2018 ASEE Mid-Atlantic Section Spring Conference: Washington, District of Columbia Apr 6

Public and Private Clouds in Hands-on, Online Learning

Dr. Peng Li, East Carolina University

Public and Private Clouds in Hands-on, Online Learning

Peng Li

East Carolina University, Greenville, NC 27858

Abstract

Adoption of cloud computing is a clear trend in online learning. What types of cloud platforms are available? What are the pros and the cons of different cloud options? We will discuss cloud options for hands-on education, based on our experience with multiple cloud platforms. There are different ways of implementing cloud computing. A public cloud is owned and operated by an external service provider. Computing resources are allocated on a pay-per-usage basis. A private cloud is owned and managed by local institutions. Computing resources are provisioned to internal users on demand. Private clouds are highly customizable and may be adjusted to better meet the needs of local users. However, setting up a private cloud requires expertise and resources. Access to public clouds can be set up quickly but users possess less control. In a diversified online learning environment, multiple cloud platforms may be necessary.

Keywords

Virtual lab, private cloud, public cloud, online learning.

Introduction

Learning by doing is important for engineering and technology education. Cloud computing leverages efficient sharing of on-demand, self-managed, virtual infrastructures, allowing swift deployment and adaptation of curriculum and laboratory experiences in step with the advances in the field. Adoption of cloud computing technologies is a clear trend in online learning.[1,2] Some educators have used cloud technologies for years. Others are still exploring ways of incorporating cloud computing in teaching and learning. There are many different approaches of implementing cloud computing.[2,3] For educators with little experience in cloud computing deployment, selecting a suitable cloud platform to use can be a challenge. What types of cloud platforms are available for educators and students? What are the pros and the cons of different cloud platforms? The goal of this paper is to provide some insights into cloud solutions for hands-on, online learning, based on our experience with multiple public cloud and private cloud platforms.

According to NIST, cloud computing is "a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction".[4] Cloud computing technologies make it possible for students to study and learn from any place and at any time. The cloud approach is different from the traditional remote lab approach in that the whole lab environment is virtualized and hosted on clouds with better accessibility, availability, and usability. There are different ways of implementing cloud computing. Private clouds are highly customizable and may be adjusted to better meet the needs of local users. However, setting up a private clouds requires expertise and

sometimes large amount of resources. Access to public clouds can be set up quickly but users possess less control. In a diversified online learning environment, multiple cloud solutions may be necessary.

In order to meet the growing demand for hands-on labs in online learning, we started exploring virtual labs in 2006. In a typical virtual lab, a virtual environment, containing one or more virtual machines, is created by the instructor for students to use. Students perform hands-on labs using the virtual environment. In the first few years, a decentralized approach was used, in which the virtual environment was installed with VMware Workstation or Oracle VirtualBox on students' personal computers. However, it is challenging for instructors to monitor the labs and provide help in a decentralized setting. Also a personal computer may not be powerful enough to support more complicated virtual environments that contain multiple virtual machines. Starting 2008, we moved to a centralized, cloud-based approach, in which the virtual environment is hosted in a private or public cloud. The students can log in to the cloud remotely to perform hands-on activities at any time and at any place. In the past ten years, we have used multiple private and public cloud platforms.

Private Cloud

A typical private cloud is owned and managed by local institutions. Computing resources are provisioned to internal users on demand. Table 1 shows the private cloud and virtualization solutions we have used, including commercial, non-free one such as VMware vCloud Director (formerly VMware Lab Manager), and free, open source one such as Virtual Computing Lab.

Virtual Computing Lab (VCL) is a free, open source cloud computing platform that dynamically provisions computing resources, such as bare-metal machines and virtual machines, to end users. It was initially developed at North Carolina State University and was released under the Apache license in 2008. A typical VCL infrastructure is composed of a VCL web front-end, a database server, one or more management nodes and a blade center with bare-metal servers. The end user logs in to the VCL web site to reserve computing resources, typically a bare-metal machine or a virtual machine loaded with Windows or Linux operating system and applications.

NETLAB+ is a remote access, scheduling system offered by Network Development Group (NDG). NETLAB+ is unique due to NDG's partnerships with leading industry training programs including Cisco Networking Academy, VMware IT Academy, EMC Academic Alliance and Red Hat Academy, among others. Hands-on labs from these programs are provided via NETLAB+ for students to use. The institution needs to set up the physical equipment to host the virtual environments. Software licensing costs and hardware costs vary, depending on the number of users.

ProxMox Virtual Environment (PVE), like Docker, is a container-based virtualization platform rather than a cloud platform. Containers virtualize the operating system instead of the hardware. Therefore, they have less overhead and are more efficient than hypervisors. The disadvantage of PVE is that it is Linux-based and can only provision Linux containers (virtual private servers) and does not support other guest operating systems. For hands-on labs requiring only a single Linux server for each student, PVE may be a good choice due to its low hardware requirement.

VMware vCloud Director (vCD) is a private cloud platform with a self-service portal which allows users to provision and control their workloads. It provides an easy-to-use interface for instructors/administrators to manage users, roles, workspaces and configurations. The instructors can create and control students' virtual machines. The instructors can also monitor students' activities and provide real-time assistance through the portal. vCD offers some powerful features not yet available on other platforms such as network fencing. The licensing cost for vCD is high. However, education institutions may receive discounted or free licenses through VMware Academic Program.

