Cloud Management Platform Selection: A Case Study in a University Setting

Francisco Cocozza, Gustavo López, Gabriela Marín, Ricardo Villalón and Francisco Arroyo

Centro de Investigaciones en Tecnologías de la Información y Comunicación (CITIC)
Universidad de Costa Rica
San José, Costa Rica
Email: {francisco.cocozzagarro; gustavo.lopez_h; gabriela.marin; ricardo.villalon; rafael.arroyo}@ucr.ac.cr

***Abstract*—Cloud Computing has gained importance in recent years. There are many implementations' analyses and evaluations of Cloud Management Platforms (CMPs) in the literature. Moreover, the context and characteristics differ drastically between implementations, depending on the user requirements and usage context. This paper presents a case study of the process we followed to select a Cloud Computing management platform to be deployed in a University. Administrative and academic requirements were gathered and studied to define the most appropriate platforms. We present an overview of available CMPs. Moreover, we show a set of comparison criteria that could be used to determine which CMP adapts best to the cloud deployment scenario.**

Keywords–Cloud Computing; VCL; OpenStack; Cloud Management Platform.

I. INTRODUCTION

Requirements for computational services in the industry and academia have grown vastly in the past decades. Cloud Computing is one of the efforts conducted by the information technologies and computational scientists to keep up with the demand. The National Institute of Standard and Technology (NIST) of the U.S. Department of Commerce defines Cloud Computing as a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction [1]. The essential characteristics of Cloud Computing are: (1) on-demand self-service, (2) broad network access, (3) resource pooling, (4) elasticity, and (5) measured quality of service [2].

There are many possible settings for a cloud, e.g., universities, commerce or multinational companies. Even though nowadays there are several public clouds offering services, we consider important to describe the process of developing a cloud depending on the context in which it will be deployed.

Our research is motivated by the fast growth of Costa Rican industry in areas related to technology and services. Since 1998 more than 100 companies focused on IT related services, established operations in Costa Rica, including companies like IBM, Sykes, Infosys, VMWare, Hewlett-Packard, Intel and others. Moreover, 5.8% of national GDP consists of IT & IT Services. Without doubt, Cloud Computing is currently undergoing a huge hype and many companies see it as the future of IT [3]. This situation creates a constant demand of trained and qualified professionals. Being the major University in the country, University of Costa Rica (UCR) needs to deal with the training of these professionals. In order to train professionals in Cloud Computing and related topics, the University required to implement its own cloud services, and create know-how on the trending topics. The main effort conducted by UCR was deploying a cloud and create coursework for the topic.

These paper aims to provide an experience report and a set of tools that could be helpful when deploying a cloud. This research project was conceptualized after a training in which the benefits of Cloud Computing for universities were presented. The presenter explained how Cloud Computing through the implementation of Virtual Computing Lab (VCL) improved IT services at North Carolina State University. Several research papers were published by VCL partners. They demonstrated how VCL reduced costs and help delivering computational resources to over 30,000 students and faculty members [4][5]. We decided to reproduce their effort at UCR.

A cloud is a complex entity wich consists of several components such as hypervisors, authentication mechanisms, file system backends, certificate authorities, data base engines, and others. An implementation of such elaborated infrastructure involves an extensive number of decisions, which cannot be covered and described on a single paper. Given that, this work is focused on the process we followed to select the software platform for managing the UCR's cloud.

In the process of creating our cloud, we found several CMPs comparisons in the literature. However, these comparisons vary greatly in terms of granularity and the concepts addressed. Therefore, we proposed a set of criteria to consider when implementing a cloud. Our criteria are based on literature review and expertise obtained while implementing UCR's cloud.

The comparison of CMPs is a complex labor, identify the characteristics of software depends on the completeness of its documentation and case studies reported. Therefore, from all the considered CMPs we applied our comparison criteria only in four preselected ones: CloudStack, Eucalyptus, VCL and OpenStack.

The rest of the paper is structured as follows. Section two shows the related work, focused on an overview of available CMPs and comparisons reported in literature. Section three gives the context in which UCR's cloud was developed. Section four shows the process we followed to select a CMP and the key features that guided our selection. Section five shows the selection process results and discussion. Finally, Section six shows some conclusions and future work.

Copyright (c) IARIA, 2015. ISBN: 978-1-61208-388-9

II. RELATED WORK

According to a recent Merrill Lynch research note [6], Cloud Computing is expected to be a $160-billion addressable market opportunity. Cloud Computing is also a prominent technology trend [7]. Given this large market, many companies with a desire for profit and other nonprofit organizations presented a variety of CMPs in the past years.

A CMP is a set of software for managing cloud environments [8]. A CMP includes self-service interfaces, provisioning system images, usage measure and workload optimization. CMP's main goal is to allow enhanced resource management and monitoring of the cloud resources.

Several CMPs are available nowadays: Abiquo [9], CloudStack [10], Eucalyptus [11], Nimbus [12], openQRM [13], Openstack [14], Open Nebula [15], Apache Virtual Computing Lab (VCL) [16], HP's CloudSystem Matrix [17], among others. In this Section, we present a brief introduction to these CMPs and their main deployment scenarios.

Abiquo [9] is a hybrid cloud management system for small and medium-sized business. Abiquo supports multiple hypervisors and is focused on enabling organizations to leverage existing virtualization technologies and public clouds. Abiquo also allows central management of resources via Graphical User Interface (GUI).

Apache CloudStack [10] is an open source software designed to deploy and manage large networks of virtual machines as Infrastructure-as-a-Service (IaaS) Cloud Computing platform. CloudStack is used by service providers running cloud services, product vendors and organizations who have used the software to deploy private clouds.

Eucalyptus is a Linux based open source software. Eucalyptus is the acronym for Elastic Utility Computing Architecture Linking Your Programs to Useful Systems. It was developed for creating private and hybrid clouds. The software is suited for enterprise clouds that supports the industry-standard. [18][11].

