Private Cloud: A Teaching Case for a Multi-campus Systems Administration Course

Iwan Tjhin
Whitireia New Zealand
Wi Neera Drive, Porirua 5022
New Zealand
+64 4 237 3100
iwan.tjhin@whitireia.ac.nz

ABSTRACT
This paper describes a private cloud system that centrally manages and runs virtual machines in protected environments. The approach attempts to address some of the limitations of teaching degree-level systems administration course using VMs running on individual lab desktop computers. The system enables the possibility for every student to learn and simulate operating environments similar to those of real-life works. The result is a high performance and availability system that has flexible and scalable configurations to support management and isolation of VMs environments in a multi-campus setting. It also enables alternative students' progress monitoring and eases submission logistics.

Categories and Subject Descriptors
K.3.1 [**Computers and Education]**: Computer Uses in Education – *computer-managed instruction (CMI), distance learning.*

K.6.4 [**Management of Computing and Information Systems]**: System Management – *centralization/decentralization.*

General Terms
Design, Experimentation.

Keywords
Teaching systems administration; virtual machine; VMware; private cloud; centralised system.

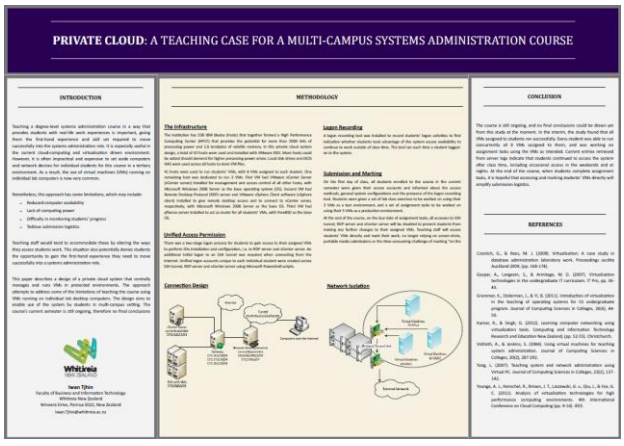

This poster paper appeared at the 4th annual conference of Computing and Information Technology Research and Education New Zealand (CITRENZ2013) incorporating the *26h Annual Conference of the National Advisory Committee on Computing Qualifications*, Hamilton, New Zealand, October 6-9, 2013. Mike Lopez and Michael Verhaart, (Eds).

1. INTRODUCTION
Teaching a degree-level systems administration course in a way that provides students with real-life work experiences is important, giving them the first-hand experience and skill set required to move successfully into the systems administration role. It is especially useful in the current cloud-computing and virtualisation driven environment. However, it is often impractical and expensive to set aside computers and network devices for individual students for this course in a tertiary environment. As a result, the use of virtual machines (VMs) running on individual lab computers is now very common. Nonetheless, this approach has some limitations, which may include reduced computer availability, lack of computing power, difficulty in monitoring students' progress and tedious submission logistics. Teaching staff would tend to accommodate these by altering the ways they assess students work. This situation also potentially denies students the opportunity to gain the first-hand experience they need to move successfully into a systems administration role.

This paper describes a design of a private cloud system that centrally manages and runs VMs in protected environments. The approach attempts to address some of the limitations of teaching the course using VMs running on individual lab desktop computers. The design aims to enable use of the system by students in multi-campus setting. The course's current semester is still ongoing, therefore no final conclusions could be drawn yet for the moment.

2. METHODOLOGY
2.1 The Infrastructure
The institution has 338 IBM blades (hosts) that together formed a High Performance Computing Center (HPCC) that provides the potential for more than 2000 GHz of processing power and 2.6 terabytes of volatile memory. In this private cloud system design, a total of 42 hosts were used and installed with VMware ESXi. More hosts could be added should demand for higher processing power arises. Local disk drives and iSCSI NAS were used across all hosts to store VM files.

41 hosts were used to run students' VMs, with 8 VMs assigned to each student. One remaining host was dedicated to run 3 VMs. First VM had VMware vCenter Server (vCenter server) installed for management and access control of all other hosts, with Microsoft Windows 2008 Server as the base operating system (OS). Second VM had Remote Desktop Protocol (RDP) server and VMware vSphere Client software (vSphere client) installed to give remote desktop access and to connect to vCenter server, respectively, with Microsoft Windows 2008 Server as the base OS. Third VM had pfSense server installed to act as router for all students' VMs, with FreeBSD as the base OS.

