December 2002

TEACHING MULTI-TIER SOFTWARE DEVELOPMENT USING VIRTUAL MACHINES

Kyle Lutes
Purdue University

Alka Harriger

Katherine Fanjoy
Purdue University

TEACHING MULTI-TIER SOFTWARE DEVELOPMENT USING VIRTUAL MACHINES

Kyle Lutes, Alka Harriger, and Katherine Fanjoy
Purdue University
kdlutes@tech.purdue.edu arharriger@tech.purdue.edu
kfanjoy@aol.com

Abstract

The increased use of distributed computing platforms has companies looking for software developers skilled in multi-tier application development. Universities need to discover ways to teach these skills while at the same time keeping costs for hardware, software, administration and floor space at a minimum. This paper reports how the VMware virtual machine software product can be used in a general-purpose computing laboratory to realize these goals. Reasons why a product like VMware can be useful, a general discussion of virtual machine technology, and its use in two semesters of an advanced programming course are described.

Keywords: VMware, virtual machine, multi-tier, client/server, computer laboratory

Introduction

The increase of application service providers and use of e-commerce systems, along with the decrease in hardware costs, has led to increased popularity in distributed computing (Moltra, 1999). Companies are looking for software developers skilled in Java, C++, and Visual Basic as well as Internet development, networking, database management, and a variety of operating systems (Computerworld, 2000). This range of topics and computing platforms is diverse, and universities may find it difficult to finance and manage the different architectures and operating systems that students need to learn about in order to fulfill industry needs. Since industry is looking for graduates who understand these topics, universities need to discover ways to provide the necessary training while keeping costs for hardware, software, administration and floor space at a minimum.

To address this problem, the Computer Technology (CPT) Department at Purdue University used the VMware virtual machine software product to provide greater flexibility in computer usage. The goal of using VMware was to provide students access to computing privileges necessary to develop enterprise applications, while at the same time ease computer lab management and reduce the amount of required physical lab space. According to the Market Data Retrieval (MDR) College Technology Review: 2000 - 2001 Academic Year Report, 75 percent of all colleges and universities use Windows as their network operating system (Microsoft Corporation, 2001). Although there is documentation from a variety of sources on using VMware under Linux, documentation for VMware under Windows is relatively nonexistent.

Universities do not always have the capital to invest in new or experimental technology such as this, and therefore could benefit from a discussion concerning the positive and negative aspects of using VMware within an educational setting. Therefore, documenting the use of VMware in the enterprise application development course may encourage other institutions using Windows to evaluate the prospect of using VMware to improve their own educational offerings.

Enterprise Application Development

The CPT Department offers a software development course, CPT 450 Enterprise Application Development, that teaches students how to develop and deploy multi-tier client/server information systems. Started in the Fall 1999 semester, this course was created to offer a greater depth of programming topics taught within Purdue's School of Technology. CPT 450 covers topics such as

Object-Oriented Programming (OOP), 2-tier and *n*-tier client/server architectures, software layering, database connectivity, multi-user concurrency issues, and application deployment and distribution. The course is currently implemented using the Microsoft Windows Distributed InterNet Architecture (DNA) enterprise environment.

Students are taught to encapsulate business and data handling logic into reusable COM components. These components connect to the Microsoft SQL Server through ActiveX Data Objects (ADO). The components are registered in Microsoft Transaction Server (MTS), or COM+, on the server. Students create web user interfaces to these components using Active Server Pages (ASP) that are run under Internet Information Server (IIS). Visual Basic (VB) is also used to create rich graphical user interfaces that connect to the same components through Distributed COM (DCOM). This architecture is frequently distributed over several computers with one machine acting as a client running the user interface, one running the web server, another as an application server running MTS/COM+, and yet another as a database server.

Unfortunately, COM components must be registered in the Windows Registry, which requires administrator security privileges, and giving students administrator privileges could potentially compromise network stability and security. Another problem was that developing web applications is best done running IIS on the same computer as the development tools. When the course was first offered, students performed coursework on computers in general-purpose, public labs, which did not include IIS and MTS. And since students did not have administrator privileges, they were not able to easily code, test, and deploy the applications they were required to develop for the course assignments. To alleviate this problem, the course was provided with its own stand-alone computer lab during the second offering of the course.

This lab contained 10 computers running NT Server 4.0, which served as student workstations. Each workstation could run its own web, application, and database server. Students enrolled in CPT 450 were given user accounts with administration privileges so that they could register their COM components. Since students could install software on these computers, the lab was connected to the Internet through a firewall.