Nimbus is a set of open source software Cloud Computing components written in Java and Python focused on providing IaaS capabilities to the scientific community [12]. Nimbus is designed to turn clusters into an Infrastructure as a Service cloud.

OpenQRM [13] is a free and open-source Cloud Computing management platform for managing heterogeneous data center infrastructures. The openQRM platform manages a data center's infrastructure to build private, public and hybrid IaaS clouds. OpenQRM is designed for companies of various kinds.

OpenStack is a free and open source Cloud Computing software platform. OpenStack software controls large pools of compute, storage, and networking resources. OpenStack is designed to be deployed in many settings. OpenStack's use is largely reported in various research papers [19][20][21]. Moreover, 35 case studies are reported by OpenStack in the use of the software [14].

OpenNebula [15] is an open source project aimed at building the industry standard open- source Cloud Computing tool to manage the complexity and heterogeneity of large and distributed infrastructures.

Apache VCL [16] is an open-source solution for the remote access over the Internet to dynamically provision and reserve computational resources for diverse applications, acting as Software as a Service (SaaS) solution. VCL was conceived as a tool for educational use and was first deployed at East Carolina University, USA.

HP CloudSystem Matrix [17] is a cloud infrastructure from Hewlett-Packard that combines storage, servers, networking and software for organizations to build complete private, public and hybrid Cloud Computing environments.

Knowing the quantity of CMPs available several researches have been conducted to compare different CMPs. However, these comparisons are too different from one another to be used together when choosing a cloud management platform. For instance Cordeiro et al. [22] compare Xen cloud platform, Eucalyptus and OpenNebula. The comparison is based on the platform's architecture, their networking management, virtual machine placement and inter-host communication. On the other hand, Voras et al. [23][24] present a set of comparison criteria based on storage, virtualization, management, network, security and support. The authors introduce several cloud management platforms and other cloud related software such as Open Nebula, Eucaliptus, Ubuntu Enterprise Cloud, OpenQRM, Abiquo, Red Hat Cloud Foundations, Edition One, OpenStack, Nimbus, mOSAIC. However, no evaluation is reported.

Wind [25] presented another comparison between Eucalyptus, OpenNebula, Abicloud and Nimbus based on these criteria: architecture, programming language, supported cloud types and hypervisors, user interface, licensing, robustness, interoperability, security and compatibility. This research presents an interesting comparison criterion, however, only four CMPs were evaluated and the more used platforms were not included. The criterion used in our research includes several of the criteria presented by Wind, but applied on other CMPs. Cerbelaud, Garg and Huylebroeck [26] compared Enomaly Elastic Computing Platform (ECP) Eucalyptus, OpenNebula and oVirt based on the following criteria: VM creation tool and repository, image storage, uploading, saving and choosing host. Steinmetz, Perrault, Nordeen, Wilson and Wang compared OpenStack and Eucalyptus; however, their measurement was solely based on performance.

Other authors have compared cloud vendors and providers. For instance, Khan, Noraziah, Herawan, and Mat Deris [27] compared Amazon EC2, Microsoft Azure, Google App Engine, Sun Grid and GRIDS Lab Aneka based on: service type, user access, virtualization capabilities, and programming framework. A similar comparison was made by Li. Yang, Kandula and Zhang [28] they compared Amazon AWS, Microsoft Azure, Google AppEngine and Rackspace CloudServer. However, the authors used a totally different set of criteria, making impossible any comparison between the two. The last choose elasticity, storage, internal communication, networking and cost as their criteria.

This section showed the variety of CMPs and cloud vendors in the market. Moreover, we introduced most of the currently available CMPs. On the other hand, we presented some comparisons available in the literature and proved that the compared CMPs and criteria used are very difficult if not impossible to compare. The main differences found were:

- Heterogeneity: we discovered different terms referencing the same concept. This can be associated with the lack of standardization on cloud platform features names.

- Granularity: some comparisons provided detailed criteria about a specific feature, while other comparisons evaluated high level characteristics of the cloud platform.
- Completeness: the criteria did not meet our expectations on the level of detail and covered features that we wanted to evaluate.

Given the lack of homogeneity and completeness, and the difference in granularity, we decided to use our own set of criteria. Our definition is based on previous comparisons and our contextual requirements.

III. CONTEXT

This section details the context and background in which the UCR's cloud was developed and presents an overview of processing power and capacity required by the cloud services to be deployed.

A. *Demographics and physical contexts*

UCR is the largest and most important Costa Rican university. Established in 1940, it has three main focuses: research, teaching and being socially responsible. It is a public university with a budget of $454 million USD for 2014 [29]. UCR's academic offer includes 244 undergraduate programs and 243 graduate programs (including masters and doctoral programs).

With over 40,000 active students, 4,292 faculty members and 3,520 administratives, UCR has large computational requirements. UCR is spread all over Costa Rica. The main campus is in the capital of the country, also offers seven secondary campuses and four local branches (smaller facilities). The IT Department in charge of supplying these requirements was starting to exceed their capacity.

UCR's physical and demographical complexities lead us to consider all these factors when creating a computational solution that meets academic and administrative requirements.

B. *University IT Services*

The services provided to the academic and administrative population of the University are in charge of the IT Department. It offers a variety of services ranging from networking and infrastructure maintenance to website hosting.

The IT Department implemented, some years ago, virtualization hosting services to fulfill the diverse university's requirements. As the number of requests increased, the services performance and capacity became a problem. The main problem was that the hardware was not able to keep up with the demand. With 140 virtual machines distributed among eight virtualization servers, the platform had reached its limit by the end of 2012. In 2013, the IT Department deployed a strategic plan for continuous improvement in which they decided to update their hardware.

The main services provided by the IT Department are: web hosting, virtualization, mail services, DNS services, network infrastructure maintenance, database management, computational equipment provisioning. Nowadays, these services are manually requested and implemented; therefore, any automation would be beneficial.

