Computer Lab to Go: A "Cloud" Computing Implementation

Marianne C. Murphy
mmurphy@nccu.edu
Computer Information Systems, North Carolina Central University
Durham, NC, 27707, USA

Marty McClelland
mmcclell@nccu.edu
Computer Information Systems, North Carolina Central University
Durham, NC, 27707, USA

ABSTRACT

We describe our experience with a pilot deployment of a virtual computer laboratory at a medium sized public university. The virtual computer laboratory pilot is a technology transfer of the virtual computer laboratory developed by North Carolina State University. The technology provides scalable, high performance computing resources requested through an internet browser and accessed through either a remote desktop connection or ssh client. This paper includes a brief review of virtualization, a review of instructional uses of virtualization, a description of "cloud" computing, the North Carolina State University implementation history, our pilot experience and lessons learned.

Keywords: cloud computing, VCL, virtualization, virtual lab, virtual machine

1. INTRODUCTION

Historically, universities have built computer labs to provide students access to computer technology and university licensed software. Currently, our university Information Technology (IT) division develops an image for the computer labs based on faculty requests for software then deploys the image to the university lab machines. Modifications or additions to the image usually happen once or twice a year with a consequence of a delay in the availability of the software for student use after the faculty request.

Information System educators have demonstrated educational benefits to providing virtual machine images particularly when student users need administrative access to applications (Harvey, Johnson, & Turcheck, 2006; Liegle & Meso, 2007; Powell et al., 2007). This paper describes an implementation of a virtual computing lab (VCL) that provides a library of machine images of different applications and platforms. Users access the VCL through a standard internet browser and select the machine image of the application / platform they desire.

Our experience is that our students embrace the opportunity to access the university computer lab from home or anywhere with a high speed internet connection. This method allows students more flexibility in meeting the demands of their coursework.

Small and medium universities tend to have limited resources but still need to adequately prepare their students with appropriate information technology skills. This requirement mandates the installation and maintenance of adequate computer facilities for the faculty and staff. The minimum requirements of such facilities include the installation of diverse operating systems, up-

grades, patches and new software installs to support instruction. Additionally, a smaller staff must still disseminate licensed software to students and faculty as well as manage authorized access to software and administrative access to systems. These activities can significantly strain a small staff and lead to delays in responding to instructor requests.

Large and well funded colleges and universities can meet these challenges more easily than smaller and secondary institutions. However, all have significant responsibility to their students that will impact these students ability to compete in a global and competitive job market.

Virtual computing offers a considerable solution to these challenges. Building a stable and robust virtual environment is a significant and expensive project. This paper addresses how one medium sized public liberal arts school was able to deploy a virtual computing lab and offers helpful solutions to other educational institutions facing the same computing challenges. This deployment illustrates a cloud computing implementation that extends access to a diverse set of computing resources. In particular, we describe the pilot program, lessons learned and benefits.

2. VIRTUALIZATION

Virtualization is used in many different contexts. In a broad context it refers to the intangible use of computer resources. Virtualization hides many of the physical characteristics of computing from the end user by providing a "virtual image".

The term virtual machine dates back to the 1960's and refers to the experimental IBM M44/44X system (Vaughan-Nichols, 2006). In this type of platform, a host control program provides a simulated computer environment (virtual machine) for its guest software.

The terms virtualization and virtual machine have evolved over the years and acquired various meanings depending on the context. Computers today have enough power to use virtualization to create many smaller "virtual machines." In other words, virtualizations allows for subdivision or sharing of many resources (Barham et al., 2003).

The role of virtualization in education, particularly computer education, is dramatically increasing. Virtualization allows for interactive as opposed to content only delivery (Gaspar et al., 2008). The authors go on to say that the impact on student learning can be significant.

In this paper, we refer specifically to the transparent access of various applications and platforms to students and faculty in every discipline. This access is provided via the Internet in a "cloud" computing implementation. Therefore, access is not limited to a particular classroom or course.

"Cloud" Computing

Cloud Computing is a term used in 2006 by Amazon for their elastic cloud computing service and often used today to describe using the Internet to allow users access to a wide array of technology enabled services (Worthen, 2008). The technological infrastructure and resources are generally owned and controlled by a third party.

