Review of existing knowledge

Introduction

This section will identify the existing implementation for the virtual labs and will investigate alternate technologies that can be implemented to improve the current model. This investigation will form the basis of the new proposal.

Current situation

Currently each lab has 20 physical client machines with access to the core university infrastructure via a 100mb LAN link with CAT6 cabling. The current specifications are outlined in Table 1.

Table 1 Hardware and Software setup for current labs

Component		Model	Manufacturer	Description
Lab Machines				
	Operating System	Windows XP 86x	Microsoft	32bit edition
	Processor	Intel Core 2 Duo E8400 2.0GHZ	Intel	Dual Core CPU
	Motherboard	unknown	unknown	With on-board sound
	Graphics Card	NVIDIA Quadro FX580 PCI-EXPRESS 512 MB	NVIDIA Leadtek PCIe	512 MB dedicated memory
	Random Access Memory (RAM)	4GB RAM	HP	Two paired 2GB dims
	DVDRW drive	32x	Stone Computers	
	Network interface card	Broadcom Netlink ™ Gigabit Ethernet	unknown	1000mbps (100 used)
	Hard drive	320 GB HDD	Fujitsu	SATA II
Network Components				
	Switch	100MB	Cisco	Switch connecting machines in lab to university core infrastructure
	Network Cable	CAT5	unknown	
Software (Used on lab Machines)				
	Virtual PC 2007	32bit edition	Microsoft	Installed on 20 lab machines for student labs sessions
Virtual machines				

	• Windows 2003 server professional (86) • Windows 2008 server professional (86) • Windows XP professional (86) • Ubuntu Linux	All 32-bit	Microsoft, Canonical	Installed locally on the 20 machines in the lab; used for student lab sessions

Currently the clients store their own instances of the virtual machines required for all labs, these VMs are impersonal to the user and revert to a default start when shutdown. This is an issue in exam situations where VMs must be left on after the session, which takes up valuable lab time for others, in order to be assessed and marked; this also makes validation of marking near impossible.

The existing infrastructure is outline in the diagram below (Figure 1).

