ASEE'S VIRTUAL CONFERENCE
At Home with Engineering Education
JUNE 22 - 26, 2020 #ASEEVC
Paper ID #31629

Developing Information Technology Labs on Google Cloud Platform

Dr. Peng Li, East Carolina University

Developing Information Technology Labs on Google Cloud Platform

1. INTRODUCTION

Information technology is evolving rapidly and instructional material must be adapted accordingly. Cloud computing leverages on-demand, efficient resource sharing of a virtual infrastructure, which allows swift deployment and adaptation of curriculum and laboratory experiences in step with the advances in the field. In addition to the in-house cloud-based (private cloud) systems in our college, we partnered with Google to offer our Information and Computer Technology students additional experience with Google Cloud Platform (GCP), a commercial, cloud-based (public cloud) platform. The cloud-based lab environments are accessible to students anytime and anywhere over the Internet. Our purpose is to capitalize on cloud computing technologies to enhance the learning experiences for both on-campus, face-to-face students and online, distance education students.

Four new laboratory modules are developed. Content such as cloud security is added to reflect the new development in the information technology field. Students can conduct up-to-date lab exercises in virtual environments hosted remotely on the cloud. This approach is different from the traditional remote lab approach in that the whole lab environment is virtualized and hosted on a commercial public cloud (GCP) with better availability, scalability and usability [1]. A student survey was conducted anonymously to evaluate the hands-on labs and lab environments. The overall feedback was positive. The whole experience may be useful for other institutions that are interested in adopting commercial, public cloud platforms in laboratory education.

2. MOTIVATION

Our Information and Computer Technology program focuses on learning by doing. Hands-on laboratory is always a vital part of the curriculum. Laboratory courses serve not only on-campus, face-to-face (F2F) students, but also online, distance education (DE) students. Student enrollment, especially in the DE section, has been increasing since early 2000s, while resources have been constrained due to state budget issues. There was an urgent need to find ways to deliver hands-on laboratory to meet the growing demand. To address the need, we deployed virtual labs gradually, utilizing virtualization and cloud computing technologies as they became available.

We started from a decentralized, individual approach in 2006, later moved to centralized, in-house, cloud-based (private cloud) solution around 2008, and then adopted commercial cloud-based (public cloud) solution such as Microsoft Azure in some courses in 2015. In a virtual lab, a virtual environment, containing one or more virtual machines, is used by students to perform hands-on labs on a computer. Deployment of virtual labs allows us to provide hands-on learning experiences more efficiently to both on-campus, face-to-face students and online, distance education students [2].

3. RELATED WORK

Virtual machines were adopted in computer science and technology education as early as 2002 [3]. A virtual machine is a software representation of a physical computer, on which real operating systems and applications can run. Virtualization allows a single physical computer to run multiple virtual machines with their guest operating systems. The technology reduces or eliminates the necessity to have multiple computers to host diverse operating systems typically deployed in physical labs. Advances in virtualization and distributed computing have facilitated the development of cloud computing services. Cloud computing is an approach for delivering applications and services, on demand, over a private or public network, by leveraging efficient resource pooling of virtual infrastructure. Using virtualization and cloud computing technologies in education could help lessen hardware costs, increase the availability of resources, and accelerate deployment of new curriculum material [4] - [8]. Some studies indicated that students conducting virtual labs performed as well as students using physical labs in a traditional environment [9], [10].

In a decentralized, individual approach, the virtual environment is hosted on students' personal computers. Examples of software we used include VMware Server, VMware Workstation/Fusion and Oracle VM VirtualBox. Free licenses for VMware software were provided by VMware Academic Program. VirtualBox was free for educational use. The decentralized (distributed) approach made it possible for students to conduct virtual labs anywhere and anytime. Not many resources are required from the college except for occasional needs for technical support. The decentralized approach is suitable for courses in which only a single virtual machine is needed, such as Windows Server Administration or Linux Web Programming. However, when a virtual environment became more complicated and contained multiple virtual machines, not all students' personal computers were powerful enough to run the environment smoothly. In a decentralized approach, typically the instructor creates the virtual environment and provides support. Not all instructors have time or skills to do it.