Decades of research in virtualization, distributed computing, grid computing, utility computing and more recently networking, web and software services have led to cloud computing (Hayes, 2008). This service provides a reduced information technology overhead for the end-user, greater flexibility, reduced total cost of ownership, economy of scale and on-demand services (Vouk et al., 2008). The principal features of a "cloud" are abstraction and hiding of complexity, use of remote resources and efficient utilization of distributed resources (Vouk, 2007). This framework allows for a reliable aggregation, sharing and allocation of software as well as computational, storage and network resources on-demand.

Virtual Computer Lab

A Virtual Computer Lab (VCL) generally describes an environment used by one group or classroom. Virtual lab environments have been used by various universities for numerous purposes. In most cases, a virtual environment has been deployed for teaching a particular course in either a traditional or distance education classroom.

Telecommunications classes often rely on hands-on experiments and communication hardware and thus have been challenging, particularly for a distant learner. The hard-

ware needed to provide technology skills such as network configuration and security is too costly or simply unavailable but virtualization offers a viable solution at practically no cost (Anisett et al., 2007). Epelbaum proposed a solution whereby students could use industry software to download various procedures to simulate a real experiment and thus providing a student with the means to learn the necessary telecommunication skills (Epelbaum, 2000). Others have built on this proposal by providing a complete virtual machine. In one classroom, computer networking is taught with extensive hands-on experience using an array of Xen virtual machines (Powell et al., 2007).

Virtual computing labs have not been limited to computer networking. Many challenges are encountered teaching web application development where the deployment of a virtual lab is created to meet these challenges (Liegle & Meso, 2007).

Building virtual computer labs for the use in a particular classroom is beneficial but may not be an effective use of the resource. This paper expands the use of a virtual environment from the classroom setting to the university setting. Such deployment provides academic faculty and students with numerous resources and flexibility as well as countless applications.

3. VIRTUAL COMPUTING LAB (VCL) INSTALLATIONS

North Carolina State University VCL Project

The North Carolina State University (NCSU) VCL is a remote access service that allows users the convenience of using a variety of resources via an internet connection. This model was originally described in 2004 by (Averitt et al., 2004). The VCL allows platform-independent access to a large array of computing configurations and applications. The VCL is a joint venture of the College of Engineering Information Technology and Engineering Computer Services group and the High Performance Computing team in the Information Technology Division. Figure 1 is a depiction of the basic infrastructure.

In 2003, NCSU was facing various computing issues on campus. Faculty was frustrated by the amount of time (sometimes up to a year) it took to have necessary software installed. Students were frustrated by hardware and software failures in the campus computer labs. Additionally, the university faced a major funding cut for its supercomputer center forcing a shutdown in 2003 (Young, 2008).

Virtual computing became the logical solution for various reasons. Most significantly, physical labs did not have to be maintained and upgraded by the Information Technology teams. Faculty could install their own software and provide necessary access to their students. Additionally, students were not confined to the time and space issues associated with a physical lab. They could use their own computers without having to purchase necessary software reducing their financial costs and allowing for better time management.

With grants from INTEL, IBM and the University of North Carolina General Administration, NCSU opened their virtual computing lab in 2004. Since that time, other universities and community colleges have followed. George Mason opened their virtual computing lab in 2008 and is primarily built on the NCSU model.

Obviously the cost benefit of such a project is significant. However, even with such savings, smaller universities may not have the funds or expertise to take on such a project. In order to expand this model across the state and beyond, NCSU and IBM instituted the "Virtual Computing Initiative". The goal of this initiative is to create a multi-institutional shared computing services community based upon the VCL model. Our institution became one of the initial participants in this endeavor. The next section describes the first year of the VCL Pilot at our institution.

VCL Project at North Carolina Central University

North Carolina Central University (NCCU) is primarily a liberal arts school with approximately 8,300 students. High performance computing applications are utilized in numerous academic programs and research institutes. Hence the need to provide a diverse computing environment to meet both its research objectives as well as student educational support is necessary.

In 2005, NCCU began to seriously investigate a cost effective methodology that would allow smaller schools with limited resources to deploy a VCL. Faculty and students were frustrated by limitations on the software and hardware available on campus.

