

## **Teaching Case**

# LINUX, Virtualization, and the Cloud: a hands-on student introductory lab

Anthony Serapiglia  
Anthony.Serapiglia@stvincent.edu  
CIS Department, St. Vincent College  
Latrobe, PA 15650

## **ABSTRACT**

Many students are entering Computer Science education with limited exposure to operating systems and applications other than those produced by Apple or Microsoft. This gap in familiarity with the Open Source community can quickly be bridged with a simple exercise that can also be used to strengthen two other important current computing concepts, virtualization and cloud computing. Students can gain firsthand experience through the creation of virtual machines utilizing free LINUX distributions with pre-installed productivity software. This enables the students to have a greater understanding of these concepts and how they may be applied to meet their own needs, or the requirements of a corporate scenario.

**Keywords:** LINUX, Virtualization, Cloud Computing, Operating Systems, Open Source Software, Open Document Format

### **1. INTRODUCTION**

Computer Science / Information Systems / Information Technology (CS/IS/IT) education is constantly striving to keep pace with the developments and evolutions of industry implemented Information Technology. For graduates with Computer Science degrees to be able to hit the ground running and garner jobs post-graduation that are above entry level positions requires the education they receive in the collegiate school to be as near as possible to the hard knocks school of the work world.

While it is understandable that it is impossible for the school curriculum to respond to every change and new fad of the computing industry, especially in smaller schools with tighter budgets and fewer resources, it is incumbent upon every school to incorporate concepts and practices that form the fundamental foundation of computing in the business world.

Three of the most prevalent topics in the IT industry since the turn of the millennium have been Cloud Computing, Open Source software, and Virtualization. Each can be taken individually, but often they are seen working together as part of a greater system. These three areas have grown and matured to where companies of any size must address the questions that each poses.

The following lab centralizes an activity that can present to CS/IT/IS students on any level a valuable experience with Virtualization, Open Source software, and Cloud Computing. With proper setup and resources, the core exercise can be accomplished in under an hour, and easily expanded as needed for longer time windows. The exercise can also be augmented by a multitude of additional assignments and tasks to expand upon the experience, or to further apply the concepts in different environmental scenarios.

In this lab, students will create virtual machines (VMs) utilizing readily available, often free,

virtualization software. These VMs will have as their operating system (OS) one of many open sourced LINUX distributions which are also free of charge. It is most beneficial to include multiple 'flavors' of LINUX to include a variety of pre-installed options for Web browsers, office productivity software, as well as look and feel of file management and desktop environments. Students will also utilize Cloud Computing concepts in managing files for ultimate delivery of the artifacts related to the assignment.

After experiencing these areas of current computing trends hands-on, students will be prepared to address further questions and scenarios projecting the concepts and applications of the technologies into various situational based assignments catered to their educational level and course of study.

## 2. BACKGROUND

For several years I have straddled the line between academia and the business world as a full time Managed Service Provider (MSP) consultant to Small/Mid-Sized business (SMB) and an adjunct professor of Computer Science and Information Systems. Having a foot in both worlds allowed me to continue to develop at a very rapid rate with the new technologies entering production in the business world, as well as relate those techniques and tools to students not yet exposed to the vagaries of that environment.

Through anecdotal evidence I became aware that many of my students were coming into Computer Science classes with a base recognition of many words associated with computers and computing, but with a very wide range of definitions of what these words meant or what actually accomplished the tasks. In an interview with Newsweek, Inventor Dean Kamen stated, "If you are not on the front edge of the wave (of innovation), it's harder to understand it, much less be involved. You used to be able to open the hood of a car and understand, that's the engine. Most people today look at technology and see magic (Kenny, 2004)." The students that I found in my class room were walking in knowing that magic was occurring, and not knowing how. They were on the tangential edge of consumption of computer devices and computing tasks, certainly not on the leading edge of implementation - let alone innovation.