Starting 2008, we experimented with centralized, in-house, cloud-based (private cloud) systems to host virtual environments in some courses. A private cloud system provisions shared resources (virtual environments) to users on demand inside an organization over the intranet. The system is managed internally. The students log in to virtual environments remotely to do hands-on labs. The in-house, cloud-based system can host complicated virtual environments with multiple virtual machines. The centralized, private cloud approach was different from the traditional, remote lab approach in that labs are performed on virtual machines and resources can be used more efficiently via scheduling software. Unlike physical computers, virtual machines can be installed, modified, stored, moved, deployed and destroyed easily. The instructors can have direct control over virtual machines and can update content quickly. On some private cloud systems, instructors can monitor users' consoles directly and help users when needed. For online classes, real-time assistance is a useful feature which is not available in our early, decentralized approach. Examples of in-house, cloud-based (private cloud) systems we used include Apache/NCSU Virtual Computing Lab (VCL) [11], VMware Lab Manager, VMware vCloud Director and NDG NETLAB+. They are installed and maintained by the University Information Technology and Computing Services or the College Technical Support Team. Some private cloud systems, e.g. VCL, are open source and free. Others, e.g. VMware vCloud Director and NDG NETLAB+, have software licensing costs. For all private cloud systems, there are costs for hardware (servers, storage, networking) and technical support. The in-house, private cloud approach did require constant software/hardware upgrades over the past few years due to increasing demand. The private cloud approach is still more cost-effective than the traditional physical lab approach in which only physical equipment and computers are used. Without the private cloud, we will not be able to offer many hands-on labs as we do now because the cost of the traditional approach will be much higher.

Around 2015, we started exploring a commercial, cloud-based (public cloud) platform - Microsoft Azure. Starting 2016, we used Google Cloud Platform (GCP) in a small number of courses. A public cloud is owned and managed by a third-party service provider such as Amazon, Microsoft and Google that provides IT resources for multiple clients over the internet. Amazon Web Services (AWS) has the largest share in the commercial, public cloud market. Some public cloud vendors provide free credits to educators and students via educational grant programs to use their systems. In 2015, Microsoft Azure University Educator Grant offered the most generous monthly credits. The Azure portal was also simple and easy to use. However, the Azure grant program ended in 2017. The main public cloud platform we are using now is Google Cloud Platform. The differences between major commercial, public cloud platforms (AWS, Azure and GCP) for educational use were explored in a conference paper in 2018 [12]. However, the commercial cloud computing market is evolving so fast that some data in the paper may need to be updated again. For different educational institutions, the selections of commercial, public cloud solutions will likely be different, based on different needs and budgets. It is notable that availability and scalability of commercial cloud-based systems, very important for business continuity, are typically higher than those of in-house, private cloud solutions.

4. LAB DEVELOPMENT

In this paper, we are going to discuss our experiences in developing information technology labs on a commercial, public cloud - Google Cloud Platform (GCP). The monthly credits were provided through Google Cloud Platform Education Grant. Each student was given $50 credit per month. The instructor/teaching assistant received $100 credit per month. In our labs, the approximate cost of running a WordPress VM instance was $24.75/month. Our students were instructed to shut down their VMs after completing the labs. Therefore, the $50 credit was enough for most students. The lab development was carried out on GCP at https://console.cloud.google.com/. The students logged in to the same site to perform labs.

The purpose of the project was to help our Information and Computer Technology students learn basic virtualization and cloud computing practices. The labs were designed as supplemental modules that adhere to course objectives. To meet the learning outcomes, the students were expected to

1) understand basic virtualization and cloud computing practices;
2) identify building blocks of cloud computing systems;
3) understand basic usage of cloud infrastructures;
4) recognize commonly used cloud computing services and security applications.

The approach serves our students well because

1) it provides students with multiple hands-on exercises to learn;
2) it does not require major adjustments to the curriculum;
3) it can be deployed and tuned quickly.