Project Goals: The overall goal of the project is to provide a comprehensive computing environment that would meet the computing needs of its students and faculty. The important objectives are:

- Reduce the time of requested software to implementation by providing an alternative to the reliance on the Information Technology division.
- Ease the overflow of the physical computer labs by providing needed applications via an internet connection.
- Make efficient use of licensed software by tying the control of usage to the number of users at any given time rather than the number of physical machines.
- Reduce student costs of purchasing needed software on their personal computers by providing access to these applications in a virtual environment.
- Provide a wide array of applications in a cost effective manner.
- Support applications for classroom use without overburdening the classroom computers.
- Support the growing Distance Education programs at the university.

The deployment needed to be cost effective and have a rapid delivery of needed applications. Additionally, reliance on the expertise and cooperation of IBM and North Carolina State University is vital.

Implementation: A collaborative effort with NCSU resulted in a roll out of the VCL at North Carolina Central University in the fall of 2007 and by spring of 2008, the VCL pilot was in full swing being used by faculty and students on various levels in the School of Business and the School of Library and Information Science. A rollout to the entire university is anticipated for the 2008/2009 academic year. Figure 2 depicts North Carolina Central University's connection to NCSU who provides the schedule software free of charge. North Carolina Central University owns and maintains their server blades.

This deployment at North Carolina Central University has significant importance for other schools facing the same technology issues. The model is a loose confederation of organizations. Currently, this confederation consists of several colleges in the University of North Carolina system and community colleges. Eventually, this network will extend to several K-12 school systems and nonprofit organizations.

Information technology is continually changing, often leaving smaller universities and poorer public school systems continually behind. Providing students at these institutions with the technology skills they need to secure future employment is challenging. Projects like VCL that can be modeled to extend to these institutions will provide significant tools to lessen this gap.

Initial Experiences/Results: The one year pilot program consisted of 100 unique users. These users utilized over 25 images across 24 server blades. The primary users were from Computer Information Systems, Decision Science, Marketing, Finance and Hospitality and Tourism. The fall semester usage was minimal; however, by January and February, activity had dramatically increased. In this two month period a total of 287 reservations were made for a total usage time of 431 hours (Seay & Tucker, 2007). The two most popular programs were SAS and SPSS. Additionally, requests for software installations increased. As needed applications became available, reservations increased.

Informal discussions were held with six professors in order to determine their experiences as well as their students. Table 1 is a list of the basic questions. Overall, everyone agreed that having the VCL is necessary and their experiences were positive. Additionally, all agreed that they would continue to use the VCL and encourage their students' usage.

Table 1: Informal Survey of Users

Questions Asked	Responses
What worked?	Availability and Software Access (see text)
What didn't work?	Identification of images and other (see text)
What to do differently?	Better images, training (see text)
Would you continue to use?	6 Yes, 0 No
Did you use it for class?	4 Yes, 2 No
Did you use it for research?	4 Yes, 2 No
What were the students' experiences?	Mixed (See text)
Did it affect student performances?	2 Yes, 2 No, 2 N/A

The VCL was used both for teaching and research. Although resources were not allocated between the two, the heaviest usage was research particularly with programs such as SAS and SPSS. In class usage was generally positive but outside of class students experienced issues that discouraged their usage. Graduate students and upper class students generally had fewer problems than students in lower level computer classes. Most of the students in the spring semester were not required to use the VCL out of class. Students were encouraged to use the VCL in place of going to the computer lab for such programs as SAS and Microsoft Office 2007 and other programs that they did not have access to on their personal computers.

Problems and Issues: The primary problem students had was the wait time for the image to be available for use. Some students had difficulty locating the appropriate image for the application they wanted to use. While the SAS image was easily identified since SAS was the image name, other images had vague labels that did not readily identify the applications included in that image. The procedure is to sign up for a reservation to a particular image. A drop down link is made available from which a user selects an image. The applications available on the image were not originally clear and therefore users could not find the application needed. Modifications were made so that the image indicated the applications that it contained (See figure 3).

Reservations were fulfilled in a very short time frame, students would log on only to find they made the reservation for the wrong image and therefore they had to make a new reservation. The students could not easily navigate in order to find the applications they needed. Professors experienced some of the same issues but were more easily able to find the applications needed and successfully use the VCL.