While this setup worked well, there were a few weaknesses. Due to the limited number of computers, the course was limited to a small enrollment to ensure each student had a computer during lab. In addition, students would often save work to the hard drive, where files would then be available for other students to find. This practice led to loss of student privacy as well as an increased risk of plagiarism of another student's work. Finally, the course instructor was responsible for installing and configuring the computers at the beginning of each semester, which required an average setup time of one to two hours per computer. Despite these problems, this solution worked well because it provided the flexibility necessary to teach the course. Unfortunately, the school was unable to justify providing the resources for this lab since few students used it and floor space was already scarce.

As a compromise, the school's network support staff suggested we implement a software application called VMware, which would provide the flexibility needed to teach this course in a general purpose, public lab. VMware offers the ability to create several "virtual machines" on one computer, called the host machine, which gives the user the ability to run multiple virtual operating systems that under the host machine's operating system. Therefore, one physical computer could actually be running several "guests" operating systems at the same time. Communication between the host and guests is possible through networking options. This solution offers great potential, since each student has the ability to operate and administrate his or her own distributed computing environment.

Since this product was relatively new, most of the literature available was concerned with basic aspects of implementation: hardware and software requirements, installation of the VMware software, and general configuration (Hall, 2000; Katz, 2000; Lynch, 2001; Rado, 2000; Railsback, 1999). Other sources offered greater levels of detail, but the implementation was specific to Linux. For example, Neih and Con Leonard (2000) detail the use of VMware and Linux to grant students privileges that allowed kernel development, which would have been prohibitively expensive otherwise. Fonvieille (2000) provided step-by-step instructions to install, configure, and set up a network using VMware for the Linux operating system, and Fuller (2000) shared some problems with installing VMware in Linux. Though this information provided a baseline of what VMware was capable of doing, it was not as helpful for fulfilling the goals of CPT 450.

The VMware documentation provided with the software offers a great deal of information concerning requirements, installation, and implementation in a Windows environment (VMWare, 2001a). It also has a relatively complete description of configuring networks using virtual machines. Still, these are not always helpful since the user is assumed to have certain permissions on the host machine. Given that the host machine in this case was already part of another network, users did not have administrative privileges to change the configuration of the host. Therefore, the solutions provided were not always applicable.

After reviewing several information technology related publications and product websites, it became evident that only a few universities, such as Columbia University and Clemson University, have some documentation regarding their experiences using virtual machine technology. Still, the applications of the technology were in desktop management and operating systems development, which is not closely related to its use in an enterprise development course. A detailed description of the implementation and outcome of using VMware in a distributed Windows programming environment could not be found in the literature. This paper addresses this gap by sharing the classroom experience of using VMware in an undergraduate course on enterprise application development.

Virtual Machine Technology

Virtual machine technology such as VMware offers users the ability to run multiple operating systems simultaneously on the same computer. An operating system (OS) running under such a system can be rebooted, modified, crashed, and reinstalled without affecting other applications running on the computer. This technology has eliminated the need for rebooting or repartitioning if the user desires to run multiple operating systems. An operating system installed under the virtual machine software is generally called a "guest OS", while the native operating system is known as the "host OS".

The concept of a virtual machine was first used for IBM mainframes in the 1960s through the VM operating system. Since mainframe machines were very expensive, the VM operating system could be configured to support multiple client operating systems, eliminating the need for purchasing multiple machines to perform multiple functions (Sims, 2000). As the cost of hardware decreased and the market became more diversified during the 1980s, virtual machines fell out of favor.

During the 1990s, the personal computer (PC) became increasingly important as a business tool. With systems such as Macintosh, Windows, and Linux competing for popularity, the software industry felt an increasing need to create applications that could run on multiple operating systems. The concept of a virtual machine resurfaced with the advent of the Java language, which runs code in a virtual machine written for a target platform (Babcock, 2000). Nonetheless, the majority of applications were written for specific platforms.

This specificity created problems for users who employed applications that ran on different operating systems, especially UNIX users who needed to use Windows productivity applications such as Word or Excel (Petreley, 2000). To solve this business problem, the concept of the virtual machine returned to favor. This may be attributed to the flexibility and convenience of benefits such as saving on hardware and floor space, simplifying management due to a reduced number of physical machines, reducing time and cost of developing and testing software on multiple platforms, and reducing set up time for operating systems (VMWare, 2001b). Since each guest OS runs independent of the host (and other guests), virtual machines can also be a tool for system administrators to test new products before they are installed on a live system (Bar, 1999).

VMware

VMware is "a virtual-machine platform that provides an abstraction of x86 PC hardware so that multiple operating systems can run unmodified and at the same time on a standard PC" (Neih & Con Leonard, 2000). Founded in 1998 by researchers from Stanford University, the developers of VMware sought to "to minimize total cost of ownership (TCO) of computing infrastructure by maximizing its manageability, flexibility, and efficiency" (VMware, 2002a). Developed for computer users desiring to run multiple operating systems simultaneously, the software is based on the company's MultipleWorlds technology, which "creates a thin software layer that sits between the Intel architecture and the operating system, virtualizing the hardware and managing all hardware resources" (VMWare, 2001b).