This lab was initially developed as part of a class in operating systems being delivered to students working towards a Master's degree in Information Systems. To my surprise, an informal poll on the first day of class showed that not a single student of the twenty-four had ever used LINUX directly to their knowledge. In response to this, I devised a series of progressive labs utilizing VMWare Workstation to introduce the students to LINUX as a desktop OS and progressing to their installation of a server system including several separate user accounts and an FTP service for 'drop box' scenarios. This initial trial worked very successfully with many comments from the students expressing gratitude for the inclusion of hands on practical tasks.

Subsequently, I have adjusted the level of instructions and tasks to accommodate undergraduate introductory classes. In the graduate classes little instruction was provided, mandating that the students find their own resources to support and direct their tasks. In the undergraduate classes a more detailed set of step by step instructions was necessary to hand hold the students through the initial setup process. However, once at the desktop of the LINUX OS, students were expected to 'find their own way' by looking for familiar words, menus, and icons. The tasks were also adjusted to the undergraduate level as more of an introduction to the possibilities of replacing costly licensed software with free open sourced equivalents.

In the past year, I have included a more formalized survey to begin the exercise. This survey has been administered to three classes, a non-major introductory course, a freshman level introductory course for majors, and a senior level special projects course for majors. As expected, results varied greatly for the senior level course and the two introductory courses. A full ten of ten had experience in using LINUX on some level in the senior level course. However, there was little difference between the two introductory level courses. In the non-major course, 24 of 24 did not have any experience with LINUX, and 23 of 24 did not have any LINUX experience in the intro course for CS majors. One-hundred percent of the students in all three classes responded that they had had some experience with both Apple OS's and Microsoft OS's. Between the two intro courses, 8% (4/48) responded that they believed that LINUX to be a command line only operating

system with no point and click graphical user interface.

The survey also included questions regarding how students access their e-mail. Answers to these questions have highlighted the need to address core concepts of client server systems to the students. Not a single student utilized a standalone/installed client (such as Microsoft Outlook or Opera for MAC) application in the traditional sense to access their e-mail. All students responded that they exclusively utilize the web interface for either school e-mail or other public services such as Gmail, Hotmail, Yahoo, AOL, etc. In a subsequent question, 56% (27/48) incorrectly believed that in utilizing the web interface to view an e-mail message, that message was automatically saved to the local device on which it was viewed. For file sharing, 62.5% (30/48) had utilized some form of Internet based file sharing with Google Docs, Windows Live, Picassa, or Dropbox. Every respondent (48/48) believed that uploading a picture to a social networking sites such as Facebook or Twitter was the equivalent of file sharing.

Prior to the lab, two of the students in the upper level class had experience with running a virtual machine (VM). Both had MAC books and were currently running Windows 7 as a VM. No other students responded affirmatively that they had previous experience with a virtual machine.

### 3. LAB EXERCISE

#### Environment:

The environment for the lab can be very flexible. The most essential piece needed is a virtualization application. For this lab I have been utilizing VMWare Workstation; however I have also performed the lab with VMWare Player and Oracle's Virtual Box. VMWare Player and Virtual Box are free; VMWare workstation is not, although a trial version is available.

I have been lucky to be in a lab that I have administrative control over. I can ensure software is pre-installed and ready. I have also performed this lab in settings that utilized Deep Freeze for protection of the machines. In these cases, the installations worked perfectly fine, as long as students did not re-boot after completing the install of the virtualization software.

LINUX distributions are many. Distrowatch (<http://www.distrowatch.com>) is a terrific

resource for tracking the most current and most popular distributions. Each version offers their own take on what they think is best, and customized to fit various environments. There are many that are for general use and provide multiple applications pre-installed for convenience. I have utilized Fedora, Knoppix, Ubuntu, LINUX Mint, Puppy LINUX, and Pear OS. All of these distributions are available for free and downloadable. Most are in the 700MB range in size, though some will grow to over 1GB. These Operating Systems (OS's) are available as disk image (ISO) files that can be loaded directly by the VM software or burned onto and then booted off of DVD's and also flash drives. It may be beneficial to have these available on a flash drive or DVD for ease of copying from computer to computer to avoid a larger class all pulling the same larger file down from the Internet at the same time. This is a good opportunity to stress familiarity with the file system of the computer they are working on. Placing these files on a network share can work, but will also create a terrific amount of network traffic that will slow everyone down. Keeping the disk image for the OS local will avoid this.