The primary issues in this pilot year are:

- The reservation system was difficult to navigate and frustrating since reservations had to be made to a particular image.
- The image names were not intuitive and therefore the user could not easily determine which image contained the application needed.
- Wait times were not generally more than 10 to 15 minutes for the initial reservation but wait times for the image load would sometimes take beyond 30 minutes.
- Images are loaded on demand as opposed to preloaded. The pros and cons of such a system need to be carefully assessed in order to extract the maximum benefit for users.
- Since this was a pilot program, the university help desk did not provide support and therefore students had limited options for seeking assistance.

These issues primarily affected student usage since students did not know where to direct questions in order to get answers quickly and therefore used alternative solutions to using a particular application.

Primary benefits:

- Software can be easily deployed and provide access to avoid delays for IT to physically configure labs, classrooms
- Simplify license management, for instance the university may only purchase

50 licenses that can allow 50 simultaneous users, rather than installing this application on 50 local computers, it can be installed on a blade allowing more than 50 users to have access to the application and controlling the reservations to only allow a maximum of 50 users at any one time

- Computers that access VCL can be low end / older machines since the high performance computing is actually performed on the server. Meanwhile, the server is physically secured. Thus, we do not need labs with high performance work stations and do not have the issues of securing high performance work stations.

Additional instructional benefits:

- Instructors can quickly install software needed for class – no delays waiting for IT to deploy or conflicts with software currently installed in the lab.
- Platform of virtual image can be different from platform of client machine – for example, a virtual image of a linux server can be delivered to a Windows desktop.
- Students can access the software from anywhere (home, university lab and laptop) with an internet connection. This feature is particularly beneficial to students who do not have the application on their personal computer and cannot easily access the university computer lab. Additionally, this is particularly significant in offering nontraditional classroom instruction such as distance education and online instruction.

Future Expectation: All agreed that the benefits of the VCL far outweighed any frustration and that the first year pilot was successful. Virtualization offers many benefits beyond the use in a particular classroom. It primarily is a computer lab that faculty and students can take with them.

Overall the pilot year was extremely successful. However, as with all "firsts", improvements for ongoing deployment and other implementations can be made. The primary lessons learned in this implementation include:

- The use of a campus authentication server with local system administrators would provide better access.
- The reservation system provides usage control but can be limiting to the user. Standard preloaded images available to users may offer more benefit since these preloaded images are available on demand.
- Image content should be clearly organized and identified. Reservations are made to the image and it was difficult for users to identify what image contained the necessary application. Modifications were made to image names (See figure 3).
- The images and content needs to be balanced. In this first year, each time a request for an application on the VCL was made, a new image was added and often the content of the image was unclear. A solution is to determine common content for a particular image perhaps by subject or user group and reduce the number of images. Not only would this method be more clear to the user but would also simplify the preloading / loading of these images.
- Resources were not allocated between research and teaching. The plan is to monitor usage over the next year and determine how these resources should be allocated.
- Broadband access is generally preferred though as little as 128kbs is sufficient for many applications.

The future success of this implementation will be determined by the users. The second year of the project is focused on increasing usage within the school of business by supporting faculty and student requests. In subsequent years the focus will be a roll out to the entire university.

4. CONCLUSIONS

The benefits of virtualization in the classroom have been clearly defined and understood. However, the benefits beyond the classroom have been minimally discussed. Clearly the benefits as a "computer lab on the go" for students and faculty are significant and usages of such VCL systems are growing. Although extremely beneficial, the

cost of an effective virtual computer lab can be enormous. These costs can be significantly reduced and shared through the use of technology transfer. In this paper, we describe a VCL that is used beyond individual classrooms. In particular, we describe the implementation and its first pilot year at a medium size public university.

Virtualization is an effective solution for colleges and universities to meet the changing demand of information technology as well as the changing academic environment. Undertaking any technology change or improvement can stretch the financial resources of even large institutions. In this paper, we offer a solution to smaller universities trying to keep pace and offer students the most advanced technologies in order to adequately prepare him or her for the global workplace. We encourage any university frustrated by the length of time to implement software and/or in need of sharing resources to investigate the use a cloud computing implementation. This type of implementation is cost effective and provides efficient use of resources.