VMware was created to provide full functionality for several guest operating systems including Windows, DOS, and Linux. In theory, any operating system that runs on an x86 architecture (such as Intel or AMD) could run as a guest OS (Neih, & Con Leonard, 2000). Supported host operating systems include Windows 2000/NT and Linux. Since the Macintosh runs on a different architecture, VMware does not support it. In addition to providing a platform on which to run applications, VMware also provides the ability to create virtual networks within one machine. Depending on the type of network setup chosen, the virtual machines can communicate with other virtual machines, the host machine, or other machines on the network.

When VMware was first released in May 1999, reviewers were skeptical of its success: "we could not imagine a use for VMware anywhere in corporate America. The falling price of PCs makes it a product without a need, and one whose need might never

materialize" (Chowdhry, 1999). These concerns proved to be unfounded considering that the software company today has reached "more than 1 million registered users and more than 5,000 corporate customers in more than 100 countries" (VMware, 2002a). There are many cases of VMware being used in industry. Marimba has increased testing capacity by 70 percent while reducing hardware costs after implementing VMware in its testing environment (Macarlo, 2002). Development companies such as Wind River, Informatica, and NetIQ have improved multi-platform development efforts through the use of VMware. Most recently, it was announced that Microsoft would use VMware to allow its sales force to perform software demonstrations that previously required several computers (Shankland, 2001).

While the educational sector can also benefit from using VMware for testing applications in multiple environments, the primary interest in academe appears to be for providing students administrative computing experience without losing security. For example, Columbia University has adopted VMware in its operating system design course due to its ability to provide students administrative rights through the guest OS (Neih & Con Leonard, 2000). Clemson University uses VMware for the management of all desktops to provide greater flexibility for students (Babcock, 2000). As the benefits of VMware become more widely known, it is believed more educational institutions will consider using this technology as a way to provide a better learning experience for students while keeping costs at a manageable level (VMware, 2000).

Obviously, there cannot be such benefits without an associated cost. VMware 2.0, which was the version first used in the School of Technology, has been labeled as a "resource hog", and Babcock (2000) recommended that machines have at least 128 MB of RAM to function well. The guest operating systems also suffer slower performance, 50 to 65 percent of a computer's normal performance, than a native operating system (Mullen, 2000). For this reason, users are advised against running a heavily used operating system, multimedia application, or games under a guest OS in VMware (Randall, 2000). In addition, the guest operating systems do not support DVD, CD-R, CD-RW, Imation SuperDisk LS120 disk drives, or IDE zip drives. This reduces the types of storage devices that may be used with a virtual machine, so a greater amount of disk space may be necessary if a large amount of storage is needed. Other shortcomings of VMware 2.0 include limited graphics support and a maximum virtual disk size of 2 GB.

Despite the shortcomings, VMware 2.0 offers four valuable options to support networking among the host and guest operating systems: *no networking*, *host-only networking*, *bridged networking*, and *custom networking*. Communication between the virtual machines and the host is run through a virtual switch that resides on the host computer.

The "no networking" feature is self explanatory, and allows the guest OS to run in complete isolation from other machines. "Host only" networking allows a guest OS to communicate with the host machine as well as other guest machines that may be running on the host, but there is no communication between the virtual machine and the external network. "Bridged networking" uses a virtual bridge to allow the guest OS to communicate with other machines on a network, not just the host. "Custom networking" refers to a network configuration other than the ones previously described, intended for use by expert users who desire more specialized networking schemes.

Other features of interest are the various disk modes offered in VMware. When creating a virtual machine, the user may choose to use one of three modes: *persistent*, *nonpersistent*, and *undoable.* In the persistent disk mode, the virtual machine will store any changes that are made, essentially behaving in the same fashion as a regular hard drive. In the nonpersistent mode, no changes are saved to the virtual machine. This feature is useful for those who want the virtual machine to always start in the same state, such as for use with software testing or demonstrations. Finally, the undoable mode allows the user the option of saving changes to the disk when the virtual machine is powered off. This may be useful when performing tasks that may need to be undone if something goes wrong. For this mode, all changes are stored in a redo log. When finished using the virtual machine, the user is prompted whether to commit the changes in the redo log to the disk, discard the redo log, or save the redo log. By default, the redo log is stored in the same directory as the disk file, but a different location may be specified in the configuration menu. If the user chooses to save the redo log, VMware will offer several options the next time the machine is powered on. These include the choice to commit the changes stored in the redo log, discard the redo log, or continue to append any changes to that redo log.

VMware 3.1, which was just released in April 2002, offers some changes and improvements over the previous versions. A new repeatable resume feature is useful in a classroom setting where the instructor may need to restart the guest OS with applications already running instead of rebooting the virtual machine.

