus 0.50 of a core to leave some CPU computing resources for the node.

5.2. CPU Usage

To measure CPU usage and confirm the maximum number of instances previously calculated, we run a test, instantiating 5 virtual machines every 5 minutes to prevent a saturation on the hard disk, until we reach a point higher than 75% of maximum CPU load. Subsequently, we instantiated one extra machine every 3 minutes, until the load exceeded 80%. For this load test, every VM instance was running a loop program to use all the CPU assigned, in order to stress the system. This scenario represents the worst case.

The results of this test are shown in Figure 1, where the vertical axis represents the percentage of CPU usage and the horizontal axis a period of 10 hours, from 00:00 to 10:00. We start the testing at 03:00, and we could see that the figure shows a period longer than the 7 hours of testing, in order to contrast with the idle state of the server.

When we instantiated 34 virtual machines, we found that we could reach the theoretical limit. After this point, we run some connection and response tests to the server, in order to check if we still had control over it with the available resources left. In the results of this test, we can see that the host reached a maximum CPU load of 81.3%, and a 70.2% on average, during the 7 hours of testing.

5.3. RAM Usage

The expected usage of RAM during the test was 8.5 GB for the 34 instances running, this is 34 VM x 0.25 GB/VM = 8.5 GB.

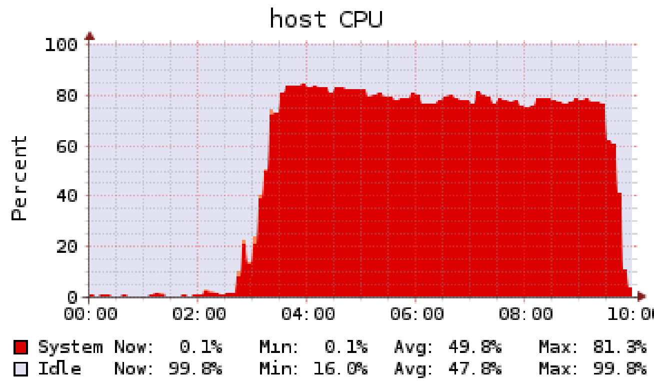

Figure 1. CPU usage, 10 hour time frame.

In Figure 2, the vertical axis shows RAM usage behavior during testing in gigabytes, memory usage reached a maximum of 8.8 GB, the calculated value plus the host RAM usage.

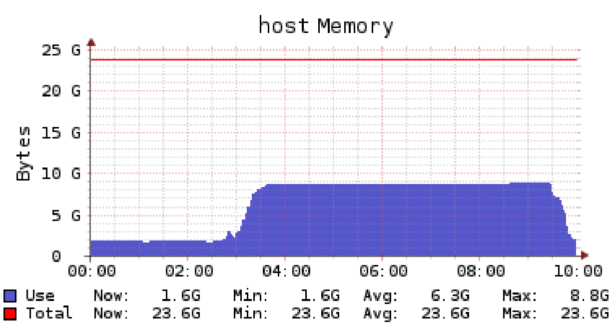

Figure 2. RAM usage during the test.

5.4. Energy Consumption

Another factor considered important to analyze was the power consumed by the host. We used this value to estimate the reduction in costs by using this type of implementation. Figure 3 shows the energy consumption during the test, where the vertical axis represents the host power consumption in Watt (W).

The node in idle state has a power consumption of 105 W. Using the 34 machines previously instantiated the consumption reaches 147.74 W. This means the system is using 42.74 W to run 34 virtual machines at 100% of CPU load, that is equivalent to 1.25 W for each one of the instantiated machines.

To better understand the power consumption impact of these results in an academic environment, we compare an-

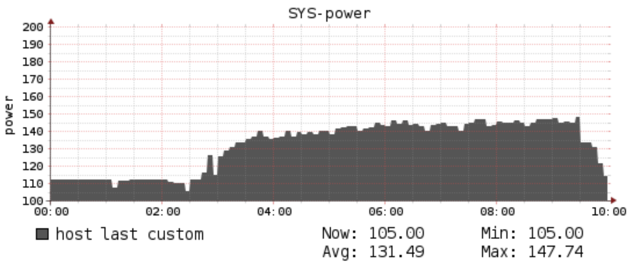

Figure 3. Power consumption of the host.

other two probable scenarios of an academic data center based on the results of the previous test.

In the second scenario, we approximate power consumption of a real server with a processor similar to the used on the instance template, the template used 0.25 of a CPU core of the host, that is comparable to a 667 MHz processor. We use an online tool for PSU calculation [3] and obtained an approximate value of 86 W per server, so, for 34 servers the power consumption will be 2856 W. This is approximately 19.33 times greater than the cloud platform.

For a third scenario, we assume that all of the 34 servers have specifications similar to our host, an actual scenario for some universities, each one with a power consumption of 105 W. In this case, the consumption to have these real servers running rises to 3570 W, a power consumption 24.16 times greater than the cloud.