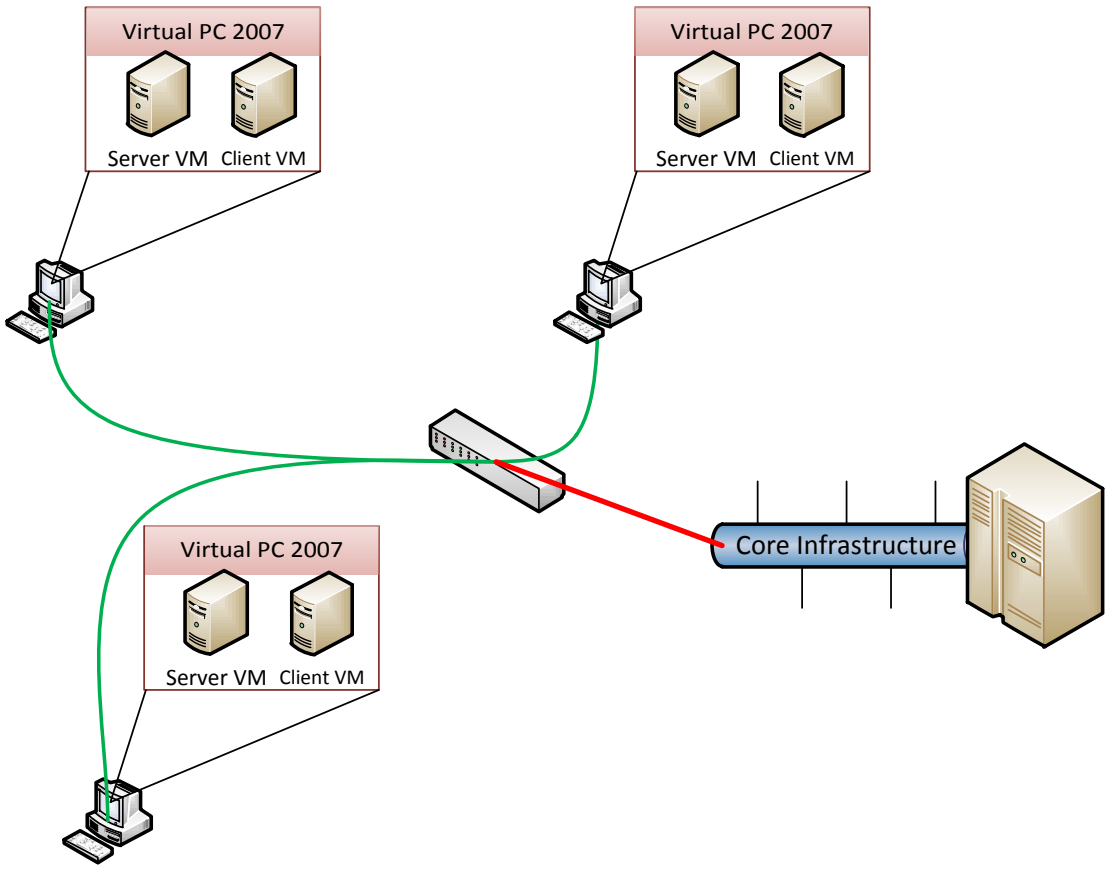

Figure 1 Existing Virtual Infrastructure

Cloud Computing

"Cloud computing is managed, shared applications, development platforms or computing infrastructure via a network" (McDonald, K (2010))

"Cloud" is a term coined in around 2004; however its principles have been around for many years, web based email services such as Hotmail are technically cloud computing as they are a resource found through a network, i.e. the Internet, other business applications such as salesforce.com have been around for a long time too (since 1994), in fact this is one of the first Software as a Service (SaaS) platforms, which is discussed later.

Microsoft's recent "To the Cloud" advertising campaign shows how far cloud computing as a term has come, it is appealing to the home consumer and advertising features of the Windows 7 operating system that utilise elements of the Internet in a cloud like manner

Cloud Service Models

There are 3 core service models of cloud computing, and there are many examples each. Many providers of cloud infrastructures coin their own models too, and Cisco's 4th model as presented on 28th February 2011, in Amsterdam is Data Centre as a Service (DCaaS), this is where the whole data centre is the service, including all the cooling, power and floor space. The 3 primary models are:

- Software as a Service (SaaS) - Software that is utilised and accessed over the cloud, this can be a simple webmail system such as Hotmail or Gmail or any front end application that is part of your business processes such as a CRM or ERP application, that is hosted by a 3rd party organisation.
- Platform as a Service (PaaS) - A development environment for SaaS services for testing, development or use, examples include IBM Test and Development cloud and Amazon Web Services (AWS).
- Infrastructure as a Service (IaaS) - Provides the provisioned physical infrastructure for an organisation, this is where you rent physical servers and utilise them however you want, from wherever you want. This is a service typically offered by larger organisations and examples include Microsoft Azure Cloud and Amazon EC2/S3 infrastructures.

(Adapted from McDonald, K (2010) - *Above the Clouds - Managing Risk in the World of Cloud Computing*)

Cloud Delivery Models

There are 3 models of delivery for cloud services, these are:

Private - Private clouds (aka, on-premises cloud) are cloud deployments inside the organization's premises, managed internally without the benefits of the economy of scale but with advantages in terms of security**.**

Public - Public Clouds are the original concept of cloud computing based on the ubiquity of the internet. This type of cloud provides all the benefits of the economy of scale, ease of management, and ever growing elasticity. The major concern about this style of deployment is security.

Hybrid - Hybrid Clouds are a deployment type that sits between the private and the public clouds. Hybrid Clouds are usually a combination of private clouds and public clouds, usually, managed using the same administration and monitoring consoles.

(Adapted from "Cloud Computing, the new IT Paradigm" http://itechthoughts.wordpress.com/2010/02/23/cloud-computing-the-new-it-paradigm/)

Cloud Decision

Cloud computing is the buzz word of the entire IT industry and is something that should not be ignored, for the purpose of this project, and bearing in mind the customer is the university I think a private cloud model should be used to deliver the virtual machines. Further research into this area has highlighted the open-source Ubuntu Enterprise Cloud (UEC) as a possible provisioning mechanism, and being open-source means we can keep costs down to a minimum.