#### Tasks:

Most often I have included the installation of the virtualization software as the initial task for the students to perform. I believe that we often take it for granted that all students are familiar and comfortable with simple installation procedures. I have often found this is not the case and many do benefit from the process.

Once the installation of the VM software has been completed, students are tasked with creating three different virtual machines one at a time. This is accomplished by first creating the shell for the VM in the software. This requires designating how much of the resources of the physical host machine can be dedicated for use by the VM and includes CPU, RAM, and Hard Drive space. A point of emphasis can be made here as to what exactly does make up the basics of any computer system. Once these resources are designated, it should be stressed to students that they have essentially created a new machine, just like the one physically in front of them, and that at this point is has all the 'physical' parts necessary but no operating system. In all of the VM software packages, they then point the virtual optical drive to the ISO file for the LINUX OS they wish to run. Upon "powering up" the new VM, the new LINUX OS will automatically begin to load. This process is

repeated for the subsequent VM's as well. It is important to note to students to make sure to power down their previous VM's before powering up a new one. More than one on at a time will begin to tax the resources of the host and noticeably slow their progress.

Within each OS a separate basic assignment is given to highlight the similarities of the OS with what they are already accustomed to. Each of the LINUX distributions will have a version of an open source office suite for word processing, spreadsheet, presentation, and possible personal database work. Each will also have a web browser program. Simple assignment can be centered around searching for current standing for the college's sports teams that are in season, importing into a spreadsheet, and including a picture that was saved from a website on the spreadsheet. Once completed, the student must e-mail the saved spreadsheet to the professor. This type of assignment will force the student to explore the OS to find the programs necessary to create the artifact, as well as to send it off for completion. The inclusion of the image file will force the student to save a file from the Internet onto the file system of the VM and find it to include into the spreadsheet enforcing the existence of a different and separate file structure from the host machine. I recommend including different assignments for different OS's with multiple variations of open source office applications and browsers for variety. Screen shots from within the VM can be included into documents. Scavenger hunt questions can be asked and answered to provide content for documents. Games can also be included.

E-mail is a preferred method of delivery for the artifacts of the assignments. This enforces the ease of access of mail systems and the use of various browsers to access them. It also reinforces the idea that even within the VM, there are still multiple connections to the outside to direct output from the system.

#### 4. CONCEPTS

The following is a list of concepts that can and should be stressed throughout the lab. These can be enforced through the specific assignments given within each VM instance, or simply highlighted with running dialog between the instructor and the class during the session.

- Components of a computer
- File system navigation

- Operating systems
- Open source software
- Output channels from a computer
- Variety of software for common tasks
- Utilization of computer resources
- Virtualization
- Cloud computing, personal
- Cloud computing, corporate
- File type in saving documents
- Compatibility of software, file types
- Commonality of menus and icons

#### 5. ASSIGNMENTS

The purpose of the preceding lab exercise is to ensure that all students have been exposed to the essential elements of virtualization, open sourced OS and application options, and virtualization. It is not meant as exhaustive training in any of these areas. By introducing and allowing the students to experience these applications personally, it will remove barriers of ignorance and allow them to move forward to greater conceptual assignments. These assignments can be centered on specific applications or scenarios as would most benefit the course in general. The following are examples of possible longer form assignments that build upon the concepts introduced in this lab:

##### Assignment 1:

You have been hired as a consultant to a Small/Mid-sized business that is less than two years old. Currently the office has 11 employees. The company has 8 PC's that are mixed between Windows XP and Windows 7. Five employees are utilizing personal laptops for work as well, three of these are MACs. The owner would like to standardize and make sure everyone is on the same page and can work together seamlessly. Provide a proposal to him that would utilize open sourced software and a LINUX based operating system.