5. ACKNOWLEDGEMENTS

NCCU's pilot VCL would not have been possible without the support of IBM and NCSU. Numerous IBM and NCSU employees generously shared their time and expertise with NCCU. In particular, Andy Rindos of IBM has encouraged and supported NCCU in the development of the VCL. Equipment for the VCL was purchased with two IBM Shared University Research awards. Sam Averitt of NCSU had the vision of the benefits to other institutions of the technology transfer of the NCSU developed VCL and has freely shared the expertise of his staff and the technologies developed by NCSU. And none of this would have been possible at NCCU without Dr. Cameron Seay who has championed the NCCU virtualization and high performance computing initiatives.

6. REFERENCES

Anisett, M., Bellandi, V., Colombo, A., Cremonini, M., Damaini, E., Frati, F., et al. (2007). "Learning Computer Networking on Open Paravitual Laboratories." *IEEE Transactions on Education, 50*(4), 302-311.

Averitt, S., Vouk, M., Schaffer, H., Sills, E., & Howell, G. (2004). "Support for Distributed Scalable High Performance and Grid-based Computing and Problem Solving with Emphasis in Science and Engineering." North Carolina State University: UNCG.

Barham, P., Dragovic, B., Fraser, K., Hand, S., Harris, T., Ho, A., et al. (2003). "Xen and the Art of Virtualization." *ACM SIGOPS, 37*(5).

Epelbaum, S. S. (2000). "Virtual Lab Experiments in Telecommunications for Distance Learning*".* Paper presented at the ISECON.

Gaspar, A., Langevin, S., Armitage, W., Sekar, R., & Daniels, T. (2008). "The Role of Virtualization in Computing Education." *ACM SIGCSE Bulletin, 40*(1).

Harvey, V. J., Johnson, R. S., & Turcheck, J. C. (2006). "Virtual Laboratory Intrusion Detection Experience for Information Systems Professionals." Paper presented at the ISECON 2006.

Hayes, B. (2008). "Cloud Computing." *ACM, 50*(7), 9-10.

Liegle, J. O., & Meso, P. N. (2007). "Evaluation of a Virtual Lab Environment for Teaching Web Application Development." *Information Systems Education Journal, 5*(7).

Powell, V. J., Turcheck, J. C., Wu, P. Y., Franzi, L. C., Johnson, R. S., Parker, I. W., et al. (2007). "VLabNet: A Virtual Laboratory Environment for Teaching Networking and Data Communications." Paper presented at the ISECON.

Seay, C., & Tucker, G. (2007). "The Virtual Computing Initiative at a Small Public University." North Carolina Central University.

Vaughan-Nichols, S. J. (2006). "New Approach to Virtualization Is a Lightweight." *Computer-IEEE Computer Society*(November), 12-14.

Vouk, M. (2007). "Cloud Computing." North Carolina State University.

Vouk, M., Averitt, S., Bugaev, M., Kurth, A., Peeler, A., Rindos, A., et al. (2008). *Powered by VCL" "*Using Virtual Computing Laboratory (VCL) Technology to power Cloud Computing." Paper presented at the ICVCI.

Worthen, B. (2008, September 23, 2008). "Overuse Clouds Buzz Term's Meaning*." Wall Street Journal,* p. B8, from http://blogs.wsj.com/biztech/2008/09/23/overuse-clouds-buzz-terms-meaning/

Young, J. R. (2008). "A Computer Lab that students use but never see." *The Chronicle of Higher Education, May, 2008*.