UEC uses a platform known as Eucalyptus, which is what Amazon's EC2 service is built upon. The structure for this requires two server machines across two different networks; with client machines (the lab PCs) having a Linux based operating system that has virtualisation capabilities provided by the KVM module.

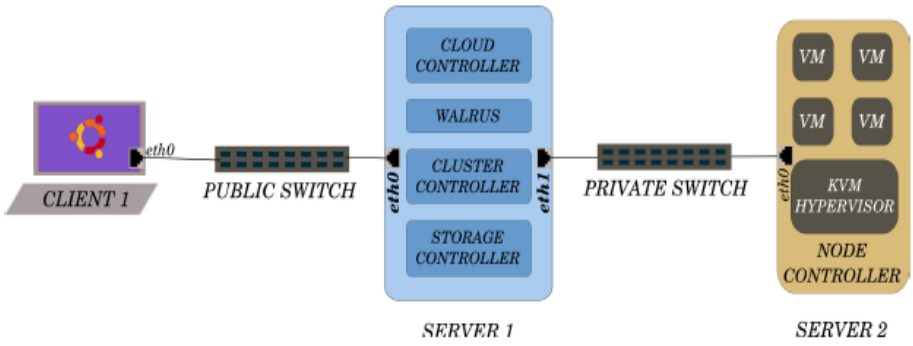

Figure 2 UEC Physical Topology, source: Eucalyptus Beginner's Guide: UEC Edition

Virtualisation

Virtualisation lets you run multiple virtual machines on a single physical machine, with each virtual machine sharing the resources of that one physical computer across multiple environments. (Source: VMWare.com)

There are many levels of virtualisation, such as bare-metal and OS-level; both have their advantages and disadvantages.

Bare-Metal virtualisation (also known as full virtualisation) is where a special kind of operating system (a hypervisor) is used instead of the traditional Windows or Linux platforms that we are virtualising, this hypervisor allows access to levels of the hardware that other methods do not allow, and as such this means that the virtual machines can be manipulated and adjusted as if they are on real hardware. This method is commonly employed at an enterprise level; my own experience comes with running VMware ESX 4 on 3 identical machines to allow for failover clustering and vMotion, which allows you to migrate live machines between physical machines, maintaining your uptime. VMware estimates a consolidation ration of 15:1 when deploying a virtual solution to replace existing physical servers; this has economic and green benefits.

OS-level virtualisation is where an application installed within a Windows or Linux operating system is used to build, manage and run virtual machines; examples of this include Virtual PC which is used in the current implementation, VMware Workstation, and Oracle's VirtualBox. One other OS-level product that exists is KVM, the Kernel Virtual Machine, this is a Linux only platform and "Using KVM, one can run multiple virtual machines running unmodified Linux or Windows images. Each virtual machine has private virtualized hardware: a network card, disk, graphics adapter, etc." (www.linux-kvm.org). this is to be the chosen virtualisation platform for the solution, as it is what the Ubuntu Enterprise Cloud uses as its platform.

Recommendation

Based on the research above a decision has been made to produce an infrastructure with UEC as a private cloud used to deploy virtual machines using KVM to the desktops, this means processing is done at the server (cloud) end and for the users there should be little impact on the local resources. The implementation must investigate how to

save changes to virtual machines as they will be created on the fly as and when they are required. There is a level of education required for the staff to manage the system once it is operational and discussions will be held throughout as to how this project can progress.

Storage is an issue, as virtual machine images can be very large, and with up to 20 instances of each at one time a lot of storage will be required, a bill of materials will be produced with basic and recommended configurations in mind. UEC can use iSCSI which is a storage controller mechanism that is cheaper to run than an enterprise level storage area network (SAN) and this will be further investigated.

I have produced the topology below to show the initial idea for the solution (Figure 3).