One of the main goals of the IT Department is to decentralize the administration of the resources they provide. This goal can be achieved by offering self-provisioning services (one of the main features of a cloud). NIST defines self-provisioning or self-service as: the ability of a user to unilaterally provision computing capabilities, without human interaction of the service provider [1]. Self-provisioning is a particularly important aspect because it enhances the perception of service's efficiency and makes the management and use of the resources easier for both, the service provider and the end user.

Any change in IT services causes the raise of disagreements. We faced this reality when implementing UCR's cloud platform. Based on the research presented by Kuo [30], we established the main challenges and opportunities for this cloud implementation (Table I).

Organizational inertia relates to the resistance to change by not only the end users, but also the people in charge of maintaining and supporting the services. The lack of experience related to Cloud Computing was also a challenge. Even though most of the people implementing the cloud are experienced professionals in virtualization, network and infrastructure areas, they did not have proper training in Cloud Computing at the beginning of this project.

Costa Rican jurisdiction on digital data is still very ambiguous. This causes security and data jurisdictional concerns needed to be addressed in order to use international cloud services. Moreover, University of Costa Rica's infrastructure may cause a bottleneck that could be a challenge when deploying cloud services.

All the challenges mentioned above were considered; however, we had an opportunity windows that we were encouraged to use. The main opportunities were: financial support of high authorities at UCR and a vision for change in the administration. Many offices inside UCR requested IT infrastructure improvements, the creation of a cloud could take advantage of the existing infrastructure and share resources in order to satisfy most of these needs. Moreover, at country level, this project represents a great opportunity. Currently the IT services industry in Costa Rica is a sophisticated and continuously growing area.

According to the Costa Rican Investment Promotion Agency (CINDE) [31], companies are constantly demanding trained and qualified professionals in areas like Cloud Computing, Big Data, Virtualization, and others. These market requirements forced UCR to improve its educational offer in Cloud Computing and related technologies. Creating a cloud could provide the expertise and new technological infrastructure required for this improvement.

TABLE I. CHALLENGES AND OPPORTUNITIES CONSIDERED IN THIS CLOUD IMPLEMENTATION

Challenges	Opportunities
- Organizational inertia. - Lack of expertise in Cloud Computing. - Security and data jurisdiction issues. - Infrastructure bottlenecks.	- Financial support for hardware acquisition. - Visionary and change-aware administration. - IT Infrastructure needs. - Growing IT services industry requiring trained and qualified professionals.

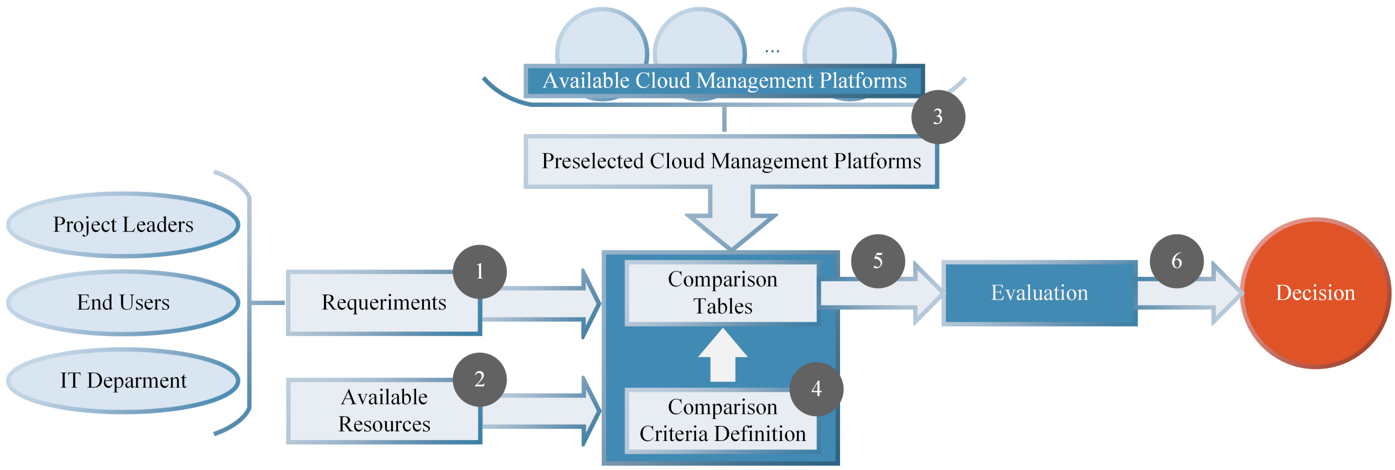

Figure 1. Process followed to select a Cloud Management Platform.

IV. Cloud Management Platform Selection Process and Implementation

This section presents the process followed to select a cloud platform used to meet UCR's computational requirements. Several cloud platforms exist. We wanted to select the most appropriate one to UCR's context, taking into account important variables given by the context, institutional regulations, legal issues and user requirements.

Figure 1 shows the components and steps involved on the process to select a cloud platform. To summarize, the process consisted of the following steps:

1) Gather the requirements and expectations for the cloud platform from different perspectives; i.e., the stakeholders.
2) Determine the resources (hardware and software licenses) already available to support the cloud platform, and determine their technical specifications.
3) Identify available cloud platform solutions and preselect a subgroup based on basic filtering, based on the most relevant criteria.
4) Define our own set of comparison criteria, based on other comparisons and information gathered in the previous steps.
5) Compare the cloud platforms based on our criteria.
6) Select the cloud platform that best fits our requirements.

The following subsections provide an in-depth description of the implementation of each one of the six steps mentioned above.

A. *Requirements (step 1)*

The CMP implementation is one of the projects of the IT Department. For all purposes, the IT Department plays the role of Service Provider (SP). One of their main goals is to provide high quality services that can meet the expectations and requirements of their users. As a SP, the IT Department has to ensure that the platform is flexible in its service delivery while keeping the users isolated from the underlying infrastructure [27].