Though there aren't definite statistics at this time, VMware 3.0 promises improved performance in disk access, graphics, I/O, and networking (VMware, 2002b). It also promises to offer support for the upcoming release of Windows XP as both a host and a guest operating system as well as Windows ME as a guest operating system.

Another highly anticipated feature is the addition of support for USB devices, DVD-ROM (though not DVD video), and CD-R/RW. In addition, the minimum disk size has been increased from 2 GB to 128 GB for IDE drives and 256 GB for SCSI drives. Previously, SCSI drives on the host machine were treated as IDE drives (VMWare, 2001a). Support for ISO images has also been added, which would allow for the creation of disk images that act like CD-ROMs, including bootable CD-ROMs, which may be used by the guest system CD drive.

Another feature of interest is the support of a new networking option: Network Address Translation (NAT). This option is valuable if a user desires to have an Internet or other TCP/IP connection but is not able to have the virtual machine on the external network. When using NAT, the virtual computer is given an IP address on a private network through a virtual DHCP server on the host computer. The NAT device will then pass information between the virtual computers and the external network.

Both versions of VMware cost $299 to download the software and $329 for a packaged distribution. In addition, 30-day evaluation licenses are available for those who wish to test the software before making a purchase. VMware offers support for its product through its website in the form of product documentation, a knowledge database, and user forums. Additional support by phone or email is also available, but it may cost a fee to use depending on the situation. In the future, VMware plans to release packages that include virtual machines with pre-installed operating systems. Packages currently planned are Windows 2000 Professional, Windows XP Home Edition and Windows XP Professional (VMware, 2002b).

Alternatives to Vmware

Given that college graduates will enter a workforce that employs an increasingly heterogeneous environment, there is a strong need to provide instruction on a wide range of hardware/software technology. Hardware-based partitioning is common in high-end Unix servers. IBM's Linux-on-the-mainframe strategy also uses this idea. (Dyck, 2002) However, the support costs for providing access to varied technology through high-end servers, mainframes, and multiple processors is obviously prohibitive for many educational organizations.

Programs that emulate many operating systems may be an affordable alternative, and many proven options exist for systems that run Linux. Applications such as Executor, Bochs, SoftWindows, and Wine were developed to provide Windows compatibility on top Linux for those Linux users who needed access to some favorite Windows applications. However, these products are not always efficient and do not provide the complete features of the desired system (Dejanovic, 2000). Another application, Win4Lin, was developed to run Windows 95 or Windows 98 under Linux (Petreley, 2000). If Linux is the underlying OS, Bushong (2002) names Win4Lin the winner over VMWare. Unfortunately, Win4Lin only provides support for a few platforms, which do not include Windows.

The majority of courses that comprise the CPT baccalaureate program use Windows applications. Therefore, the CPT solution required that Windows be the underlying OS in its computer laboratories. When CPT's experiment with virtual machine technology began, the only emulator software available for the Windows operating system was VMware Workstation 2.0. Since that time, other emulator alternatives for Windows have become available. For example, Virtual PC from Connectix was released in the US in June 2001. Connectix was founded in 1988 and has been at the forefront of virtual machine (VM) development from the beginning. The initial product, Virtual used VM technology and was the first implementation of virtual memory for personal computers. As a follow-up to Virtual, Connectix introduced Virtual PC in 1997, which combined the VM technology with a processor emulator to enable Macintosh users to run Microsoft Windows and Windows applications. The recently introduced Virtual PC for Windows has already earned PC Pro's Excellence Award, making it a viable contender for the virtual machine technology market for Windows. (Connectix, 2001)

The CPT Experience

When attempting to add information to a body of knowledge, it is important to determine whether such information is already readily available in addition to learning the most often used methodologies to perform the desired research. Case studies are useful when the object of study is impacted by the context in which it is studied, therefore ideal for a holistic, in-depth investigation (Feagin, Orum & Sjoberg, 1991). Myers (2001) recommends that the case study research method is well-suited to IS research, since the focus is on the study of information systems in organizations. The widely recommended methodology of Yin (1994) provides a detailed, step-by-step process to complete a case study (Myers, 2001; Stewart & Gable, 1999; Tellis, 1997).

The point of the case study design is to help the researcher focus on why the topic is being investigated and what types of results should be expected. Since the interest in developing case study methodologies is relatively new, there is little standardization in the types of case study designs available. Given that there are such a variety of design definitions, the case study researcher should choose the type that best describes the goal of the study, though it would be ideal to choose a design developed by the researcher whose methodology is being used. There is no evidence requiring the categories to be mutually exclusive, and since a study may have a more than one goal, it is feasible that it could be described with more than one category. Multiple-case designs use several cases to solidify confidence in findings. These designs follow replication rather than sampling logic (Yin, 1994). Since these designs are not based on statistics, the generalization of results from either design does not necessarily represent the general population (Yin, 1994, Stake, 1995).