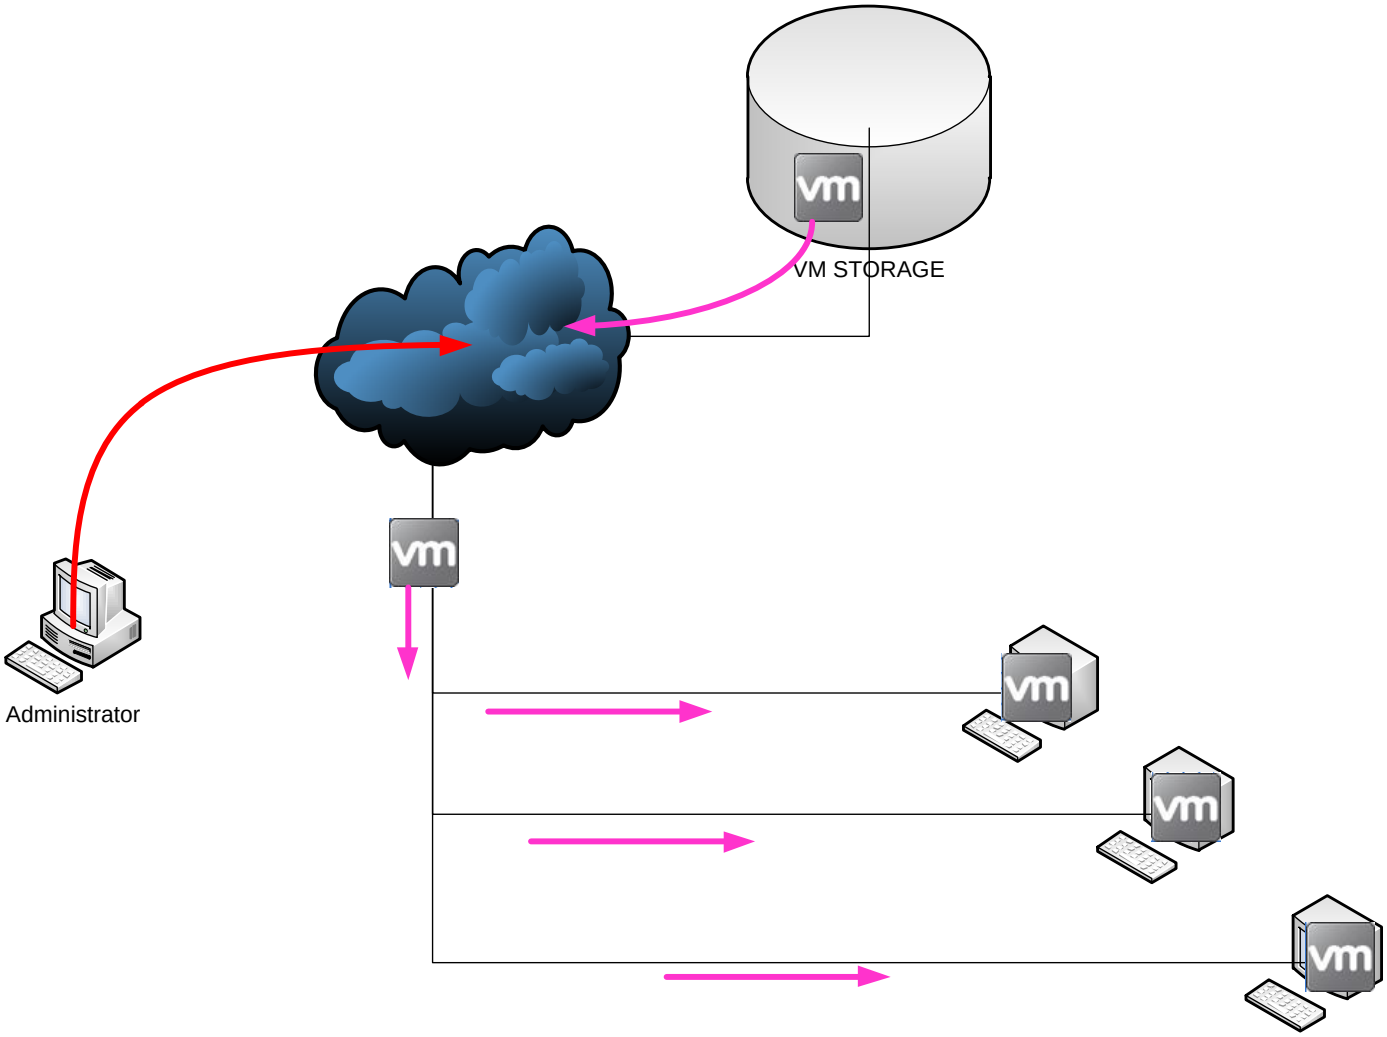

Figure 3 Proposed Solution

Primary Research

I have conducted a questionnaire to gather the thoughts of the students who currently use the virtual labs. From this questionnaire my aim was to produce a requirements list for the final solution and to use this as a measure in my evaluation.