Profiles for users in a university setting are highly diverse. Therefore, we determined that the requirements for this project should be gathered from different perspectives. As shown in Figure 1, the "requirement component" includes the feedback of three different groups: (1) IT Department, (2) Cloud Computing project leaders and (3) end users.

For collecting the requirements from the IT Department, different approaches were used. To understand the internal processes of the IT Department we established periodical technical meetings. After we had a clear vision of the internal functioning of the IT Department, we met with specialized technicians that could contribute in the definition of the cloud. The topics discussed in the meetings were widely varied, ranging from network topologies to virtualization technologies. The main goal of the meetings was to shape and refine the required cloud services.

The Cloud Computing project managers and researchers involved in this project were the ones considered to establish the requirements for academic purposes. The team in charge of the project includes seven computer scientists (4 Ph.D., 2 M.Sc, and one B.Sc.).

For the end users perspective, a different approach was used. In this case, a survey was applied to the local IT administrators of every department and faculty to gather their opinion and expectations on the new services. With the input provided on this survey, we identified some of the most requested services for the cloud. Figure 2 shows the main services requested and the percentage of the end users that wanted each service. The perspectives mentioned were consolidated to establish the following requirements:

The general requirements are:

- Licensing: the cloud platform must have an Open Source licensing model.
- CMP Database Management System: the selected DBMS must be one of the following: MySQL, MariaDB and/or PostgreSQL, which are Open Source DBMS.

The academic requirements are:

- Lab deployment support: the cloud must offer the ability to deploy virtual computer laboratories [32] for academic courses.

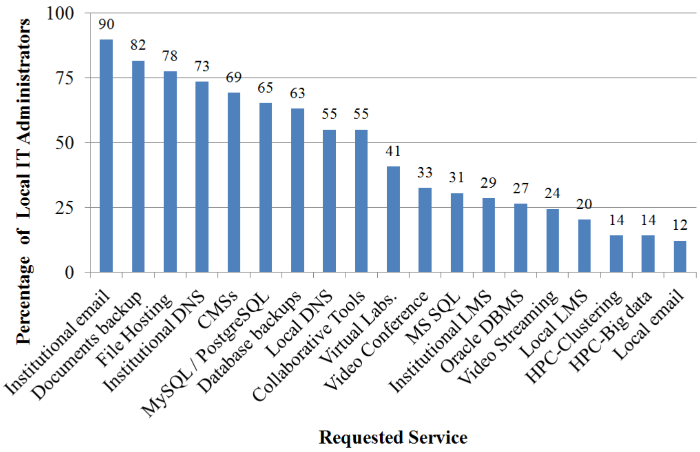

Figure 2. Cloud services requested by users (local IT Administrators)

- Bare-metal provisioning: the cloud operative system should have built-in support for automated bare-metal provisioning. Bare-metal provisioning is the process of installing an Operating System (OS) or Type 1 hypervisor directly on a computer's hardware [33].
- Scheduling of resources reservation: it covers a typical academic scenario, where the same type and number of virtual machines are needed on a specific time and on a regular basis.

Finally, the administrative requirements are:

- High availability for critical services.
- Infrastructure compatibility: the cloud platform must be aligned with the current hardware and software. For example, the currently licensed hypervisors must be supported by the cloud platform. Also, integration with specific appliances is preferred.
- Service model: the cloud platform must provide the most basic service model, Infrastructure as a Service (IaaS). This was defined by the IT Department and us (the researchers) as the initial stage of the cloud service. Based on IaaS, other service models like Software as a Service (SaaS) or Platform as a Service (PaaS) could be offered in the future.
- Platform management: the selected platform must provide features that facilitate its use and administration.
- API completeness: the services and functionality provided by the CMP's API must be extensive and exposed using standards, e.g., REST designed [34].

B. Available Resources (step 2)

The second step in the procedure is to determine the available hardware. In a cloud, the hardware plays a key role on the success and reliability of the provided services.

The IT Department at UCR assigned its newest hardware for the implementation of this project. The process followed to select the hardware is out of the scope of this paper. However, it is worth mentioning that the decisions related to hardware specifications were based on technical studies and evaluations. The main characteristics of the hardware dedicated to the cloud platform are:

- Servers: 75 blade servers (6 enclosures). Each blade has:
  - 128 GB of RAM
  - Dual 8 Core at 2,4 GHz
  - 2 146 GB SAS discs
- Storage appliance: the particular architecture of the storage system consists of three different layers, similar to the hierarchical memory model of any modern computer. These layers are composed by NL-SAS, SAS and SSD discs, which in total provide about 400 TB of RAW storage capacity.
- Network appliances: two dedicated advanced switches that offer a throughput of 24.3 Tbps per appliance, Software Defined Network (SDN) capabilities, layer 2 and layer 3 support for large deployments, and high availability, supported by their redundant architecture.

The selection of the CMP must be aligned with the hardware on which is going to be deployed; trying to take advantage of the specific characteristics that each component can provide. The final decision can also be substantiated on low level features, like hardware compatibility or integration with CMP.

At the software side, the University already has under its acquisitions some licenses for proprietary software, ranging from DBMS to virtualization platforms. The IT Department has deployed an Oracle database cluster, which is actively used by most of the University's software systems. In the case of virtualization software, the University has enough VMware vSphere licenses. These specific platforms should influence the selected cloud platform, in order to take advantage of already acquired software and know-how.

The following section describes the procedure to identify the available cloud platforms and how a subset of those solutions was chosen for the next step, the comparison of the software.

C. Preselection Process (step 3)

No matter how good the hardware is, its value depends on the software that manages it. Therefore, the selection of the software is crucial. Selecting the software that will run the cloud involves the research of available cloud platforms.