Following the methodology provided by Yin (1994), the first step was to determine the actual questions of the study. For this case, the questions included the following:

- Why was VMware used for CPT 450?
- How was it implemented?
- How did those involved with the system react to its usage? This includes learning about student and instructor experiences using VMware, setup and maintenance issues, and the type of money and space savings that were realized.

Based on the refinement of these questions, the research design was developed, which included the specific types of data collection required. In addition to the research design, the case study design needs were determined. Originally, it was planned to only document the use of VMware during the Spring 2001 semester. Since the implementation of multi-tier programming was not realized, some additional work was done during the Fall 2001 semester. This was regarded as a continuation work from the previous semester, and therefore was treated as a single case. Since the goal was to examine the procedures and explain the impact of the experience of using VMware for the enterprise application development course, this case was defined as explanatory with descriptive elements.

Data collection for this project relied heavily on documentation, archival records, and interviews for sources of data. This included VMware usage instructions, assignment descriptions or course descriptions, and installation and configuration documents used in previous semesters. Interview sources included the course instructor, TCN administrators, and students who have completed the course. Once the data was collected and organized in the case study database, the findings were analyzed. For this study, the data was organized into relevant categories that best represented answers to the study questions.

Spring 2001 Semester

CPT 450 students first used VMware in the Spring 2001 semester when it was installed in one of the department's general purpose computing labs. In order to achieve better performance, the amount of memory for computers in the lab was increased to 512 MB at a cost of approximately ten thousand dollars. Since the size of a virtual machine could easily exceed 1 GB, a Network Attached Storage (NAS) device, which could store approximately 300 GB and cost about fourteen thousand dollars, was purchased to serve as storage for the student virtual machines. Professors supplied the network administration staff with a list of students enrolled in CPT 450, and they then created a secure folder for each student on the NAS device.

To provide communication between host and guest operating systems, the host-only networking option was chosen. Bridged networking was not allowed due to security concern, because bridged networking would give each virtual machine an IP address on the network, which would potentially overwhelm the network. In addition, this option would give students the ability to install malicious software on the virtual machines, which could be used to gain password information or harm other machines on the network. VMware was also configured so virtual machines could not be created using the raw disk option. This option allows VMware to access the local disk/partition as a raw device rather than a file on a file system (VMWare, 2001a). This gives a guest operating system the power to manipulate information on the physical hard drive, which could lead to a compromise in security. To facilitate data transfer between the guest and host machines, the host machine was able to map the guest machine's drive using the Windows *net use* command. This allowed the contents of the virtual machine drive to be viewed from the host machine, but the drives of the host machine could not be seen by the virtual machine. This caused some usability issues when transferring files between machines since it was easy to confuse the host and guest operating systems.

The host-only networking option did generate some limitations in teaching the class. When CPT 450 was held in a stand-alone lab, the instructor would perform an exercise to illustrate how record locking could occur when multiple users tried to update the

same database record at the same time. In previous semesters, each student had a computer on the network and it was possible for each to attempt to update the record on the server. With the use of VMware, each student has his or her own network on the host and other students cannot access machines outside of their own. For this reason, the record locking demonstration was not possible using VMware.

Since various software packages were required for each virtual machine, the network support staff needed to devise a way to deliver the installation files to the users of VMware. Providing physical CDs for each student was a financial and logistical problem, so installation files were provided through the network. Unfortunately, VMware could only read CDs located on the local CD-ROM drive. To work around this, an empty virtual machine disk file was created for each installation CD, and the contents of the installation CDs were copied to these virtual disks. VMware could then be configured to map these virtual disk files as additional disks, which allowed the necessary operating systems and applications to be installed.

Unfortunately, using VMware for CPT 450 did not turn out to be the ideal solution. The network was unexpectedly slow in the lab where VMware was installed, and it often took up to 30 seconds to log into the host computer. Switching between the host and guest operating systems required a similar amount of time. Some of these problems could be attributed to a problem VMware caused with the network. When installed on a computer, VMware installs a network adaptor called *VMware Host Networking Bridge*. This service caused a slowdown in network transmission, leading to slower login and file transfer times. Though it is not known why this occurred, this problem did not appear when the host was upgraded to Windows 2000 for the Fall 2001 semester. In addition, the amount of memory on the host was not sufficient to run more than two virtual machines at the same time. However, even using only two guests on the host simultaneously led to performance poor enough that it was unacceptable for use in the class. Therefore, the virtual machine was only used as a web server for student web assignments, and the use of VMware for creating *n*-tier client server programs was not realized. Due to these problems, the course instructor designed the assignments so it was possible for a student to complete them on a home computer without the use of VMware. The instructor felt it was unfair to mandate use of the VMware lab program because of its slow performance.