The questions asked were:

1. What year of study are you?
2. Do you like the labs being in a virtual environment
3. Do you find the virtual machines easy to locate and launch?
4. What concerns do you have with these virtual machines?

5. What would you consider an acceptable start-up time for each virtual machine?
6. Do you have any ideas on how the virtual labs can be improved from an infrastructure perspective (i.e. not the lab content but how the virtual machines operate)

I had 38 responses, with 20 completing the optional open ended question 6.

The results are shown below.

1. 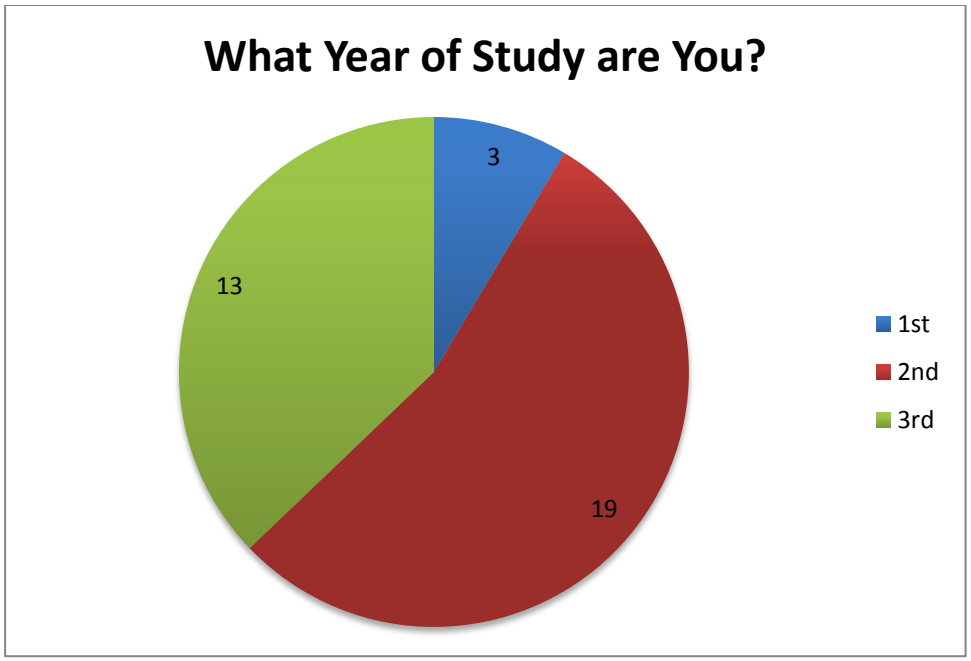

2. 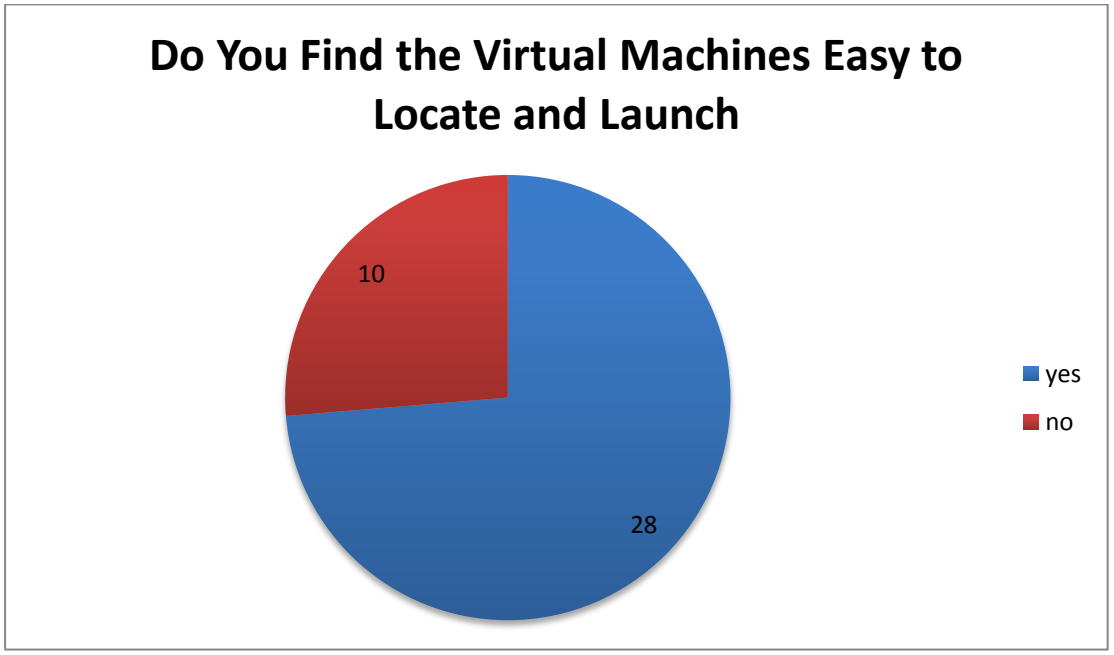

3. 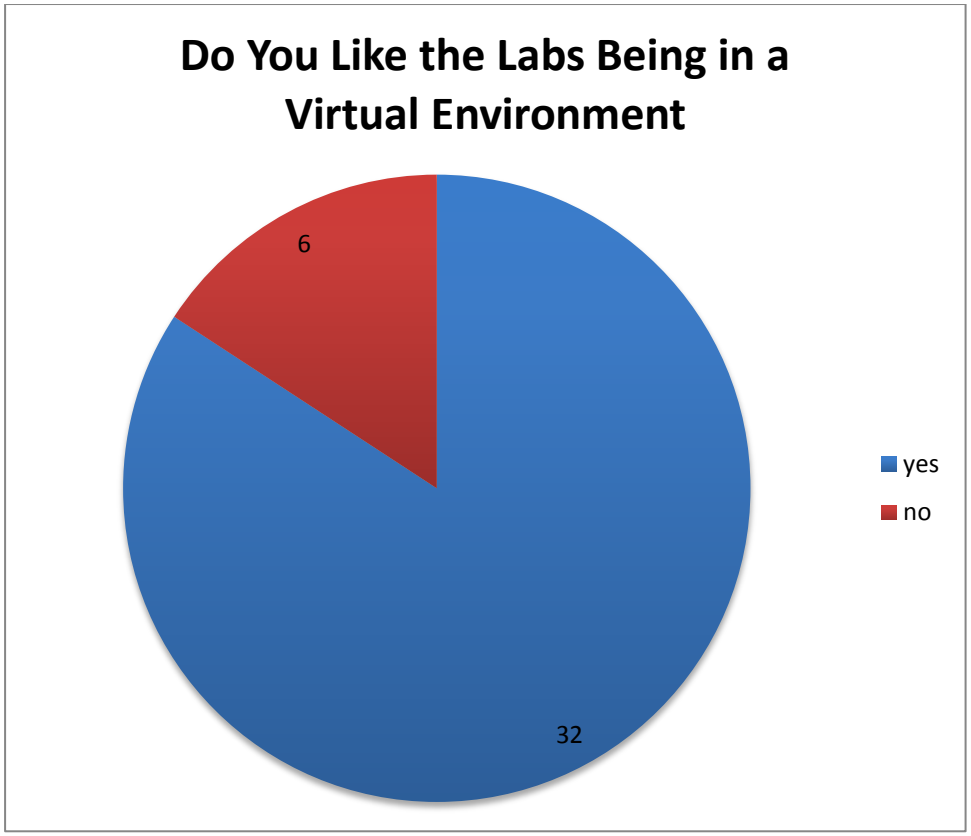

4. 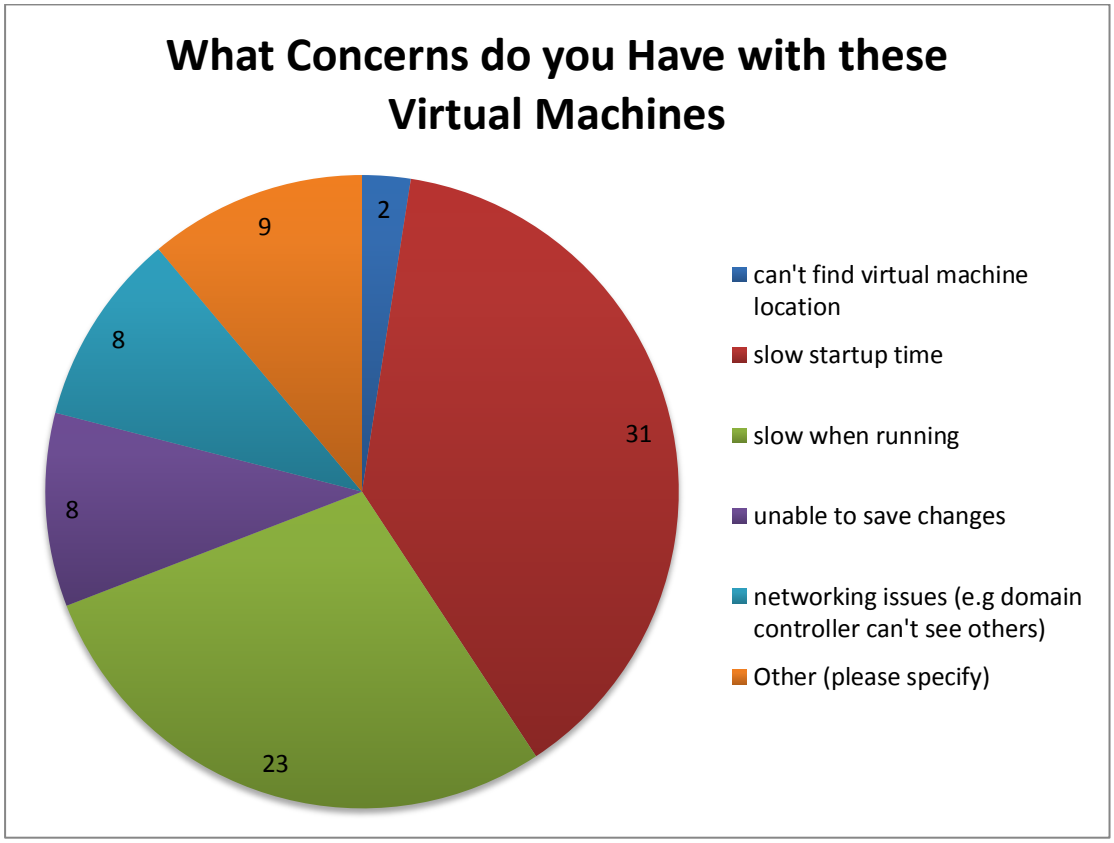

The “other” responses were (9 responses):

- Taking an exam, in 40 minutes on virtual machines that take at least 10 minutes to restart. You do the math
- Changes are still there from previous person. Shouldn't need to enter a new password every lab.
- When starting step by step configuring the lab then suddenly stops and i don't know the issue and i have to restart the whole process again and waste me time and achieving nothing.
- Can’t locate them in all of the rooms when you need to practise for exams
- Whilst I don’t find it hard to find the VMs, they could be more clearly signposted for beginners
- No comment

- Computers restarting
- No concerns
- No direction from Tutor, should at least be actively teaching/explaining the lab to the class. Not waiting until we don't understand.