An initial review was done to determine the existing options and trends. We identified a variety of options. These options includes all the CMPs mentioned in the Related Work (Section II).

At this stage, we were not sure if all alternatives were suitable for our context, so, a filtering process was performed. We defined the following criteria needed to be met by UCR's cloud:

- CMP must provide IaaS model.
- Licensing model of the CMP must be open source: University's internal regulations dictate that the use of open source software should be preferred [35].
- CMP must be able to use our own hardware and infrastructure.
- CMP must be perceived as a solution with good evidence of usage and support: the access to documentation and case studies found for each of the platforms were used as parameters for this filter.

Copyright (c) IARIA, 2015.

- Selected CMP must be robust and well known: external references to platform software and its use in industry (by large organizations) were factors used to measure the popularity of the product.
- Application scenario: the platform software must fit the specific purposes of the University. As stated before (Section IV-A), administrative/enterprise and academic/educational requirements must be met.

The preselected platforms were picked by the project research group and some of the IT Department's staff directly related with the project. At the end of the filtering process, the preselected CMPs were CloudStack, Eucalyptus, Openstack and VCL.

Openstack, CloudStack and Eucalyptus were chosen mainly because of the numerous evidence showing its use in enterprise scenarios. VCL is a CMP more focused on academic or educational areas, given the context and requirements on which the cloud is going to be deployed, this platform was of special interest for our evaluation.

Once the preselected platforms were established, we proceed to evaluate them using the comparison instrument that is going be described on the next section.

D. Comparison Criteria (step 4)

Based on the criteria presented by other comparisons [22][23][24][25][26][27][28], a preliminary set of criteria was defined. Some features or characteristics of concern were not present on other existing cloud platform comparisons. We added some new evaluation criteria according to our interests and evaluation purposes.

The criteria were grouped in two different categories: technical features and management features. Figure 3 shows detailed used criteria.

The technical features involve authentication mechanisms, networking, storage and other technical aspects. These features were also divided into eight categories. The general features involve the architecture of the platform, the supported DBMSs and the language in which it was developed. Authentication features incorporate all security protocols. Computational features addressed the supported hypervisors and deployment techniques, and high performance computing support.

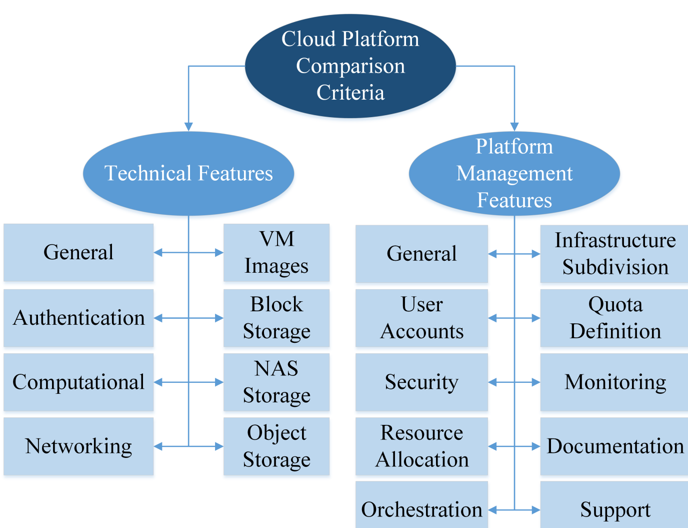

Figure 3. Criteria classification structure

Networking features are related with IP version management, traffic isolation, and remote access capabilities. Virtual machine features were related with virtual machine images management and supported formats. Block, network attached and object storages evaluated all the available characteristics and actions for storage.

Management capabilities involve general features such as licensing, management interfaces, compatibility with other clouds through APIs and virtual machine administration. Management also involve user accounts and security features (privileges, roles and permissions).

Resource allocation, orchestration and infrastructural management features manage the distribution and scheduling of virtual machines and networking and storage components. They also managed the segregation of the infrastructure.

Quota definition establishes the usage limitation of resources for each user depending on roles and other aspects. Monitoring features involve the measurement of the cloud platform components. Finally, documentation and support features gives an idea of the usability and reliability of the evaluated cloud platforms.

E. Platforms Evaluation (step 5)

On the previous step a set of criteria was defined. The features selected for comparison were based on the parameters defined by other comparisons, and by our specific requirements. Once the comparison was made, a particular cloud platform was assigned to each of the project researchers. Each researcher was responsible for collecting the information of a specific platform. There was a cross validation between researchers to verify the gathered information. Every inconsistency was validated and resolved by the group.

The evaluation presented in [36], is a comparison table that shows which requirements are met and the main characteristics of each platform. The cloud software platforms included in the comparison were preselected on step 3 (Section IV-C), and the criteria used to evaluate those platforms were defined in step 4 (Section IV-D). The evaluation of characteristics was the last step of the procedure.

Tables II and III are segments of the complete CMPs comparison. Table II shows some of the comparison criteria related to technical features, and Table III refers to part of the evaluated management features. The extract in this paper only considered features in which Eucalyptus, OpenStack and CloudStack differ. The full version of the tables [36] shows not only all the criteria evaluated but also VCL is incorporated in the table.

With all the collected data (requirements, hardware specification and context variables) the final decision was made. The decision was made by the implementing team; composed by the seven researchers previously mentioned, the head of the IT Department, three Cloud Computing project leaders, and several collaborators (approximately 20 people).