The development procedure contains four steps:

1) The objectives of the labs are defined. For instance, an objective is defined as follows: "After completion of this lab, the students will understand how to adjust firewall rules to secure containers and services."
2) Lab manuals and lab answer sheets are created. Each lab manual contains detailed instructions for students. In the lab answer sheet, there are fill-in questions and discussion questions. In addition, screenshots and log files are required in some labs as evidences that certain tasks have been completed.
3) The new labs are tested on Google Cloud Platform and the lab solutions are developed for instructor use.
4) The labs are evaluated by faculty/students and then revised. Most revisions are minor.

Four new laboratory modules were developed, including:

- Lab 1: Google Cloud Platform – Getting Started
- Lab 2: Instance Groups and Monitoring
- Lab 3: Container Security
- Lab 4: Suricata IDS/IPS

In Lab 1, students practice basic usage of Google Cloud Platform by deploying a web app WordPress. In Lab 2, students use Google Cloud Platform to create an instance group and conduct performance monitoring. In Lab 3, students install and use Docker on Google Cloud Platform and adjust firewall rules to secure containers and services. In Lab 4, students install Suricata on Google Cloud Platform and use Suricata as packet analyzer, intrusion detection system and intrusion prevention system.

These labs are created from scratch and are not adapted from labs we developed previously for the in-house, cloud-based systems. However, the experiences and lessons we learned previously were helpful. If a virtual lab requires only a single virtual machine, it can usually be ported to commercial, public cloud such as GCP easily. If the virtual lab uses multiple virtual machines or specialized virtual environments, it may be difficult to move the lab from one platform to another platform. For example, on our in-house, private cloud system VMware vCloud Director, some labs require nested virtual machines (VM inside VM) with complicated networking configuration. They cannot be easily ported to a commercial, public cloud such as GCP. In the labs we developed, Lab 1 and Lab 2 are closely related to Google Cloud Platform itself. Lab 3 and Lab 4 can be ported to other platforms relatively easily because they do not depend on a particular platform.

The portable lab modules were inserted in two courses ICTN 4700/1 Virtualization Technologies and ICTN 4200/1 Intrusion Detection Technology as bonus labs. An anonymous, voluntary, student lab survey was administered using Qualtrics in ICTN 4200/1, as shown in Table 1. Out of 69 students, 18 students took the survey. The response rate was 26.1%. The student feedback was generally positive. No control groups were used in the study. The following were some comments from students:

“While using Google Cloud for these last two labs, I was surprised by just how many option are available to run in Google Cloud. If the prices are comparable to those of VMware, Google Cloud could be the go to source for running VMs.”

“Having real labs to perform your work on really helps the online student.”

“It blows my mind that Google and AWS are so good that my VM could be sitting in another state but yet still functions as if I were in front of it.”

“I've loved experimenting with Google's Cloud Platform. It was nice to use a real service that people are using daily to deploy VMs.”

Table 1: Student Lab Survey in ICTN 4200/1

Question	**Strongly Disagree**	**Disagree**	**Neither Agree nor Disagree**	**Agree**	**Strongly Agree**	**Total**
The labs facilitated my understanding of course topics.	1	0	0	10	7	18
The lab exercises were well organized.	1	0	1	9	7	18
The cloud lab environments were easy to use.	1	0	2	9	6	18
I was able to study at my own pace.	1	1	0	6	10	18
The hands-on exercises in the cloud lab environment were as effective as those in physical computer labs.	1	1	1	7	8	18

5. CONCLUSION

To meet the growing demand for hands-on labs in our information and computer technology program, we used decentralized, individual virtual environments, centralized in-house cloud-based (private cloud) systems and later centralized, commercial, cloud-based (public cloud) systems such as Google Cloud Platform. Four virtual labs were developed for Google Cloud Platform. The student feedback was generally positive. Each solution has its own advantages and disadvantages. The decentralized approach and the centralized approaches can be complementary. For an information technology course requiring a simple virtual machine, the decentralized solution or the commercial, public cloud solution can be considered. If the instructor needs to have direct control over students' virtual machines or to have specialized, complicated virtual environments, the in-house, private cloud solution may be a better choice. Compared to our in-house, private cloud systems, Google Cloud Platform (GCP) has better availability, scalability and usability. GCP, like some other commercial, public cloud platforms, provides much better availability and scalability than our in-house, private cloud systems. GCP has a 99.99% service level objective (SLO) for covered service with no scheduled downtime. Our in-house, cloud-based systems have scheduled and unscheduled downtime several times a year. The high availability of commercial, public cloud-based systems such as GCP is especially important for business continuity when unexpected disasters occur. GCP is much more scalable with built-in load-balancing and scaling features due to gigantic resources owned by Google. The scalability of our private cloud system is limited by the hardware we have. For example, our VMware vCloud Director is hosted by only three physical servers. According to student feedback, GCP is also more user-friendly than our in-house, cloud-based (private cloud) systems.