5. 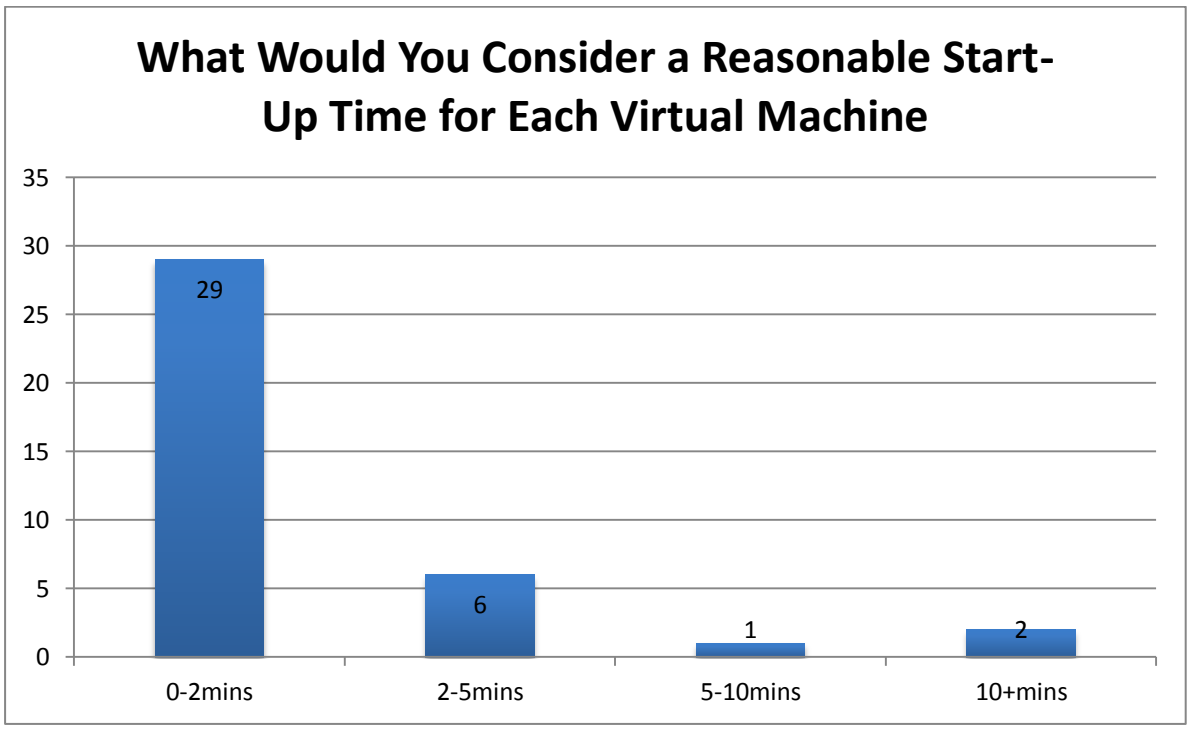

6. Written Responses shown below, 20/38 participants answered, 2 answers omitted.
   - If the machines were linked to the student disk it would be easy to save work, and maybe more powerful machines or better memory would help with start-up times?
   - Placing it on a better infrastructure would solve most issues
   - Smoother operation when multitasking in the virtual machine environment.
   - GET BETTER COMPUTERS! 2. You need at least 8 gig of RAM per each machine to run 2+ VM on it. 3. Just get a server rack, fill it up with server blades and use the computers as dumb terminals to connect to the server that will run all the virtual machine.
   - Remove all passwords. Quicker PC's
   - They need to be faster
   - Perhaps finding a method to shrink the VHD's to reduce start-up time?
   - Reduce the start-up time
   - Use VMware instead of virtual pc
   - I don’t know how this setup. But should be able to reset once the work has done on it, and should be load balanced in speed, and performance that not one virtual work faster and others slower.
   - More instructions could be given and important topics covered in the labs
   - Just generally faster, as when re booting during a test takes a lot of time, and you aren’t able to do anything until they reboot themselves.
   - More options to play with
   - Please let the lecturers explain the lab aim and objectives first, and then do the lab
   - The start up should be faster
   - The start-up time should be less than a minute because this allows the student to have more time to do the work
   - Having separate machines to store the server OS's.
   - Allow us to save current progress on a working VM

Primary Research Analysis

It is important to analyse the results gained above, in order to better understand what the end-users (students) want out of a system. Whilst the "client" for this project is the university staff, their clients are the students, and it's their experience that is going to be enhanced.

56 responses suggest that the speed of the virtual machines is a primary concern to most users. This is one of the problems that the cloud computing model will attempt to address, as the computing is done within the cloud not on the client machine.

Analysis of the written responses suggest that the students have a decent grasp of the technologies involved with provisioning the current virtual environment, and they are aware of the power required to correctly implement the labs in this way.

Some responses have concerns regarding the setup if the actual virtual machines, including a couple of references to passwords used to log in to them. This highlights the fact that even if the provisioning is improved, there is a human element involved with ensuring the setup is satisfactory, and perhaps the university staff could use this questionnaire as a basis for them improving their "service".