##### Assignment 2:

What is the current market share breakout of operating systems? Find four measures of this statistic. Explain how each of the measurements is derived. Pick four operating systems and provide an overview of their history, their strengths, weaknesses, and their public perception.

**Assignment 3:**

Your new client is a food marketing company that produces pre-made appetizers for several restaurant chains. There are five regional offices across four states. Multiple sales people and managers are in the field supplying and selling products to restaurants. You need to develop a system of portable file sharing that can work across multiple OS platforms including MAC Air Books, Windows based laptops, iPads, and Android based tablets. Identify four leading cloud file sharing services. Compare their features and costs. Make a recommendation as to which you think will work best for this company. Make sure to include in your recommendation the following: cost, collaboration features, security, multiple OS integration, multiple file type compatibility, reliability, offline options, and backup/recovery options.

**Assignment 4:**

You have just been hired as a quality assurance specialist at a small software development company. The company develops an installable application for Windows based PCs. With each revision of the software, the application need to be tested in various environments, including multiple versions of Windows and multiple versions of Windows with multiple other common software applications installed on Windows. Scenarios also include fresh installs, as well as upgrades from existing installations. The previous person in charge of this task had

14 PCs that they would need to re-install every time a new revision of the company software came out. Your boss wants you to re-think this system and make it more efficient.

**6. REFERENCES**

Bittman, Thomas J. Server Virtualization Trends, 2012. Retrieved May 23, 2012, from [http://blogs.gartner.com/thomas\\_bittman/2012/03/21/top-five-server-virtualization-trends-2012/](http://blogs.gartner.com/thomas_bittman/2012/03/21/top-five-server-virtualization-trends-2012/)

Stone, Brad. (2006, December 3). Stuff That Will Take Your Breath Away. Newsweek, from <http://www.thedailybeast.com/newsweek/2006/12/03/stuff-that-will-take-your-breath-away.html>

VIZARD, M. (2009). GAUGING LINUX MOMENTUM. Baseline, (96), 11.

[http://stats.wikimedia.org/archive/squid\\_reports/2012-4/SquidReportOperatingSystems.htm](http://stats.wikimedia.org/archive/squid_reports/2012-4/SquidReportOperatingSystems.htm)

<http://www.distrowatch.com/>

<http://www.netmarketshare.com/>

<https://www.virtualbox.org/>

<http://www.vmware.com/>

## Appendix A: Lab Instructions

### Ingredients:

- Virtualization Software (Oracle Virtual Box) - <https://www.virtualbox.org/wiki/Downloads> (Download the Windows version.)
- "LIVE OS" versions of LINUX distributions: LINUX Mint, Puppy LINUX (<http://distrowatch.com/>)
- A workstation, either a Windows or Apple PC, to host your virtual machines

### Installation of Virtualization Software:

- Acquire install package for Oracle Virtual Box (95MB install file)
- Run the installer (Each instruction line is a new window in the installer)
  - o Click next past the welcome screen
  - o Take the default location for installation (C:\Program Files\Oracle\VirtualBox)
  - o Choose whether you would like shortcuts placed on desktop or quick launch bar
  - o Reply "Yes" to the warning about the networking interface install
  - o Reply "INSTALL" to the ready to install prompt
  - o The installation will proceed – there will be several prompts to OK component installations from Windows. Choose either OK or Install to proceed.
  - o Leave the check in the box to start the program after installation. Click "Finish" when the installer is complete

### Creation of first virtual machine:

- Open Oracle VM VirtualBox Manager
- From the Machine menu, choose "New"
- In the new window that appears, provide a name form your machine
  - o Choose "Linux" as Type
  - o Choose "Debian (64 bit)" for Version
- In the next window, set the amount of memory to 1024MB
- In the next window, take the default, which creates a virtual hard drive at the recommended 8GB size
- In the next window, choose the VMDK option for Hard drive file type
- In the next window, take the default to Dynamically allocate the hard drive
- In the next window, take the default that confirms the 8GB drive size