After the completion of the Spring 2001 semester, CPT 450 students were asked to respond to questions regarding their experience using VMware. These questions were given to determine whether the problems seen with VMware were indeed problems, and also to determine any particular strengths or weaknesses with using VMware. According to the responses, students had an overall positive experience using this program despite its problems. Most were able to learn the software quickly and felt that an in-class demonstration of the software was the most effective teaching tool. All of the students who responded felt that VMware was an appropriate tool for teaching the course topics, and most were impressed with the idea of running multiple operating systems while maintaining stability in the host machine.

Students also emphasized some of the problems with using VMware. All but one respondent felt that slow performance was an issue. One respondent noticed the problem was aggravated when the virtual machine file was stored on the network, and another complained of the time required to switch between the host and virtual machine. Even so, most felt the slow performance was tolerable. Another problem was computer availability. This was an issue for five of the eight respondents, and those who did not complain either completed work at home or only used VMware during class periods. Both of these problems may have been aggravated if the assignments were written to require the use of Vmware.

Several suggestions were made to help improve the students' experience with VMware. Three of the respondents desired the ability for the virtual machines to connect to the Internet, which was not available due to the use of host-only networking. Another wished for improved graphics support, and another requested that images be made easily available in case one needed to start over. To help alleviate availability problems, one respondent suggested that access to labs running VMware be restricted to those students who actually needed to use the software rather than be advertised as general purpose, open labs.

Students also suggested improving documentation for using VMware. Two suggested creating a document that explained how to run VMware along with important features and functions of VMware. Additional indications of the need for better documentation were that one student was unable to recreate processes demonstrated in class, while another did not realize that information could be shared between the host and guest machines.

Fall 2001 Semester

For reasons previously explained, the use of VMware was reexamined for the Fall 2001 semester. During the summer between the Spring and Fall semesters, new computers were installed in the VMware lab. With a 1.7 GHz processor, 1 GB of memory and

Windows 2000 as the host OS, the new machines promised better performance than the ones used in previous semesters. Since the solutions implemented up to this point had not been satisfactory, emphasis was placed on how to achieve the goal of practical *n*-tier programming.

For the Fall 2001 semester, three virtual machines were created and configured in a similar manner as the Spring and Fall semesters of 2000. The first was running Windows NT Server 4.0 and was configured as a domain controller. A domain named CPT450 was created for use in the class. The second computer was running Windows NT Server 4.0 as well, but was set up as a stand-alone computer that was a member of the CPT450 domain. The final guest, a Windows 98 computer, had already been created the previous semester. This machine was also added to the CPT450 domain. To reduce some of the confusion when switching between machines, each was given a unique desktop background color so they could be easily distinguished from one another.

Each of these machines was also given its own IP address. When VMware was installed, it specified a range of IP addresses that could be used to create a virtual network on the host. Since all the guest machines were in the same domain, they were able to share information through the network neighborhood. Alternatively, the "net use" command could be used to establish a connection between a guest OS and the host or between guest machines.

In order for components running on different machines to communicate, the use of DCOM was required. In the previous semester, this would not work due to problems with permissions between the machines. When all the machines were added to the same domain, however, the use of DCOM was possible. This was a great step toward enabling the teaching of *n*-tier client server development. Now one computer could pose as the database server, another could house the components needed to execute the programs, while another could run as a web server and store a student's web files. Also, an additional virtual machine could act as a client running a Windows user interface for the application or accessing a web interface for the same application.

Graphics performance was questionable when using VMware during the Spring 2001 semester. Often, closing a window within a virtual machine would result in a black square where the window used to be. The screen could be forced to update by sweeping the mouse pointer over the affected area, but the occurrence of this problem hindered usability of the virtual machine. After researching this problem on the VMware site, using the AutoFit option allowed the sizing of the virtual machine window to the same dimensions as the virtual machine desktop. This, along with the increased performance of the machines used in the Fall 2001 semester, appeared to have eliminated this graphics problem.

As a response to the lack of robust instructions regarding the use of VMware, detailed documents were created for the students. This information included accessing the NAS device, starting the virtual machine, and networking between the virtual machine and the host. Students were given this document during class, and were able to follow along during a demonstration of the program. In addition, a handout was created to teach the students how to set up a multi-tier application. During an in-class exercise, students used these instructions to setup and configure a sample *n*-tier client/server architecture. It is hoped that this hands-on experience will solidify the concept of creating multi-tiered applications.