To translate these results to a monetary cost for a period of one month, we use the price per kWh for a month [1] given by the energy service provider, for the voltage level installed at the Information Technologies building of the campus, which had a value of USD $ 0.13 per kWh. Assuming a 30-day month, equivalent to 720 hours, we obtained the following values for each scenario:

Scenario 1:
1 Server virtualizing 34 machines = $13.82 /month
Scenario 2:
34 servers with low specifications = $273.68 /month
Scenario 3:
34 servers, with host specifications = $334.15 /month

6. Conclusions and Future Work

We were able to implement a Cloud Computing Service using freely available tools, and fulfilling a need that the university had. We tested that it is possible to share resources with members of the community without impacting on the resources already allocated to them, in an way that is economically viable, reproducible and easily scalable.

From the results we got, we have demonstrated that an academic cloud using the Infrastructure as a Service model,

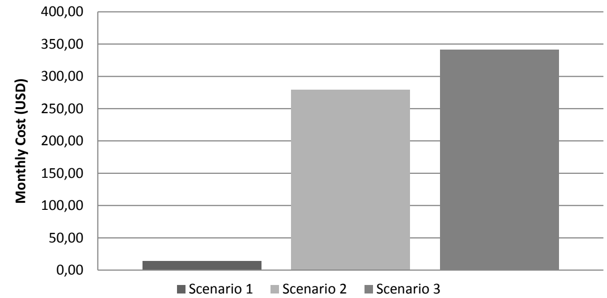

Figure 4. Monthly cost of the energy consumption of the host.

could get better efficiency over its underutilized hardware, saving as much as 24 times only in energy costs.

The development of cloud computing services using the Infrastructure as a Service model for academic communities, using free or open software, represents a great opportunity to use, in an efficient way, the underutilized resources they could have, without having to spend money for licensing costs.

The cloud computing paradigm represents a change in the way an academic institution may do its hardware investments. The increasingly high costs of underutilized computing data centers, represented in expenses of electric consumption, air conditioning, adequacy of facilities and the physical security of equipment, make this kind of project a low-cost solution to make efficient distribution of computing resources.

For the future, we are going to improve our network security model by increasing the network isolation of the VMs and automating some other aspects of the management. We also expect to extend the model over a large academic infrastructure to widely validate our tests.

References

[1] E. de Santander S.A. E.S.P. Tarifas aplicadas al mes de febrero de 2012. http://www.essa.com.co/essa/wrm/tariffs_documents/TARIFAS_ESSA_FEBRERO_2012.pdf, March 2012.

[2] E. Diaz Carreño. Modelo y prototipo de servicios de Computación en la Nube para estudiantes y profesores de la Escuela de Ingeniería de Sistemas e Informática de la Universidad Industrial de Santander. Biblioteca Universidad Industrial de Santander, 2012. Undergraduate Thesis.

[3] Extreme Outer Vision. eXtreme Power Supply Calculator. http://extreme.outervision.com/psucalculatorlite.jsp, March 2012.

[4] J. Geelan, M. Klems, R. Cohen, J. Kaplan, D. Gourlay, P. Gaw, D. Edwards, B. de Haaff, B. Kepes, K. Sheynkman, O. Sultan, K. Hartig, J. Pritzker, T. Doerksen, T. von Eicken,

P. Wallis, M. Sheehan, D. Dodge, A. Ricadela, B. Martin, B. Kepes, and I. W. Berger. Twenty-One Experts Define Cloud Computing.

[5] R. P. Goldberg. *Architectural Principles for Virtual Computer Systems*. PhD thesis, Harvard University, 1972.

[6] G. Gruman. Gartner Says Cloud Computing Will Be As Influential As E-business. *Gartner Newsroom*, 2008.

[7] A. Khajeh-Hosseini, D. Greenwood, and I. Sommerville. Cloud migration: A case study of migrating an enterprise IT system to IaaS. In *Cloud Computing (CLOUD), 2010 IEEE 3rd International Conference on*, pages 450–457, 2010.

[8] R. Maggiani. Cloud computing is changing how we communicate. In *Professional Communication Conference, 2009. IPCC 2009. IEEE International*, pages 1–4, 2009.

[9] P. Mell and T. Grance. The NIST Definition of Cloud Computing. Technical report, Sep 2011.

[10] J. Staten. Is Cloud Computing Ready For The Enterprise? Technical report, Forrester Research, Inc., Cambridge, USA, 2008.

[11] M. Steinbauer, I. Khalil, and G. Kotsis. Building an on-demand virtual computing market in non-commercial communities. In *Proceedings of the 28th Annual ACM Symposium on Applied Computing*, SAC '13, pages 351–356, New York, NY, USA, 2013. ACM.

[12] L. M. Vaquero, L. Rodero-Merino, J. Caceres, and M. Lindner. A break in the clouds: towards a cloud definition. *SIGCOMM Comput. Commun. Rev.*, 39(1):50–55, dec 2008.

[13] I. Waßmann, D. Versick, and D. Tavangarian. Energy consumption estimation of virtual machines. In *Proceedings of the 28th Annual ACM Symposium on Applied Computing*, SAC '13, pages 1151–1156, New York, NY, USA, 2013. ACM.