

CHOOSING CLOUD COMPUTING SOFTWARE FOR AN ONLINE VIRTUAL LABORATORY SUPPORTING OPERATING SYSTEMS COURSES

¹T.SUSHMA,²PRASANNA POGAKU,³THALLA DAYAKAR

^{1,2,3}Assistant Professor

Department of CSE

Vaagdevi Engineering College, Bollikunta, Khila Warangal, Warangal, Telangana, India, 506005

ABSTRACT

In order to complement the Operating Systems course, the article offers a survey on cloud platforms appropriate for a virtual online laboratory using Linux online environments. In contrast to using pre-made IaaS (Infrastructure as a Service) cloud services from providers, the study supports the decision to create specifically tailored cloud environments tailored for educational needs. It also justifies the choice to use private cloud as a deployment model and IaaS as a service model. An analysis is conducted of the relevant works on cloud platforms for operating system education. The paper also reviews the writers' earlier studies on virtualized environments and technologies for the Operating Systems and Cisco CyberSecurity Operations courses. The fundamental and supplementary prerequisites for cloud computing software for the online virtual laboratory that supports the Operating Systems course have been clarified. Lastly, the work compares the cloud computing platforms Eucalyptus, OpenStack, CloudStack, and OpenNebula and supports the choice of the first and second choice among these cloud computing software platforms.

Keywords: Linux, private cloud, IaaS, operating systems, and virtual online laboratory

1. INTRODUCTION

Most of the Operating Systems (OS) courses include practical assignments on real OSs. In many cases, these assignments require giving students administrative access to their isolated instance of the OS. Also, OS courses usually consider Linux OS, although the majority of computers in university labs, as well as students' laptops more often come with Windows installation. For this reason, different virtualization technologies may be used.

Virtual online laboratories became popular for teaching various courses [1] including courses in OS. However, OS course belongs to the courses with needs that are considerably harder to meet than the needs of most of the courses, which usually could be taught using cloud services of Software as a Service (SaaS) model.

The purpose of the article is to make a survey on cloud platforms applicable for a virtual online laboratory containing Linux online environments in the Operating Systems course, to compare these platforms and to select the most suitable platforms.

2. RELATED WORK

The available works on cloud platforms for teaching operating systems describe the experience of using a wide range of tools, including cloud services given by provider and specially tailored cloud environments adapted for educational needs.

Rajaei and Aldakheel [2] are convinced that the Operating Systems course is among the courses that benefit from using cloud-based environments the most. The authors suggest applying Amazon AWS instances to give students the ability to learn how different scheduling algorithms behave, how virtual memory is managed etc. Gaffar and Hajjdiab [3] describe the experience of working with processes, threads, pipes and sockets using cloud-based laboratories (Ubuntu Linux instances on Amazon AWS). Bhatia et al. [4] focuses on the design of cloud for higher education institution, its proof-of-concept implementation methodology (private cloud, based on OpenStack platform). The authors also formulate the resource requirement model, which estimates the amount of resources needed for a specific number of virtual machines (VMs).

Malan [5] describes the experience of moving Harvard College's introductory

computer science course (CS50) into Amazon Elastic Compute Cloud (Amazon EC2). Markova et al. [6] shows and analyses the ways of implementation of cloud services, based on SaaS, PaaS (Platform as a Service) and IaaS models, for teaching courses in mathematics, natural science and information technologies. The authors also define the most important advantages of those implementations including the possibility of using modern parallel programming tools as the foundation of cloud technologies.

3. THE SURVEY OF CLOUD COMPUTING SOFTWARE FOR VIRTUAL ONLINE LABORATORY SUPPORTING THE OPERATING SYSTEMS COURSE

3.1. Previous research

While teaching different university courses in operating systems and computer networking we have accumulated the experience of using diverse virtualization tools. More details on the Operating Systems course experience are given in the study [7]. The experience of teaching Cisco CyberSecurity Operations course is presented in the work [8].

The study [8] describes using a virtual cloud laboratory based on Apache CloudStack and EVE-NG Community for teaching the CCNA Cyber Operations course at Cisco Networking Academy. The work shows the design and implementation of the virtual cloud laboratory.