The selection process described in Section IV let us define key factors to take into consideration. The evaluation of the cloud management platforms using the comparison tables

TABLE II. TECHNICAL FEATURES

General Features			
Feature	**Eucalyptus**	**OpenStack**	**CloudStack**
Version under revision	3.4.1	Havana, version 8 (2013.2.x)	4.2.0
Development language	Java, C, Python, Perl	Python	Java
Supported DBMS	PostGreSQL	Drizzle, Firebird, Microsoft SQL Server, MySQL, Oracle, PostgreSQL, SQLite, Sybase	MySQL
Supported protocols and backends	Local users, IAM	LDAP, Kerberos, Local Users	LDAP, local users
Computational Features			
Feature	**Eucalyptus**	**OpenStack**	**CloudStack**
Supported Hypervisors	ESXi, KVM	KVM, QEMU, Xen, ESXi/VC, Hyper-V, PowerVM, Docker	VMware, KVM, XenServer y Xen Cloud Platform
Scheduling methods	EuQoS	By filters and weights, random (filtered), customized	Not Found
Networking Features			
Feature	**Eucalyptus**	**OpenStack**	**CloudStack**
IP v6 management	No	Not Found	No
Projects Traffic Isolation method	VLAN	VLAN, GRE, VLAN+GRE	VLAN
Remote Desktop access protocols	VNC	VNC (noVNC), SPICE	VNC
VM Images Features			
Feature	**Eucalyptus**	**OpenStack**	**CloudStack**
Supported disk formats	Raw, VMDK, VDI, QCOW2, ISO	Raw, VHD, VMDK, VDI, ISO, QCOW2, AKI, ARI, AMI	QCOW2, VHD, VMDK
Supported container formats	emi	Bare, ovf, aki, ari, ami	Not Found
Supported backends	Posix Filesystem	Local disc and NFS, Swift, S3, RBD	NFS, Swift, S3
Image caching	Yes	Yes	No
Create Image from VM snapshot	No	Yes	Yes
Block Storage Features			
Feature	**Eucalyptus**	**OpenStack**	**CloudStack**
Supported backends	NFS, DAS, EBS, iSCSI	NFS, LVM, Ceph, GlusterFS, and specific drivers (VNX, DELL, etc)	iSCSI, NFS
Clone a volume	Yes	Yes	No
Object Storage Features			
Feature	**Eucalyptus**	**OpenStack**	**CloudStack**
Backend	EBS, Walrus	Swift	Swift (OpenStack) and Amazon's S3
Access method	REST-Ful	REST-ful API & python-swiftclient	Just for VM Images and Volume Snapshots
Segmentation support (zones, groups)	No	Zones Rings Regions	Zones Rings Regions
Authentication integration support	GNU	Apache 2.0	Apache 2.0
Web UI integration	Yes	Yes	No

TABLE III. MANAGEMENT FEATURES

General			
Feature	**Eucalyptus**	**OpenStack**	**CloudStack**
Licensing	GNU	Apache 2.0	Apache 2.0
REST-ful API	Yes	Yes	No
User Accounts			
Feature	**Eucalyptus**	**OpenStack**	**CloudStack**
Permissions granularity	Groups, Users	Tenant, Users	Domain, Account, Domain Administrators, Projects
User with multiple projects or groups	No	Yes	Yes
Security			
Feature	**Eucalyptus**	**OpenStack**	**CloudStack**
Least privileged access design	Yes	Yes	Not Found
Fine granularity permissions definition method	JSON	JSON	On web UI
Centralized permission control	No	No	On web UI
Resource Allocation			
Feature	**Eucalyptus**	**OpenStack**	**CloudStack**
Virtual Machines block provisioning	Yes	Yes	Not Found
Orchestration			
Feature	**Eucalyptus**	**OpenStack**	**CloudStack**
Complex architectures definition support	Yes	Yes	No
Autoscaling support	No	Yes	Yes
Web UI integration	No	YEs	No
Architecture definition formats	Ansible	Heat Orchestration Template	No
Infrastructure Segregation			
Feature	**Eucalyptus**	**OpenStack**	**CloudStack**
Logic division	Availability Zones	Availability Zones >Host Aggregates	Zones >Pods >Clusters
Physical division	Regions	Cells and Regions	Regions
Quota Definition			
Feature	**Eucalyptus**	**OpenStack**	**CloudStack**
Volume quantity	Yes	Yes	Not Found
# of floating IPs	Yes	Yes	Not Found
# of security rules	Yes	Yes	Not Found
Monitoring			
Feature	**Eucalyptus**	**OpenStack**	**CloudStack**
Web UI integration	No	Yes	Yes
Centralized module	No	Yes	Yes
Alarms definition support	No	Yes	Not Found
Computational resources monitoring	No	Yes	Yes
Networking resources monitoring	No	Yes	Yes
Images monitoring	No	Yes	No
Object Storage monitoring	No	Yes	No
Volumes monitoring	No	Yes	Yes
Energy monitoring	No	Yes	No
Visualizations included	No	Yes	Yes
Extendible (new metrics)	No	Yes	No
Documentation Support			
Feature	**Eucalyptus**	**OpenStack**	**CloudStack**
End user guide	Yes	Yes	Not Found
Support forum	Yes	Yes	No
Available certifications	Yes	Yes	No

made us realize that none of the preselected CMPs complied with all University's requirements. In the next section we describe how we addressed this issue.

V. Results and Discussion

The selected platforms were Openstack and VCL. Openstack was selected to fulfill the administrative requirements, and VCL to address the requirements associated to the academic area of the university. The decision to deploy two cloud platforms was taken primarily because none of them completely met the defined requirements (Section IV-A).

A characteristic of Openstack that stood out was the architecture design, specifically the method used for communication among its various components. The approach of a message queue as a general resource for inter-platform interaction was perceived as a key feature and a highly desired characteristic.

Also, cloud platform integration with our specific hardware was an important factor taken in account. Openstack block storage offers a driver to directly interact with our storage appliances (EMC VNXs).

Similarly, the SDN capability of Openstack matches the SDN support offered by our networking appliances. Openstack's Object Storage, known as Swift, offers scalability, which is required for the platform, given the amount of data to be stored. This is reflected on Figure 2, which indicates that backup service was one of the most required services by the end users. Swift is distributed, this means that data can be geographically distributed and replicated among the university's campuses.