APPENDIX A: Figures

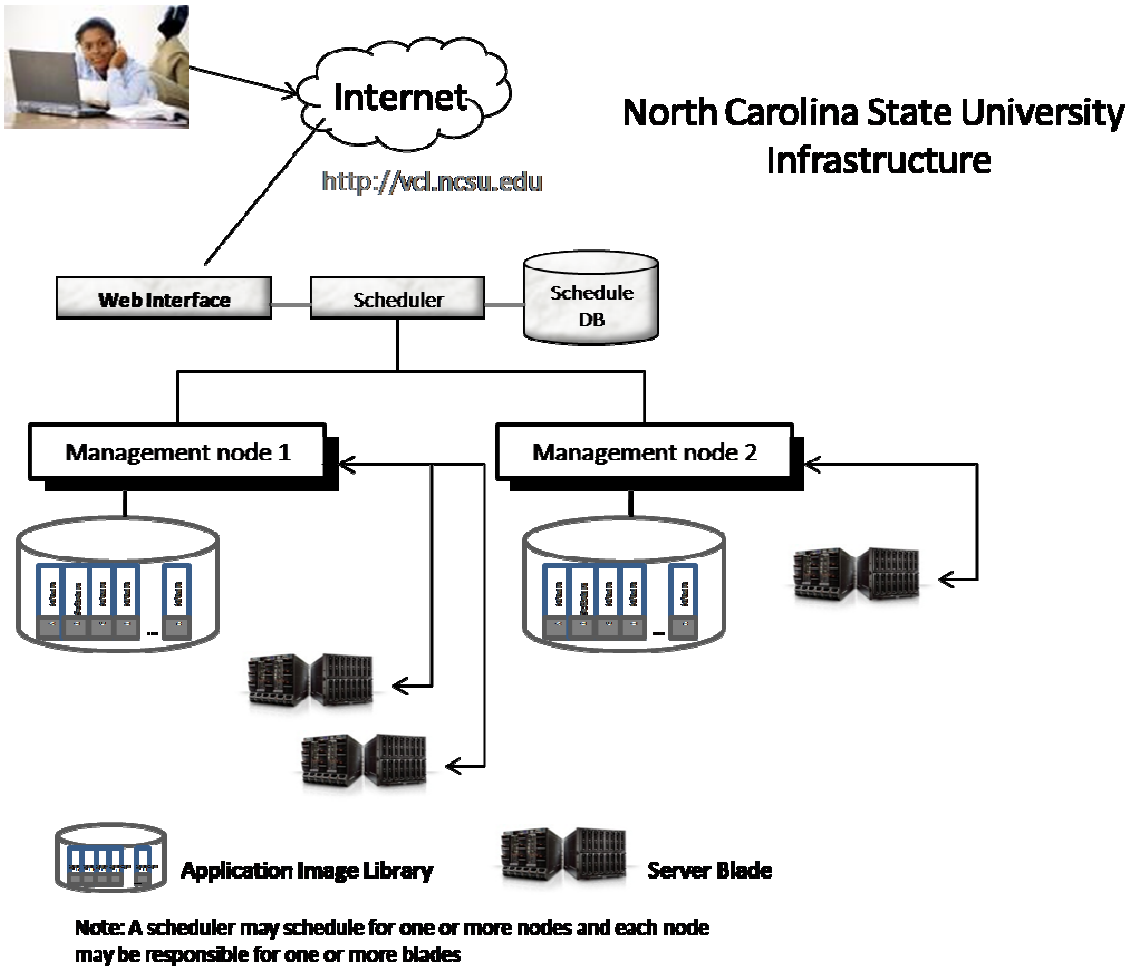

Figure 1: Basic Architecture of the NCSU VCL

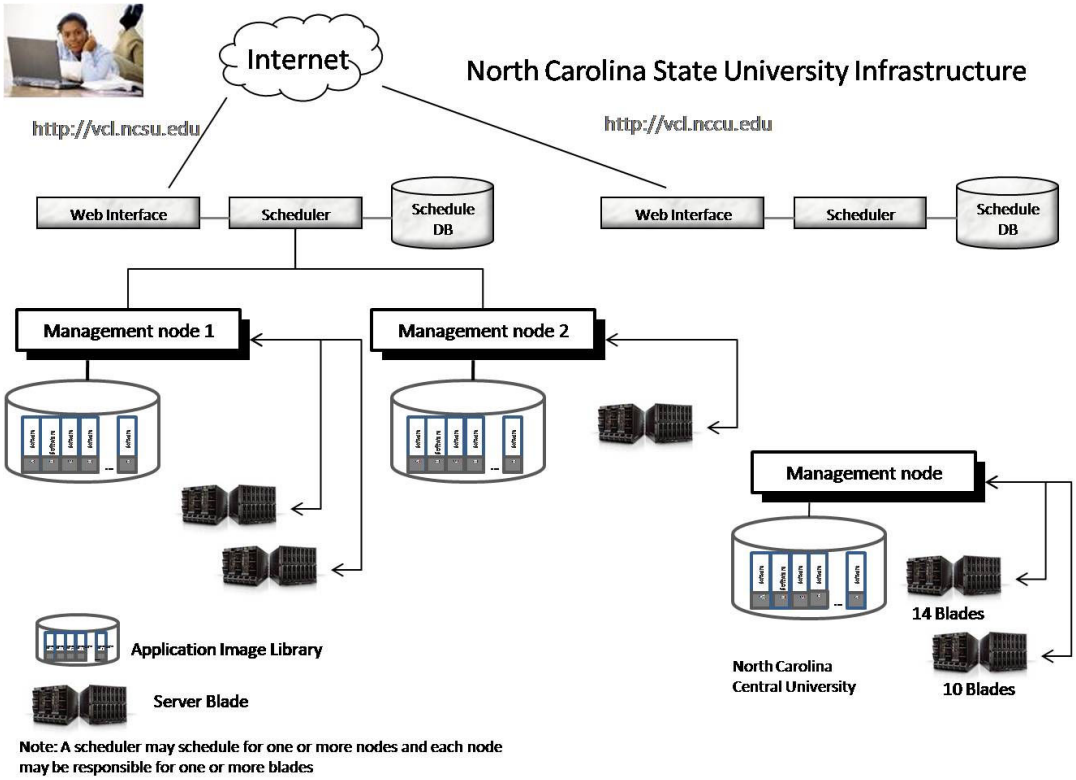