The laboratory provides the opportunity to create the sufficient number of virtual machines, to change the computing power, to simulate real computer systems, networks and network topology, to save the virtual machines state between reboots, to combine students' subnets into one virtual network, to work remotely through a virtual private network, as well as to support students' learning and to control their learning outcomes.

As regards courses in operating systems, we have been utilizing different tools over the years, including Oracle VirtualBox hypervisor, Amazon EC2 t2.micro instances, Cloud 9 IDE, virtual machine from NDG Linux Essentials [9] course and Webminal online environment as well as bare-metal Linux installation without virtualization on some of the students' laptops. To meet the

needs of the course and individual students' needs several virtualization tools could be combined according to the methodology for using Unix-like OS virtualization technologies in training bachelors of Informatics and the varied approach to applying these technologies [10, 11].

In 2021 we are supporting a two-semester Operating Systems course for students of Computer Engineering, Cybersecurity, Software Engineering, Computer Science specializations at Zhytomyr Polytechnic State University (Ukraine). We currently use VirtualBox and a Docker-based virtual machine from NDG Linux Essentials online course with intermittently utilizing Amazon EC2 and various free online environments in cases when the main virtualization tools are not available.

In the work [7] we analyzed Linux ready-made online virtual environments and made the comparison of standalone online Unix/Linux terminals in terms of their use for the Operating Systems course. As a result, list of major features of Linux online environment include the support of most Linux command (including administrative ones), administrative privileges for students inside virtual environments, basic networking operations support, the ability to upload/download files to/from virtual environments (and/or saving virtual environment state between reboots), guest OS updating and bash-scripting support.

The requirements for Linux online virtual environments for teaching the operating systems course have been also formulated and include realism, relevancy, availability, stability, scalability and security [7].

We anticipated facing several challenges when working on Linux online virtual environments designing and implementation, namely the challenge of a wide choice, the challenge of implementing, the challenge of network isolation, the challenge of a "playground" and the challenge of a transition [7].

The next stage of the research is to investigate virtualization technologies suitable for Linux online environments, including private open-source cloud platforms, and to select applicable tools. At this stage of the study, we are dealing with the challenge of a wide choice. Different virtualization tools and platforms are

available, so selecting the best suitable set of tools is a complex task

3.2. Elaborating the requirements for cloud computing software for virtual online laboratory in the Operating Systems course

Previously analyzed ready-made Linux online environments included standalone online Unix/Linux terminals (JSLinux, Copy.sh, Virtual x86, Webminal, Linuxzoo, JS/UIX, Weblinux, Browsix, CB.VU), online IDEs which include Unix/Linux terminals (Codio, Cloud 9, Codeany where, Paiza.io) and full-function Unix/Linux virtual environments in the cloud (Amazon EC2, Google Cloud Platform, Microsoft Azure) [7].

Later at least one of the mentioned standalone online Unix/Linux terminals (Webminal) changed the terms of its availability, so we had to move to other environments as alternative virtualization tools. This case illustrates the need to select more stable solutions for teaching OS course. Besides, free utilization of standalone online terminals has serious functional limitation (i.e., sudo access, networking).

Online IDEs with Unix/Linux terminals are better suited for programming courses while we need the environment for practising administrating commands.

Full-function Unix/Linux virtual environments in the cloud are given by cloud service providers according to Infrastructure as a Service (IaaS). They are ready-made and need comparatively small resources for setting up and further administrating. These services also give sudo access in guest OSs, keep virtual environments state between reboots, support virtual networking and guest OSs updates.

But the usage of these services is usually non-free, need students' credit card number for registration and have other limitation.

The other disadvantage of all cloud services given by the provider is that the terms of use may be changed for both free and paid services. Changes could affect the price, availability, maximum amount of allocated resources. In the case of Cloud9 IDE the service was initially offering a free Ubuntu VM, a full-function terminal with sudo access and an IDE for collaborative programming for all the registered users.

But later because of numerous security violations from these VMs, the provider changed the terms. After the changes, users are able to connect VMs from other services only.

So, individually designed virtual Linux environments for the Operating Systems course should provide the students with opportunities close to those given by providers of IaaS cloud services, but the solution must be more provider independent. Although we still appreciate compatibility between analyzed cloud computing software and providers' cloud service platforms.