Finally, Openstack has a solid base of community members, which enhances its perception as a top Cloud Operative System. This is evidenced by the contribution and involvement of big players such as Hewlett-Packard, Red Hat, Canonical, and others [14]. All the previously mentioned reasons made Openstack the selected platform to fulfill the administrative requirements.

For VCL, the main reason for its selection was the specific features that support the particularities of an academic scenario. This functionality was not supported by Openstack. As stated on Section IV-A, deployment of virtual labs was a mandatory functionality, this is covered by VCL. Also, VCL offers features for reservation of virtual machines based on given time and day. VCL has been already implemented in other universities; this is the case of North Carolina State University (where VCL was first implemented) [5]. We perceived this acceptance as an evidence of the compliance of this platform with academic requirements.

Even when the comparison tables provide an exhaustive description of each of the evaluated platforms, the final decision should not rely upon this as the final instrument. Instead, we suggest a conscious research and understanding of the preferred platforms. The main issue is that the tables do not provide an scaled evaluation, instead, they only provide a general overview that should be discussed to make a decision.

We presented a set of criteria used in order to select the best cloud platform for our scenario. These criteria could be used as a starting point for a cloud platform selection process, like the one presented in this paper. In [36], we show an in-depth vision of the platforms assessment. These comparison could be expanded/modified in order to suit specific interests and requirements depending on the implementation context.

The value of the presented process and comparison tables used is not the actual result of the selection process (the selection of Openstack and VCL), because it will certainly change depending on context and specific requirements. The real contribution of this research is the process and a set of criteria that could be used as a starting point to compare cloud platforms focused on the context.

VI. Conclusions and Future Work

We presented an implementation of a cloud platform to support administrative and academic requirements in an university setting. Our implementation was based on six steps followed to select the most appropriate cloud platform.

The first step, gathering requirements was addressed by performing a survey to final users, asking the IT Department for their administrative requirements and setting academic requirements with a panel of researchers. The second step helps to identify available resources (hardware and software), this step also allows to determine if new technologies are required. The third step we performed was the identification of all available platforms. Since the next steps in the selection process were intense and required lots of work, a basic filtering process was performed. We checked compliance with mandatory requirements: open source licensing model, implementation of IaaS and compatibility with the available hardware. Once we had a set of platforms that adapt to our mandatory requirements, a comparison between them allowed the selection of the best one(s). UCR's cloud uses OpenStack to support administrative requirements and VCL to support academia related requirements.

Every implementation of a cloud is very different. Thus, a standard definition of criteria is not possible; however, the criteria that we presented on [36] shows how we incorporated other works criteria and our own contextual requirements. The final steps are the evaluation of each possible platform with the criteria and the selection of the best one(s). We integrated various available CMPs comparisons available in literature and extended them for CloudStack, OpenStack, Eucalyptus and VCL.

Our future work consists of evaluating the process with other organizations to create a more robust process description and set of criteria. Other CMPs could be incorporated to the comparison tables to expand the selection possibilities. Taking into account that several comparisons are available in the literature but they widely differ in concepts granularity and completeness, we propose that a base criteria must be used as a comparison and any additional data should be added to the tables. If the tables presented in this research are expanded with more CMPs we could create a base line to compare CMPs based on a standard set of criteria.

Acknowledgment

We thank the IT Department at UCR for its collaboration on this research project. This work was supported by CITIC (*Centro de Investigaciones en Tecnologías de la Información y Comunicación*) at *Universidad de Costa Rica*, grand No. 834B3145 and by the *Escuela en Ciencias de la Computación e Informática* at *Universidad de Costa Rica*.

REFERENCES

[1] U. D. of Commerce, "National institute of standards and technology," 2014, URL: http://www.nist.gov/ [retrieved: 07, 2014].

[2] C. Baun, M. Kunze, J. Nimis, and S. Tai, Cloud Computing: Web-Based Dynamic IT Services, 1st ed. Springer Publishing Company, Incorporated, 2011.

[3] S. Kumar, S. Saravanan, and S. Mukherjee, "Recommendations for implementing cloud computing management platforms using open source," International Journal of Computer Engineering and Technology (IJCET), vol. 3, 2012.

[4] M. Vouk, "Cloud computing #2014; issues, research and implementations," in Information Technology Interfaces, 2008. ITI 2008. 30th International Conference on, June 2008, pp. 31–40.

[5] H. Schaffer, S. Averitt, M. Hoit, A. Peeler, E. Sills, and M. Vouk, "Ncsu's virtual computing lab: A cloud computing solution," Computer, vol. 42, no. 7, July 2009, pp. 94–97.

[6] D. Hamilton, "Cloud computing seen as next wave for technology investors," 2008, URL: http://www.financialpost.com/money/story.html?id=562877 [retrieved: 01, 2015].

[7] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, "Cloud computing and emerging it platforms: Vision, hype, and reality for delivering computing as the 5th utility," Future Generation Computer Systems, vol. 25, no. 6, 2009, pp. 599 – 616.

[8] M. Caballer, D. Segrelles, G. Molto, and I. Blanquer, "A platform to deploy customized scientific virtual infrastructures on the cloud," in Science Gateways (IWSG), 2014 6th International Workshop on, June 2014, pp. 42–47.

[9] A. Inc., "Abiquo anycloud," 2014, URL: http://www.abiquo.com/ [retrieved: 07, 2014].

[10] A. S. Foundation, "Cloudstack," 2013, URL: http://cloudstack.apache.org/ [retrieved: 07, 2014].

[11] I. Eucalyptus Systems, "Eucalyptus," 2014, URL: https://www.eucalyptus.com/eucalyptus-cloud/iaas [retrieved: 07, 2014].

[12] U. of Chicago, "Nimbus," 2014, URL: http://www.nimbusproject.org/ [retrieved: 07, 2014].

[13] openQRM Enterprises, "openqrm," 2014, URL: http://www.openqrm-enterprise.com/ [retrieved: 07, 2014].