Conclusions and Recommendations

Though its use in CPT 450 is continuing to evolve, it is evident that using VMware for teaching multi-tier client\server application development can be successful in a limited fashion. In general, there are several things that need to be considered before using virtual machine technology in this manner: computer performance, computer availability, and virtual machine networking options. During the Spring 2001 semester, performance was a great hindrance to the use of VMware. With only 512 MB of memory on the host machine, there was not enough memory to distribute among multiple guest computers and still achieve workable performance. When two machines, both with 128 MB memory, were running on the host machine, switching between virtual machines was extremely slow. With the upgraded computers used for Fall 2001, performance issues were much improved. Though there is a decrease in performance when multiple machines start at the same time, it is possible to quickly switch between machines, minimize a machine, or use the host machine without any noticeable performance penalty. It is not yet known if this could be attributed to the great increase in memory and processor speed, the change to Windows 2000, or a combination of both. The performance improvements realized for the Fall 2001 semester make VMware a much more usable solution than in the previous semester. From experience, it is recommended that in order run multiple virtual machines, each virtual machine be configured with at least 256 MB and the host processor have a speed of at least 1 GHz. In addition, a host operating system of Windows 2000 provided greater stability and performance than Windows NT 4.0.

Computer availability is another issue that should be addressed when considering the use of VMware. Since this program requires substantial computer resources, it is only installed on a limited number of lab computers. For the Spring 2001 semester, one lab (24 computers) ran VMware. Teaching schedules monopolized the lab for much of the day, so open lab times were often restricted from 5 p.m. to 11 p.m. To further exacerbate the problem, since this lab had the newest, fastest computers, it was a favorite among students even if they had no use for VMware. This situation often made it difficult for the CPT 450 students to find a free computer. The VMware computers used in Spring 2001 have been installed in another lab, bringing the total to 48 computers running VMware. However, the newer computers installed for Fall 2001 far outperform those used in the previous semester, so students prefer to use the better lab since it does not have the performance issues seen with the slower computers, and the older computers are often left unused. In addition, because of its hardware and operating system requirements, it is not feasible to expect students to have the ability to run VMware from their home computers.

Another consideration when considering this technology is the availability of networking options for the virtual machines. Due to the security reasons explained previously, host-only networking was the configuration used for this class. This type of networking did not allow Internet access within the guest machines, and several students complained about the lack of this option. Though it was not absolutely necessary for this course, access to the Internet would provide a greater level of convenience for the user. With the introduction of the NAT device in VMware 3.0, there is a solution to this complaint.

In general, there is potential in using virtual machine software for education. Though there have been problems with its use, VMware has proven to be a reasonable solution to the problems of teaching a course such as CPT 450. Each student has the ability to configure and use a virtual machine according to his or her needs, without the problems of restricted permissions or the fear of causing problems to the rest of the network. In addition, it is possible to set up a network between the virtual machines in order to demonstrate *n*-tier client/server architectures. Though one student cannot share information with the virtual machine of another, there is still enough functionality available through using VMware to label this software as a compromise over using the space-intensive stand-alone lab.

It is evident that computer performance is the most important aspect when considering the feasibility of implementing VMware as an educational tool. Without adequate processor speed and memory on host machines, using VMware may prove to be more frustrating than it is worth. Next, the number of computers provided for students is significant. It is important to evaluate student-to-computer ratios before implementing this software so that students are not given the additional hardship of competing for an available computer when attempting to complete assignments. Overall, while VMware can offer a solution to many problems, additional problems can be caused by not sufficiently investing in the technology to run this software.

References

Babcock, C. “VMware welcomes guest OSes," *Interactive Week* (7), 1 May 2000, p. 86.

Bar, M. “Oracle8i Server on Linux,” *Byte.com* [online]. 4 October 1999 [cited 17 May 2002]. Available from <http://www.byte.com/documents/s=147/byt19990928s0008/>.

Bushong, C. *Computerworld: Windows Emulators for Linux: VMware, Win4Lin Face Off* [online]. 4 April 2002 [cited 20 May 2002]. Available from <http://www.computerworld.com/softwaretopics/os/story/0,10801,69849,00.html>.

Chowdhry, P. “VMware technology could emerge as a force,” *PC Week* (16), 12 April 1999. p. 14.

Computerworld. *The most wanted skills for 2001* [online]. 4 December 2000 [cited 17 May 2002]. Available from <http://www.computerworld.com/news/2000/story/0,11280,54574,00.html>.

Connectix. *The Technology of Virtual PC: A Connectix White Paper* [online]. 13 September 2001 [cited 20 May 2002]. Available from < http://www.connectix.com/downloadcenter/pdf/vpcw_wp/vpcw_overviewwp_sep1301.pdf >.

Dejanovic, R *Interview with Reza Malakzadeh from Vmware* [online]. . 24 January 2000 [cited 18 February 2001]. Available from <http://linuxticker.com/article/641.html>.

Dyck, T. *Is Simulation as Good as The Real Thing?* [online]. 22 April 2002 [cited 20 May 2002]. Available from <http://www.eweek.com/article/0,3658,s=702&a=25774,00.asp >.

Feagin, J., Orum, A. & Sjoberg, G. (Eds.). *A case for case study,* Chapel Hill, NC: University of North Carolina Press, 1991.