Figure 2: VCL implementation at North Carolina Central University

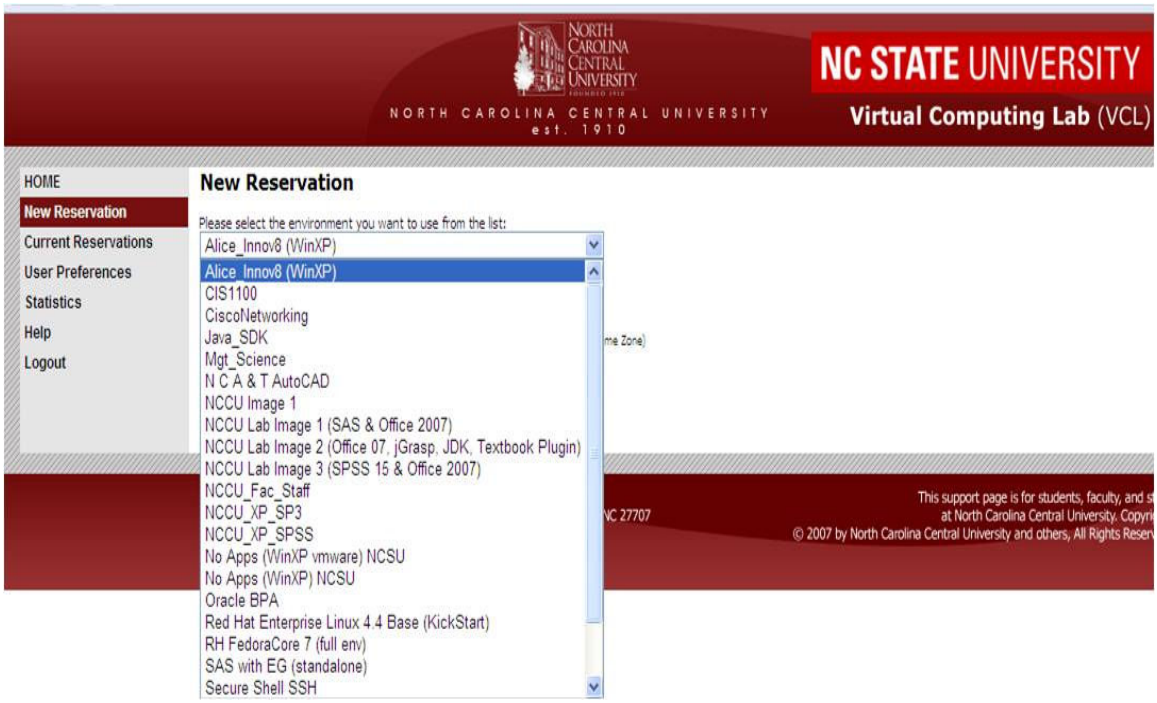

Figure 3: The reservation system