### Choose OS and start machine:

- Before powering on the machine, set direct the machine to the proper boot image
  - o In the Oracle VM VirtualBox Manager, select he machine listing you just created
  - o Click on the "Settings" gear icon
  - o In the new window, select "Storage" from the left side menu
  - o In the middle pane, select the listing for "Controller: IDE"
  - o To the right of the "Controller: IDE" listing are two disk icons, choose the first wich shows only 1 platter
  - o In the window that comes up announcing that you are about to add a new CD/DVD drive, select "Choose Disk"
  - o Browse to the location you have the ISO file for LINUX MINT and select the ISO file
  - o Place a check in the box or "Live CD/DVD"

- Start the new virtual machine
  - o Press the "ICON" icon to start the machine
  - o Through the boot process, several advisory screens will appear with information about pointer capture and color settings. "OK" past these.
  - o As a "live OS" you will find yourself at a desktop in just a few short moments without the need to login.

### **Exercise 1: From MINT**

- Once at the desktop
- Open a word processing program (look around – you'll find it!)
- Insert a picture related to Girl Scout Cookie Thin Mints with an appropriately witty remark.
  - o To accomplish this, you must also find an internet browser, find the image, save the image to the virtual hard drive of the virtual machine, and retrieve the image from that location to insert into your document...
- Save as an "Open Document" format (.odt) file name "yourname\_Mints.odt"
- From the internet browser open your school e-mail. Send the document to me as an attachment.
- Exit from the applications and power off the virtual machine.

### **Exercise 2: From Puppy**

- Create a virtual machine for Puppy just as the first
- Upon power on, you will need to make several selections during the initial boot sequence. Puppy is not a "LIVE OS" and is in fact installing itself to the new virtual disk.
  - o Select USA
  - o Select en\_US
  - o GMT-5 for time zone
  - o Local clock type
  - o Xorg for video type
  - o Probe for video type
  - o 800x600 screen size
  - o Verify 800x600
  - o Finish
- Once at the desktop, Open a spreadsheet program (Gnumeric Spreadsheet Editor)
- Create a quick spreadsheet. Copy the schedule for the school's baseball team
- Include a picture of a baseball on the spreadsheet
- Save the spreadsheet as "yourname\_Puppy".
- Open a web browser, log into your school E-mail
- Mail the spreadsheet to me.

### Appendix B: Student Survey

1 – Have you ever used LINUX?      Yes                      No

(For the following – use a scale of 1 to 10, 1 low to 10 high)

2 – How comfortable would you be creating a document in a program other than Microsoft Word \_\_\_\_\_

3 – How hard would it be for you to adjust to a new Operating System? \_\_\_\_\_

4 - Have you ever used a computer with the following Operating System?

Windows \_\_\_\_\_ MAC OSX \_\_\_\_\_ MAC IOS \_\_\_\_\_ Android \_\_\_\_\_ LINUX \_\_\_\_\_

5 – Rank the following common tasks with how important they are in your everyday computer usage:

1 - Web Browsing                      \_\_\_\_\_

2 - Word Processing                      \_\_\_\_\_

3 - E-Mail                                      \_\_\_\_\_

4 – Printing                                      \_\_\_\_\_

5 – Spreadsheets                      \_\_\_\_\_

6 - File Sharing                              \_\_\_\_\_

7 – Music                                      \_\_\_\_\_

8 – Video                                      \_\_\_\_\_

6 – Which Operating System do you have on your personal computer? \_\_\_\_\_

7 – What is your most commonly used web browser?

1 – Internet Explorer

2 – FireFox

3 – Google Chrome

4 – Safari

5 – Other: \_\_\_\_\_

8 – In what program do you work with your e-mail?

Stand Alone Client (Outlook, Opera)? \_\_\_\_\_ Web Browser \_\_\_\_\_ Phone \_\_\_\_\_

9 – If you utilize a web browser for e-mail, where do you believe are these messages saved? \_\_\_\_\_

10 – Do you use any online file storage systems/services? If yes, which?

11 – Do you believe that uploading a picture to a social networking sites such as Facebook or Twitter was the equivalent of file sharing? \_\_\_\_\_