2.2 Connection Design

The RDP server was set as the central contact point for all connections between the VMs and students' client computers, accessed through RDP thin-clients. This created a single-point incoming access management configuration that helped to simplify network isolation. Remote sessions from RDP thin-clients to RDP server could be established either directly from the Faculty's local area network or via Faculty's Secure Shell (SSH) tunnel connection over the Internet. Once connected to the RDS server, students could launch vSphere client from within the remote session to connect to vCenter server to gain control over their assigned VMs.

pfSense server was logically positioned in the middle of all other VMs in the system, routing data traffics between hosts, vCenter server, RDP server and all students' VMs. This created a single-point data traffic management. With 1 of its virtual network adapters (vNICs) listening to all Virtual Local Area Network (VLAN) tags (trunk mode), it enabled easier VLAN data traffics routing and DHCP assignment, if needed.

The logical configuration of the private cloud system is shown in Figure 1 below.

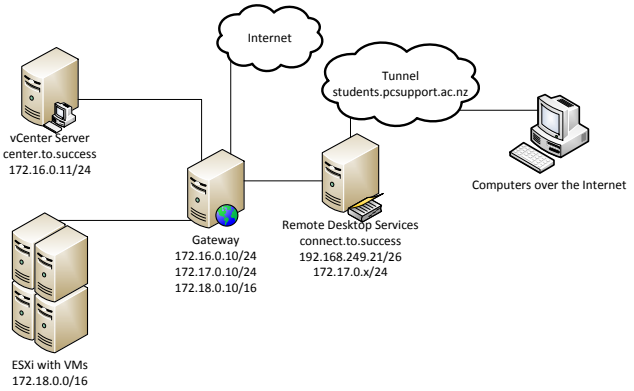

Figure 1. Logical configuration of the private cloud system, showing relationship between different nodes within and outside of the system.

VMs operations, such as client OS installation would need to done directly through this vSphere client. However, access to installed and running client OSs may be done via RDP connection or web browser, whichever supported by the installed client OSs.

2.3 Network Isolation

Each student's set of VMs were grouped to 2 unique VLAN tags, one acted as a test environment using 3 VMs and another was used to simulate a production environment using 5 VMs. This resulted in a large overall operating environment with many VLANs. This setup enabled configurations in such a way that students from one campus could use the system concurrently with students from other campuses without network interference.

2.4 Unified Access Permission

There was a two-stage logon process for students to gain access to their assigned VMs to perform OSs installation and configuration, i.e. to RDP server and vCenter server. An additional initial logon to an SSH tunnel was required when connecting from the Internet. Unified logon accounts unique to each individual student were created across SSH tunnel, RDP server and vCenter server using Microsoft Powershell scripts.

2.5 Logon Recording

A logon recording tool was installed to record students' logon activities to find indication whether students took advantage of the system access availability to continue to work outside of class time. The tool ran each time a student logged-on to the system.

2.6 Submission and Marking

On the first day of class, all students enrolled to the course in the current semester were given their access accounts and informed about the access methods, general system configurations and the presence of the logon recording tool. Students were given a set of lab class exercises to be worked on using their 3 VMs as a test environment, and a set of assignment tasks to be worked on using their 5 VMs as a production environment.

At the end of the course, on the due date of assignment tasks, all accesses to SSH tunnel, RDP server and vCenter server will be disabled to prevent students from making any further changes to their assigned VMs. Teaching staff will access students' VMs directly and mark their work, no longer relying on screen-shots, portable media submissions or the time-consuming challenge of marking "on the spot" in class time.

3. CONCLUSION

The course is still ongoing, and no final conclusions could be drawn yet from this study at the moment. In the interim, the study found that all VMs assigned to students ran successfully. Every student was able to run concurrently all 8 VMs assigned to them, and was working on assignment tasks using the VMs as intended. Current entries retrieved from server logs indicate that students continued to access the system after class time, including occasional access in the weekends and at nights. At the end of the course, when students complete assignment tasks, it is hopeful that accessing and marking students' VMs directly will simplify submission logistics.