[14] R. C. Computing., "Openstack," 2014, URL: https://www.openstack.org/ [retrieved: 07, 2014].

[15] O. Project, "Opennebula project," 2002, URL: http://opennebula.org/ [retrieved: 07, 2014].

[16] Apache, "Vcl," 2012, URL: https://vcl.apache.org [retrieved: 07, 2014].

[17] L. Hewlett-Packard Development Company, "Hp helion cloudsystem," 2015, URL: http://www8.hp.com/us/en/cloud/cloudsystem.html [retrieved: 01, 2015].

[18] D. Nurmi, R. Wolski, C. Grzegorczyk, G. Obertelli, S. Soman, L. Youseff, and D. Zagorodnov, "The eucalyptus open-source cloud-computing system," in Cluster Computing and the Grid, 2009. CCGRID '09. 9th IEEE/ACM International Symposium on, May 2009, pp. 124–131.

[19] D. Dongre, G. Sharma, M. Kurhekar, U. Deshpande, R. Keskar, and M. Radke, "Scalable cloud deployment on commodity hardware using openstack," in Advanced Computing, Networking and Informatics-Volume 2, ser. Smart Innovation, Systems and Technologies, M. Kumar Kundu, D. P. Mohapatra, A. Konar, and A. Chakraborty, Eds., vol. 28. Springer International Publishing, 2014, pp. 415–424.

[20] J. Vilaplana, F. Solsona, J. Mateo, and I. Teixido, "Sla-aware load balancing in a web-based cloud system over openstack," in Service-Oriented Computing – ICSOC 2013 Workshops, ser. Lecture Notes in Computer Science, A. Lomuscio, S. Nepal, F. Patrizi, B. Benatallah, and I. Brandić, Eds., vol. 8377. Springer International Publishing, 2014, pp. 281–293.

[21] S. Saibharath and G. Geethakumari, "Design and implementation of a forensic framework for cloud in openstack cloud platform," in Advances in Computing, Communications and Informatics (ICACCI, 2014 International Conference on, Sept 2014, pp. 645–650.

[22] T. Cordeiro, D. Damalio, N. Pereira, P. Endo, A. Palhares, G. Gonçalves, D. Sadok, J. Kelner, B. Melander, V. Souza, and J.-E. Mångs, "Open source cloud computing platforms," in Grid and Cooperative Computing (GCC), 2010 9th International Conference on, Nov 2010, pp. 366–371.

[23] I. Voras, B. Mihaljevic, and M. Orlic, "Criteria for evaluation of open source cloud computing solutions," in Information Technology Interfaces (ITI), Proceedings of the ITI 2011 33rd International Conference on, June 2011, pp. 137–142.

[24] I. Voras, B. Mihaljevic, M. Orlic, M. Pletikosa, M. Zagar, T. Pavic, K. Zimmer, I. Cavrak, V. Paunovic, I. Bosnic, and S. Tomic, "Evaluating open-source cloud computing solutions," in MIPRO, 2011 Proceedings of the 34th International Convention, May 2011, pp. 209–214.

[25] S. Wind, "Open source cloud computing management platforms: Introduction, comparison, and recommendations for implementation," in Open Systems (ICOS), 2011 IEEE Conference on, Sept 2011, pp. 175–179.

[26] D. Cerbelaud, S. Garg, and J. Huylebroeck, "Opening the clouds: Qualitative overview of the state-of-the-art open source vm-based cloud management platforms," in Proceedings of the 10th ACM/IFIP/USENIX International Conference on Middleware, ser. Middleware '09. New York, NY, USA: Springer-Verlag New York, Inc., 2009, pp. 22:1–22:8.

[27] N. Khan, A. Noraziah, T. Herawan, and M. Deris, "Cloud computing: Analysis of various services," in Information Computing and Applications, ser. Lecture Notes in Computer Science, B. Liu, M. Ma, and J. Chang, Eds., vol. 7473. Springer Berlin Heidelberg, 2012, pp. 397–404.

[28] A. Li, X. Yang, S. Kandula, and M. Zhang, "Cloudcmp: Comparing public cloud providers," in Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement, ser. IMC '10. New York, NY, USA: ACM, 2010, pp. 1–14.

[29] U. de Costa Rica, "Oplau," 2014, URL: http://oplau.ucr.ac.cr/ [retrieved: 07, 2014].

[30] A. Kuo, "Opportunities and challenges of cloud computing to improve health care services," J Med Internet Res, vol. 13, no. 3, Sep 2011, p. e67.

[31] CINDE, "Services sector," 2015, URL: http://www.cinde.org/en/sectors/services [retrieved: 01, 2015].

[32] S. Averitt, M. Bugaev, A. Peeler, H. Shaffer, E. Sills, S. Stein, J. Thompson, and M. Vouk, "Virtual computing laboratory (vcl)," in Proceedings of the International Conference on Virtual Computing Initiative. IBM, 2007, pp. 1–16.

[33] M. Rouse, "Bare-metal provisioning," 2014, URL: http://searchdatacenter.techtarget.com/definition/bare-metal-provisioning [retrieved: 10, 2014].

[34] R. T. Fielding and R. N. Taylor, "Principled design of the modern web architecture," in Proceedings of the 22Nd International Conference on Software Engineering, ser. ICSE '00. New York, NY, USA: ACM, 2000, pp. 407–416. [Online]. Available: http://doi.acm.org/10.1145/337180.337228

[35] U. de Costa Rica, "Centro de informática, normativas," 2012, URL: http://ci.ucr.ac.cr/node/105 [retrieved: 07, 2014].

[36] C. de Investigaciones en Tecnologías de la Información y Comunicación, "Cloud platforms comparison: a technical report," 2014, URL: http://citic.ucr.ac.cr/cloud-platform-comparison-technical-report/ [retrieved: 10, 2014].