Fonvieille, M. (February 2001). “Using VMware 2”. Retrieved April 29, 2001 from the World Wide Web: http://www.freebsdzine.org/200102/vmware2.php3.

Fuller, J. *The Newbie takes on VMware for Linux* [online]. 2000 [cited 1 May 2001]. Available from <http://www.linuxorbit.com/features/newbie3.php3>.

Hall, M. *VMware 2.0: virtually magnificent* [online]. 10 April 2000 [cited 24 February 2001]. Available from <http://www.linuxplanet.com/linuxplanet/reviews/1704/6/>.

Katz, L. "VMware Workstation 2.0". *Windows Magazine* (7), February 2000. pp. 117 – 119.
Lynch, C. *Review: VMware 2.0.3 for Windows NT, Windows 2000, and Linux* [online]. 9 January 2001 [cited 17 May 2002]. Available from <http://www.8wire.com/articles/?aid=1522>.
Macarlo. *Marimba Deploys VMware Software* [online]. 31 May 2000 [cited 20 May 2002]. Available from <http://www.macarlo.com/vmwaremarimba3105.htm>.
Microsoft Corporation. *Schools make Microsoft Windows the no. 1 platform for education; Microsoft expands education group, announces new learning product* [online]. 11 January 2001 [cited 11 May 2001]. Available from <http://www.microsoft.com/PressPass/press/2001/Jan01/01-11EduPlatPR.asp>.
Moltra, D. "For client/server, think thin," *Network Computing* (10), 28 June 1999. pp. 46 – 54.
Mullen, P. *Duke of URL: Inside VMWare [2.0.2 Review]* [online]. 30 September 2000 [cited 20 May 2002]. Available from <http://linuxtoday.com/news_story.php3?ltsn=2000-09-30-002-21-NW-DT-SW>.
Myers, M. *References on case study research* [online]. 30 January 2001 [cited 24 April 2001]. Available from <http://www.auckland.ac.nz/msis/isworld/case.htm>.
Neih, J. & Con Leonard, O. "Examining Vmware," *Dr. Dobb's Journal: Software Tools for the Professional Programmer* (25), August 2000. pp. 70 – 75.
Petreley, N. "Am I the only one who doesn't know why to use VMware?" *Infoworld* (22), 11 December 2000. p. 60.
Rado. *Emulating a complete PC with Vmware* [online]. 24 January 2000[cited 20 May 2002]. Available from http://linuxtoday.com/news_story.php3?ltsn=2000-01-24-014-04-PS
Railsback, K. "VMware runs Windows under Linux," *Infoworld* (21), 7 June 1999. p. 43.
Randall, N. "One machine, many OSs," *PC Magazine* (19), 18 April 2000. p. 64.
Shankland, S. *VMware wins deal with Microsoft* [online]. 7 September 2001 [cited 17 September 2001]. Available from <http://news.cnet.com/news/0-1003-200-7091332.html?tag=mn_hd>.
Sims, D. *Interview: Mendel Rosemblum of Vmware* [online]. 18 February 2000 [cited 17 May 2002]. Available from <http://www.oreillynet.com/lpt/a/76>.
Stake, R. *The art of case research.* Thousand Oaks, CA: Sage Publications, 1995.
Stewart, G. & Gable, G. *Applying the case study and action research methods to post-graduate studies of enterprise processing system implementations* [online]. 1999 [cited 2 May 2001]. Available from <http://www2.fit.qut.edu.au/InfoSys/ism/Papers/GSt99-6.pdf>.
Tellis, W. "Introduction to case study," *The Qualitative Report* [On-line serial] 3(2), July 1997 [cited 20 May 2002]. Available from <http://www.nova.edu/ssss/QR/QR3-2/tellis1.html>.
VMWare. *Columbia University Students Gain Valuable Experience in Operating System Design Using VMware* [online]. 2000 [cited 20 February 2001]. Available from <http://www.vmware.com/solutions/stories/columbia.html>.
VMware. *Getting started guide: VMware 2.0 for Windows NT and Windows 2000* [online]. 2001a [cited 20 September 2001]. Available from <http://www.vmware.com/pdf/gettingstarted_2_win.pdf>.
VMware. *MultipleWorlds Technology* [online]. 2001b [cited 22 February 2001]. Available from <http://www.vmware.com/company/multipleworlds.html>.
VMware. *Simplifying Computing Infrastructure and Expanding Possibilities* [online]. 2002a [cited 20 May 2002]. Available from <http://www.vmware.com/company/>.
VMware. *VMware Workstation 3.1 FAQs* [online]. 2002b [cited 20 May 2002]. Available from <http://www.vmware.com/products/desktop/ws_faqs.html>.
Yin, R. *Case study research: Design and methods* (2nd ed.), Beverly Hills, CA: Sage Publishing, 1994.