Eligible virtualization software includes cloud computing software as well as environments that are not necessarily intended to run in the cloud. An example of the latter is Proxmox VE, the application of which for teaching the OS administrating is described in the study [12].

However, in the work [7] we marked the need for scalability. With a scalable system, it is possible to start with a comparatively simple solution and gradually enhance it, adding new features, larger hardware capacities and meeting the needs of the course, which change over time. Although some degree of scalability can be reached with a wide range of tools, cloud platforms were initially designed with scalability in mind. Therefore, in further research, we will concentrate on cloud platforms.

According to the NIST Definition of Cloud Computing [13] the major *cloud deployment models* are the following:

- *private cloud* (exclusively used by a single organization, but not necessarily owned by this organization);
- *community cloud* (exclusively used by a specific community of customers, who belong to organizations that have shared concerns);
- *public cloud* (available for the general public, although its using is not necessarily free of charge);
- *hybrid cloud* (a composition of two or more different cloud infrastructures).

As a virtual online laboratory for the Operating Systems course would be used exclusively by the university, the most suitable cloud deployment model is a private cloud.

Also, understanding that studying product description, documentation analysis and test

installation may not show all the important aspects of the product usage in each particular case, we decided to select two pieces of cloud computing software, which would be the platform of the first choice and the platform of the second choice. The platform of the second choice may be used if any serious and unexpected obstacles will be discovered with the platform of the first choice.

Taking into account all the above-mentioned factors the requirements for cloud computing software for virtual online laboratory in the Operating Systems course should be as follows.

The basic requirements for cloud computing software for virtual online laboratory in the Operating Systems course:

- private cloud deployment model support;
- Linux guest support;
- distribution under free and open-source software licenses;
- virtual networking support;
- web interface for students with ability to create, administrate, delete virtual instances, and also to establish and configure network connections between these instances;
- detailed documentation;
- integration with authentication protocols (LDAP).

The additional requirements for cloud computing software for virtual online laboratory in the Operating Systems course:

- simple basic installation for beginners;
- Windows guests support.

3.3. The comparison of free and open-source cloud computing software for virtual online laboratory in the Operating Systems course

We analyzed four popular cloud computing software available under free and open-source licenses – Eucalyptus, OpenStack, CloudStack and OpenNebula. The major characteristics of Eucalyptus, OpenStack, CloudStack and OpenNebula important to this study are given in table 1.

Eucalyptus. Eucalyptus [14] is an open-source IaaS cloud computing software for building private and hybrid clouds. The product is developed by Eucalyptus Systems and distributed under GNU GPL v3 license with community support and an option of paid support. A distinctive feature of Eucalyptus is its compatibility with Amazon

AWS. Eucalyptus uses similar instances types, tools, virtualization technologies, terminology and supports AWS APIs. Eucalyptus also has a FastStart solution for beginners. FastStart is intended to be run from CentOS 7.9 minimal installation and requires few unused IP addresses in the subnet. Eucalyptus supports easy download of ready-made Linux images (CentOS, CentOS Atomic Host, Fedora, Fedora Core OS, Ubuntu) from a command line.

Table 1
The comparison of free and open-source cloud computing software for virtual online laboratory in the Operating Systems course.