Table 1: Private Cloud and Virtualization Platform Comparison

	VMware vCloud Director	**NDG NETLAB+**	**Virtual Computing Lab**	**ProxMox Virtual Environment**
Pros	• Market share leader in Virtualization • Easy-to-use UI • Real time monitoring/assistance • Diversified operating systems and features supported • Integration with active directory	• Easy-to-use UI • Hands-on labs ready for Cisco Networking Academy and some other IT academies • Real time monitoring/assistance for users	• Free, open source software • Diversified operating systems supported • Long provision time • User management	• Container-based technology • Low software licensing • Low hardware costs
Cons	• High software licensing costs • High hardware costs • Professional setup and maintainence needed	• High software licensing costs • Relatively high hardware costs • Professional setup and maintainence needed	• High hardware costs • Real time monitoring for users difficult to immplement • Fewer features • Professional setup and maintainence needed	• Only Linux virtual private servers supported • Fewer features • UI not easy to use

We also tested some other private cloud platforms such as OpenStack and OpenNebula but did not deploy them as production systems. Both OpenStack and OpenNebula are capable, free, open source cloud solutions for enterprise use. However, for educators and students, they are not as user-friendly as commercial solutions such as VMware vCloud Director.

Public Cloud

A public cloud is owned and operated by an external service provider. Computing resources are allocated to users on a pay-per-usage basis. Examples of public cloud we used include Amazon Web Services, Microsoft Azure and Google Cloud Platform, as shown in Table 2.

Amazon Web Services (AWS) has been providing public cloud services since 2006 and is the dominating market share leader for over 10 years. It offers a very large, mature set of features and services, sometimes difficult for new users to navigate.

Microsoft Azure is second place in the public cloud market. It is a capable platform with fewer features and services than AWS. Azure is closely integrated with other Microsoft products such as Visual Studio. Therefore, it attracts users who focus on Microsoft technologies. Although Azure is not strong in supporting open-source or third-party tools, hosting Linux virtual machines for hands-on labs is not an issue.

Google Cloud Platform (GCP) is a distant third in cloud market share terms with fewer features and services than AWS. GCP can handle traditional workloads and applications but it especially attracts users of cloud-native applications, which are particularly designed to run in a cloud environment. Also it is claimed that GCP provides better support for open source tools and applications.

All three vendors (AWS, Azure and GCP) offer IaaS (Infrastructure as a Service IaaS), including compute, storage and networking resources as a service[5]. To host virtual, hands-on labs, we used Amazon Elastic Compute Cloud (EC2), Virtual Machines in Azure and Compute Engine in GCP. The basic procedure is similar: 1) The educator applies for education grants and receives voucher codes from the public cloud providers; 2) Students are given instructions to redeem the codes, to set up accounts and to use cloud portals; 3) Students log in to cloud portals, create virtual machines and perform hands-on labs on the virtual machines.

Table 2 shows the availability of education grants from the three vendors. Microsoft used to offer an Azure educator grant but the program was discontinued. Instead, Azure with limited features is now integrated with the Microsoft Imagine program (formerly DreamSpark).

The prices of cloud instances are comparable among the three vendors and have been dropping over the past few years. Here GCP is used as an example to show the typical cost of running a virtual machine on a public cloud. For a standard instance (a CentOS virtual machine with 1 vCPU, 3.75 GB memory and 10 GB virtual hard disk) on GCP, the operation cost per month is approximately $25. For a micro instance (a CentOS virtual machine with 1 shared vCPU, 0.6 GB memory and 10 GB virtual hard disk), the cost per month is about $4. Therefore, the $50/month credit for a student can be used for two non-stop standard instances or about 12 non-stop micro instances. AWS offers much less credit through their education grants. However, for most hands-on labs, there is no need to keep the virtual machines running all the time. It is a good idea to instruct students to stop their virtual machines once the labs are completed to save resources.

The major problem we encountered when using the public cloud is that the educators have less direct control. For example, the private cloud system VMware vCloud Director provides a console to allow educators to create, manage and monitor students' roles, permissions and virtual environments. If a student needs assistance when doing a lab, the instructor can view the student's console and provide help quickly. The experience is similar to the one in a real, physical computer room. However, in a public cloud, the students have full control over their own virtual machines and the instructor usually does not have control over them. As a result, it is very difficult for the instructor to monitor or manage students' activities.

Another consequence is that it is difficult for the instructor to deploy customized virtual environments on a public cloud. On a private cloud system such as VMware vCloud Director, the instructor can create a complex virtual environment containing multiple, diversified virtual machines, save the environment as a template and deploy the template for student use quickly. When adjustments are needed, the template can be changed and re-deployed easily because the instructor has full control over part of the system. On a public cloud, a customized virtual environment with multiple virtual machines is difficult to deploy and modify.