However, commercial, public cloud platforms also have shortcomings. First, students create their own virtual machines on the public cloud. Instructors have little control over students' virtual machines and cannot monitor them easily. Second, with limited monthly credits, students cannot run complicated virtual environments. It is uncertain that the grant program will continue to offer free credits in the future. Third, students create their own accounts and therefore user management is a problem.

In the future, we plan to develop more labs on commercial, public cloud systems and use Virtual Private Network (VPN) to connect students' virtual machines with a central server to provide better support and monitoring when needed. We are also considering integrating automatic assessment scripts through the central server on the public cloud to provide immediate feedback, which has been done successfully in some labs on our in-house, cloud-based systems.

REFERENCES

[1] D. Puthal, B. P. S. Sahoo, S. Mishra and S. Swain, "Cloud Computing Features, Issues, and Challenges: A Big Picture," *2015 International Conference on Computational Intelligence and Networks*, Bhubaneshwar, 2015, pp. 116-123.

[2] P. Li, "Centralized and decentralized lab approaches based on different virtualization models," *Journal of Computing Sciences in Colleges*, vol. 26, no. 2, pp. 263-269, 2010.

[3] M. Stockman, "Creating remotely accessible virtual networks on a single pc to teach computer networking and operating systems," *Proceedings of the 4th conference on Information technology curriculum*, pp. 67-71, 2003.

[4] B. Stackpole, "The evolution of a virtualized laboratory environment," *Proceedings of the 9th ACM SIGITE conference on Information technology education*, 2008, pp. 243-248.

[5] A. Gasper, S. Langevin, W. Armitage, R. Sekar, T. Daniels, "The Role of Virtualization in Computing Education," *ACM SIGCSE Bulletin*, vol. 40, no. 1, pp. 131-132, 2008.

[6] Y. A. M. Qasem, R. Abdullah, Y. Y. Jusoh, R. Atan and S. Asadi, "Cloud Computing Adoption in Higher Education Institutions: A Systematic Review," in *IEEE Access*, vol. 7, pp. 63722-63744, 2019.

[7] D. G. Chandra and M. D. Borah, "Cost benefit analysis of cloud computing in education," *2012 International Conference on Computing, Communication and Applications*, Dindigul, Tamilnadu, 2012, pp. 1-6.

[8] J. A. González-Martínez, M. L. Bote-Lorenzo, E. Gómez-Sánchez, R. Cano-Parra, "Cloud computing and education: A state-of-the-art survey," *Computers & Education*, vol. 80, pp. 132-151, Jan. 2015.

[9] B. Stackpole, J. Koppe, T. Haskell, L. Guay, and Y. Pan, "Decentralized virtualization in systems administration education," in *Proceedings of the 9th ACM SIGITE Conference on Information Technology Education*, Cincinnati, OH, USA, 16-18 October, 2008, pp. 249-254.

[10] J. Haag, C. Witte, S. Karsch, H. Vranken, and M. van Eekelen, "Evaluation of Students' Learning Behavior and Success as a Prerequisite for Modernizing Practical on Campus Networking Courses in Higher Education," *Journal of Higher Education*, 4, no. 2, pp. 83-90, 2014.

[11] 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 the Virtual Computing Initiative*, pp. 1-6. 2007.

[12] P. Li, "Public and Private Clouds in Hands-on, Online Learning," *2018 ASEE Mid-Atlantic Section Spring Conference*, Washington, District of Columbia, April, 2018.