Characteristics	Eucalyptus	OpenStack	CloudStack	OpenNebula
Cloud service model	IaaS	IaaS	IaaS	IaaS
License	GNU GPL v3	Apache License 2.0	Apache License 2.0	Apache License 2.0
Cloud deployment model	private and hybrid cloud	private, public and hybrid cloud	private, public and hybrid cloud	private, public and hybrid cloud
Based on virtualization technology	virtual machines (KVM, Xen, VMware hypervisors)	bare-metal virtualization, virtual machines, virtual containers (various virtualization platforms)	virtual machines (KVM, VMware vSphere, XenServer/XCP, VirtualBox), virtual containers (LXC)	virtual machines (VMware vCenter, KVM, VirtualBox), virtual containers (LXD), microVMs (AWS Firecracker)
Host OS	Cent OS 7.9, Red Hat Enterprise Linux 7.9 (both 64-bit)	Windows, Solaris, Linux OSs and VMware ESX hypervisor	Linux, Windows	Linux (including Cent OS and Ubuntu)
Guest OS	all guests supported by the basic virtualization software	all guests supported by the basic virtualization software	Linux, Windows	all guests supported by the basic virtualization software
Name of guest OS environment	instances, virtual machines	instances, virtual machines, virtual containers (depending on basic virtualization software)	instances, virtual machines, virtual containers (depending on basic virtualization software)	instances, virtual machines, virtual containers, microVMs (depending on basic virtualization software)
Interface for administrator	web-based console and administration CLI	web-based UI (Horizon)	web-based UI (CloudStack UI)	Sunstone GUI, FinEdge GUI, CLI
Interface for guest OS access	web-based console, Euca2ools, AWS CLI	Horizon, OpenStackClient	web-based UI (CloudStack UI)	Sunstone GUI, CLI
Virtual networking	AWS VPC	neutron	AWS-style networking, Nicira NVP	Linux bridge networks, iBMC IQ networks, VXLAN networks, OpenSwitch networks
Integration with LDAP	yes	yes	yes	yes

Eucalyptus contains an administration CLI, a web-based console for instances management (figure 1), and supports command line tools for connection to instances, which currently are Euca2ools (Cent OS, RHEL, Mac OS X hosts) [15] and AWS CLI (Linux, Windows, macOS hosts and Docker) [16].

System requirements for Eucalyptus include Intel / AMD CPU with 4 cores 2 GHz each, minimum 16 RAM for each virtual machine, minimum 100 Gb of storage on each host, but recommends 200 Gb on Node Controller host running Linux virtual machines, 500 Gb on Storage Controller hosts and 500 Gb on Walrus object storage gateway as well (if used). Machines that host components (except Cloud Controller and Node Controller) must support UDP multicast



Depending on the scenario, OpenStack minimally requires 1 physical host with 16 GB RAM, multi-core CPU and 50 GB of free disk space (for single-node deployment); 2 physical hosts each with 16 GB RAM, multi-core CPUs and 50 Gb of free disk space (for multi-node deployment);



for IP address 239.193.7.3. All services must be installed on physical machines, not virtual [14].

OpenStack. OpenStack is a free and open-source cloud computing platform with a focus on open standards. OpenStack was founded and initially developed by Rackspace Hosting and NASA (the influence of the latter is particularly notable in the names of components). Currently, the platform is managed by OpenStack Foundation, which is a non-profit organization. The developers of OpenStack call it a cloud operating system [18]. The platform is compliant with many tools, operating systems and technologies. OpenStack uses Apache License 2.0. OpenStack comes with Horizon Web UI (figure 2), OpenStack SDK, and also supports third party tools like Kubernetes, CloudFoundry, Terraform.

Figure 2: OpenStack dashboard [19].

6 physical servers each with 8 Gb RAM, IPMI BMCs, dual NICs, support of HA architectures, network switch and Internet gateway (for data centre cluster deployment) [22].

It works with a wide range of hosts, including Windows, Solaris, Linux OSs and Vmware ESXi hypervisor. Supported Linux hosts include Cent OS, Debian, Fedora, openSUSE, RHEL, Ubuntu and other distributions [20].

It also should be noted that minimal deployment for the Wallaby series (released in June 2021) consists of Keystone identity service, Glance image service, Placement service, Nova compute service, Neutron networking service and will be most likely also include Horizon dashboard and Cinder block storage service [23].

For new users the Training Labs scripts are available, providing an automated deployment of the cluster using VirtualBox or KVM VMs, hosted on a desktop PC or a laptop with Linux, MacOS or Windows OS, minimum 8 Gb RAM and 20 GB of free storage [21].

CloudStack.

CloudStack [24] has been initially developed by Cloud.com (purchased by Citrix). Further development is currently being implemented by the Apache Foundation.

System requirements for OpenStack vary greatly because of the high flexibility of the platform and depend on the objective of building the cloud, the components selected and other factors.

CloudStack is an open-source IaaS cloud software designed for large networks, but applicable to smaller networks as well. It is distributed under Apache License 2.0.

For example, Ubuntu OpenStack installation could be deployed according to three main scenarios: single-node deployment, multi-node deployment and data-centre cluster deployment.