Table 2: Major Commercial Public Cloud Comparison

	Amazon Web Services	**Microsoft Azure**	**Google Cloud Platform**
Pros	• Dominating market share leader over 10 years • Most mature, enterprise-ready • Deepest capabilities for serving large institutions	• Second in market share • Capable and broad platform • Integration with other Mircorsoft products and services	• More open-source-centric and DevOps-centric • Designed for cloud-native clients • Deep discounts and flexible contracts
Cons	• Extensive options requiring expertise to implement • Difficult to use for new users • Complex cost management	• Less enterprise-ready for production applications • Lack of support in some open-source and third-party tools.	• Distant third in market share • Fewer features and services • Less global reach
Compute Services	• Elastic Compute Cloud (EC2) • Container services	• Virtual Machines • Additional services	• Compute Engine • Kubernetes Engine • App Engine
Education Grants Available	Yes. $200/year for member or $75/year for non-member educators. $100/year for member or $40/year for non-member students. Annually renewable.[6]	No longer offered. But Microsoft Azure with limited features is now offered through the Microsoft Imagine program.[7]	Yes. $100/month for educators. $50/month for students per course. Annually renewable.[8]

In addition to the three major commercial public cloud vendors (AWS, Azure and GCP), there are many other public cloud providers/projects. For example, GENI (http://www.geni.net) provides a virtual lab infrastructure for networking and distributed systems research and education. Some hands-on computing and networking exercises are available on GENI for classroom use. Cloudlab (https://cloudlab.us) is an infrastructure for testing and building clouds. GENI and Cloudlab are sponsored by National Science Foundation and are free to use with certain conditions. However, the scope of features and services is relatively limited on them.

Conclusions

In the past ten years, we have tested and used multiple private and public cloud platforms to assist hands-on, online learning. Setting up a local, private cloud system is not an easy task. In addition to the financial costs, a capable, local technical support team is required to install and maintain the system. User experiences are very good with private cloud systems, especially when complicated virtual environments are needed for hands-on labs.[9] Educators have more control on private cloud systems and can deploy and update exercises quickly. The students can receive real-time assistance when there are issues.

The initial cost of using a public cloud is relatively low. However, educators have less control over students' activities because the cloud is owned and managed by an external service provider. It is difficult to provide assistance when needed. To mitigate this problem, we are working on a project to connect students' compute nodes in the public cloud with a central server using IPSec channels. The central server will be used by the instructor to access students' nodes and to deploy scripts to assess hands-on labs automatically[10].

For hands-on labs which require only a single virtual machine, container-based solutions, such as ProxMox VE (OpenVZ) and Docker, should be considered. They are easy to set up and do not require much resources. These types of labs can also be hosted on a public cloud such as GCP or AWS. For hands-on labs using a complex virtual environment with multiple virtual machines, a private cloud may be a better choice. Private cloud and public cloud can be complementary and be used together to improve online learning experience.

References

1 T. Ercan, "Effective use of cloud computing in educational institutions." Procedia-Social and Behavioral Sciences 2, no. 2 (2010): 938-942.

2 S. Marston, Z. Li, S. Bandyopadhyay, J. Zhang, and A. Ghalsasi. "Cloud computing—The business perspective." Decision support systems 51, no. 1 (2011): 176-189.

3 B. Rimal, E. Choi, and I. Lumb. "A taxonomy and survey of cloud computing systems." In INC, IMS and IDC, 2009. NCM'09. Fifth International Joint Conference on, pp. 44-51. IEEE, 2009.

4 P. Mell, and T. Grance. "The NIST definition of cloud computing." (2011). [Online]. Available: https://csrc.nist.gov/publications/detail/sp/800-145/final. [Accessed March 9, 2018]

5 L. Leong, R. Bala, C. Lowery and D. Smith. "Magic quadrant for cloud infrastructure as a service, Worldwide" Gartner (2014). [Online]. Available: https://www.gartner.com/doc/reprints?id=1-2G2O5FC&ct=150519. [Accessed March 9, 2018]

6 AWS Educate. [Online]. Available: https://aws.amazon.com/education/awseducate. [Accessed March 9, 2018]

7 Microsoft Azure in Education. [Online]. Available: http://edudownloads.azureedge.net/msdownloads/Microsoft_Azure_in_Education.pdf. [Accessed March 9, 2018]

8 Google Cloud Platform Education Grants. [Online]. Available: https://cloud.google.com/edu. [Accessed March 9, 2018]

9 P. Li, "Virtual lab approaches for information and computer technology education," In *Online Learning for STEM Subjects: International Examples of Technologies and Pedagogies in Use*, M. Childs and R. Soetanto, Ed. Routledge, 2017, pp. 112-126.

10 P. Li and L. Toderick. "An automatic grading and feedback system for E-learning in information technology education". ASEE Annual Conference and Exposition, Conference Proceedings. 122. ASEE, 2015.

Peng Li

Peng Li received a Ph.D. in Electrical Engineering from University of Connecticut. His professional certifications include CISSP, RHCE and VCP. Dr. Li is currently an Associate Professor at East Carolina University. He teaches undergraduate and graduate courses in programming, computer networks, information security, web services and virtualization technologies. His research interests include virtualization, cloud computing, cybersecurity and integration of information technology in education.