CloudStack supports Amazon AWS APIs and Open Cloud Computing Interface [25]. A web-based UI (CloudStack UI) is available (figure 3). Minimal system requirements for CloudStack management server, database and storage include 64-bit x86 CPU, 4 GB RAM, 250 GB of free storage (500 GB recommended), 1 NIC, static IP address and fully qualified domain name at each physical host. The management server may be installed on a virtual machine. Hosts, where hypervisors and virtual machines would run, must support Intel-VT or AMD-V hardware-assisted virtualization, have 64-bit x86 CPU,

4 GB RAM, 36 GB of free storage, 1 NIC [26].

OpenNebula.

OpenNebula [27] is an open-source IaaS platform to build and manage enterprise clouds. It is developed by OpenNebula Systems and OpenNebula community, distributed under Apache License version 2. OpenNebula is supported by its community, although a paid subscription for enterprise usage is also available (OpenNebula EE).



Figure 3: Apache CloudStack dashboard [25].

OpenNebula supports integration with Amazon AWS and Microsoft Azure. The product contains built-in UIs (Sunstone GUI, FireEdge GUI) and also works with third-party tools, including Terraform, Kubernetes, Ansible and Docker [28]. The Sunstone GUI is shown on figure 4.

For new users, the miniONE tool is provided. This tool helps to deploy OpenNebula cloud based on KVM virtual machines, it also installs and configures all necessary components to manage and run the virtual machines. System requirements for OpenNebula miniONE include a physical or virtual server with 4 GiB RAM, 20 GiB of free storage and open ports for SSH (22), Sunstone (80) and FireEDGE (2616) [29].

All the analyzed platforms have detailed online documentation with manuals and step-by-step instructions is available.

Bedi et al. [30] after giving a comparison of Eucalyptus, OpenStack, CloudStack and Open Nebula cloud platforms arrived at the conclusion that despite all the listed systems could be successfully used for cloud deployment, OpenStack and Eucalyptus are more suitable for infrastructure provisioning, while CloudStack and OpenNebula would show better results on datacentre virtualization.

Taking into account all mentioned above, we have come to a conclusion to select the Open Stack cloud platform as the platform

of the first choice because of its high flexibility and conformance with open standards. OpenStack also provides Training Labs scripts, which would be particularly useful at the early stage of the virtual laboratory implementation and may be also applied for students' extra-curricular and scientific work.

The other three platforms (Eucalyptus, CloudStack and OpenNebula) also look promising and meet most of the requirements. These platforms have been selected as platforms of second choice.



Figure 4: OpenNebula dashboard [27]

4. CONCLUSIONS

Because of the effort, the survey on cloud platforms may now be used to support the Operating Systems course in a virtual online laboratory using Linux online environments. The experience of employing tools based on various technologies, such as cloud services provided by providers (Amazon AWS) and specifically designed cloud environments adapted for educational objectives (OpenStack, CloudStack), is described in related works on cloud platforms for teaching operating systems.

The report provides a succinct overview of earlier studies on virtualization environments and technologies. The fundamental specifications for the cloud computing software for the online virtual laboratory that supports the Operating Systems course have been refined. These include support for Linux guests, private cloud deployment models, distribution under free and open-source software licences, support for virtual networking, a web interface that gives students the ability to create, manage, and delete virtual instances as well as set up and configure network connections between them, thorough documentation, and integration with authentication protocols (LDAP). The

prerequisites also include Windows guest support and a basic installation for novices.

Eucalyptus, OpenStack, CloudStack, and OpenNebula are the four well-known cloud computing programmes that are compared in this article and are all licenced under free and open-source licences.

The study came to the conclusion that the OpenStack cloud platform should be utilised as the first option due to its flexibility, support for open standards, and availability of Training Labs scripts for students to use for after-school activities and experimentation. Thinking about how Eucalyptus

REFERENCES

- [1] A. Striuk, O. Rybalchenko, S. Bilashenko, Development and using of a virtual laboratory to study the graph algorithms for bachelors of software engineering, CEUR Workshop Proceedings 2732 (2020) 974–983.
- [2] H. Rajaei, E. A. Aldakheel, Cloud computing in computer science and engineering education, in: Proceedings of 119th ASEE Annual Conference & Exposition, ASEE Conferences, San Antonio, Texas, 2012, pp. 422–431. doi:10.18260/1-2--21076.
- [3] M. A. Gaffar, H. Hajjdiab, Teaching operating systems concepts using the cloud, World Transactions on Engineering and Technology Education 16 (2018) 156–161. URL: [http://www.wiete.com.au/journals/WTE&TE/Pages/Vol.16,%20No.2%20\(2018\)/10-Hajjdiab-H.pdf](http://www.wiete.com.au/journals/WTE&TE/Pages/Vol.16,%20No.2%20(2018)/10-Hajjdiab-H.pdf).
- [4] G. Bhatia, I. A. Noutaki, S. A. Ruzeiqi, J. A. Maskari, Design and implementation of private cloud for higher education using openstack, in: 2018 Majan International Conference (MIC), 2018, pp. 1–6. doi:10.1109/MINTC.2018.8363161.
- [5] D. J. Malan, Moving CS50 into the Cloud, in: 15th Annual Conference of the Northeast Region of the Consortium for Computing Sciences in Colleges, Hartford, Connecticut, 2010. URL: <https://cs.harvard.edu/malan/publications/ccsne10-malan.pdf>.
- [6] O. M. Markova, S. O. Semerikov, A. M. Striuk, H. M. Shalatska, P. P. Nechypurenko, V. V. Tron, Implementation of cloud service models in training of future information technology specialists, CEUR Workshop

Proceedings 2433 (2018) 499–515. URL: <http://ceur-ws.org/Vol-2433/paper34.pdf>.

[7] O. Holovnia, Linux online virtual environments in teaching operating systems, CEUR Workshop Proceedings 2732 (2020) 964–973. URL: <http://ceur-ws.org/Vol-2732/20200964.pdf>.

[8] N. Balyk, Y. Vasylenko, V. Oleksiuk, G. Shmyger, Designing of Virtual Cloud Labs for the Learning Cisco CyberSecurity Operations Course, CEUR Workshop Proceedings 2393 (2019) 960–967. URL: http://ceur-ws.org/Vol-2393/paper_338.pdf.

[9] NDG, Linux Essentials – Online Courses & Labs Training, 2022. URL: <https://www.netdevgroup.com/online/course/s/open-source/linux-essentials>.

[10] O. M. Spirin, O. S. Holovnia, Using Unix-like operating systems virtualization technologies in training the bachelors of Computer Science, Information Technologies and Learning Tools 65 (2018) 201–222. doi:<https://doi.org/10.33407/itlt.v65i3.2055>.

[11] O. S. Holovnia, Methodology for using Unix-like operating systems virtualization technologies in training bachelors of Informatics, Ph.D. thesis, Institute of Information Technologies and Learning Tools of the NAES of Ukraine, Kyiv, Ukraine, 2019.

[12] A. Batiuk, D. Vankevych, G. Zlobin, Using virtualization technologies in courses “System administration OS Linux”, Electronics and information technologies 3 (2013) 220–225.

[13] P. Mell, T. Grance, The NIST Definition of Cloud Computing, Technical Report 800-145, National Institute of Standards and Technology (NIST), Gaithersburg, MD, 2011. URL: <http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf>.

[14] Eucalyptus Documentation, 2020. URL: <https://docs.eucalyptus.cloud/eucalyptus/5/>.

[15] Euca2ools 3.4.1 Reference Guide, 2018. URL: <https://docs.eucalyptus.cloud/eucalyptus/4.4.3/euca2ools-guide-3.4.1.pdf>.

[16] Amazon Web Services, Installing, updating, and uninstalling the AWS CLI version 2 – AWS Command Line Interface, 2022. URL:

<https://docs.aws.amazon.com/cli/latest/userguide/install-cliv2.html>.

[17] Navigate the dashboard, 2018. URL: https://docs.eucalyptus.cloud/eucalyptus/4.4.5/shared/console_dashboard.html.

[18] Open Source Cloud Computing Platform Software – OpenStack, 2021. URL: <https://www.openstack.org/software/>.

[19] Openstack dashboard – openstackdocstheme 2.3.1.dev10 documentation, 2021. URL: https://docs.openstack.org/openstackdocstheme/latest/demo/dashboard_demo.html.