



**NAVAL
POSTGRADUATE
SCHOOL**

MONTEREY, CALIFORNIA

THESIS

**THE NEXT GENERATION OF LAB AND CLASSROOM
COMPUTING—THE SILVER LINING**

by

Milan Vukceвич

December 2016

Thesis Advisor:
Second Reader

Glenn Cook
Albert Barreto

Approved for public release. Distribution is unlimited.

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE December 2016	3. REPORT TYPE AND DATES COVERED Master's thesis	
4. TITLE AND SUBTITLE THE NEXT GENERATION OF LAB AND CLASSROOM COMPUTING— THE SILVER LINING			5. FUNDING NUMBERS	
6. AUTHOR(S) Milan Vukceovich				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. IRB Protocol number <u> N/A </u> .				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release. Distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) The models for providing computing services have changed over the recent years, thereby allowing many lab-computing options. However, currently at the Naval Postgraduate School (NPS), thick client computers provide computing services to students in the learning and resource centers (LRC). Due to various budget restrictions, decreased manpower, and the directives set forth by the Navy, the existing LRC solution at NPS has become increasing difficult to maintain efficiently. These reasons allow the following research questions to be asked: what viable options are available to provide the same level of capabilities in the LRCs at NPS and for the viable options, does a cost comparison show which solution would be preferable for the labs and classrooms at NPS? To answer the questions, the NPS Clouddlab virtual desktop infrastructure (VDI) solution, as well as the computing solutions at three universities, was selected as the basis for comparison. The research method used includes a qualitative methods approach that utilizes case studies to perform the analysis. Analysis performed in the research looks to find the most effective solution in terms of cost, manpower, and availability of the systems in question. The results showed, in terms of hardware and software costs and manpower that Stanford University has the most cost effective solution provided by their private cloud solution. In regards to availability, the greatest system availability was at the Naval War College, California State University, Monterey Bay, and Stanford University. All three solutions were available 99.9% of the time. The recommendations made were to implement a private cloud computing solution similar to the technology used at Stanford University, implement a bring your own device (BYOD) policy at NPS, and to expand the NPS Clouddlab solution both in terms of licensing and into the LRCs using BYOD and thin clients.				
14. SUBJECT TERMS lab, classroom, thick client computer, LRC, Clouddlab, virtual desktop infrastructure, VDI, hardware cost, software cost, manpower, availability, cloud computing, private cloud, bring your own device, BYOD, thin client, software as a service, SaaS, virtualization, personal computer, information technology, IT, client-server architecture			15. NUMBER OF PAGES 129	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU	

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release. Distribution is unlimited.

**THE NEXT GENERATION OF LAB AND CLASSROOM COMPUTING—THE
SILVER LINING**

Milan Vukcevic
Civilian, Department of the Navy
B.S., San Jose State University, 1994

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN INFORMATION TECHNOLOGY MANAGEMENT

from the

**NAVAL POSTGRADUATE SCHOOL
December 2016**

Approved by: Glenn Cook
Thesis Advisor

Albert Barreto
Second Reader

Dan Boger
Chair, Department of Information Sciences

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

The models for providing computing services have changed over the recent years, thereby allowing many lab-computing options. However, currently at the Naval Postgraduate School (NPS), thick client computers provide computing services to students in the learning and resource centers (LRC). Due to various budget restrictions, decreased manpower, and the directives set forth by the Navy, the existing LRC solution at NPS has become increasingly difficult to maintain efficiently. These reasons allow the following research questions to be asked: what viable options are available to provide the same level of capabilities in the LRCs at NPS and for the viable options, does a cost comparison show which solution would be preferable for the labs and classrooms at NPS? To answer the questions, the NPS Cloudlab virtual desktop infrastructure (VDI) solution, as well as the computing solutions at three universities, was selected as the basis for comparison. The research method used includes a qualitative methods approach that utilizes case studies to perform the analysis. Analysis performed in the research looks to find the most effective solution in terms of cost, manpower, and availability of the systems in question. The results showed, in terms of hardware and software costs and manpower that Stanford University has the most cost effective solution provided by their private cloud solution. In regards to availability, the greatest system availability was at the Naval War College, California State University, Monterey Bay, and Stanford University. All three solutions were available 99.9% of the time. The recommendations made were to implement a private cloud computing solution similar to the technology used at Stanford University, implement a bring your own device (BYOD) policy at NPS, and to expand the NPS Cloudlab solution both in terms of licensing and into the LRCs using BYOD and thin clients.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
A.	LAB AND CLASSROOM COMPUTING.....	1
B.	PROBLEM STATEMENT	1
C.	RESEARCH QUESTIONS	2
D.	BENEFITS.....	2
E.	RESEARCH DESIGN AND METHODOLOGY	3
F.	THESIS ORGANIZATION.....	4
II.	LAB AND CLASSROOM TECHNOLOGIES (BACKGROUND AND DEFINITIONS).....	7
A.	COMPUTING BACKGROUND.....	7
B.	THE PC REVOLUTION	10
C.	OVERVIEW CLIENT-SERVER ARCHITECTURE	13
D.	OVERVIEW OF THICK CLIENT COMPUTING	15
E.	OVERVIEW OF CLOUD COMPUTING.....	17
F.	ESSENTIAL CHARACTERISTICS OF CLOUD COMPUTING	19
G.	CLOUD COMPUTING SERVICE MODELS.....	20
1.	Software as a Service	21
2.	Platform as a Service	22
3.	Infrastructure as a Service.....	22
4.	Desktop as a Service.....	23
H.	CLOUD COMPUTING DEPLOYMENT MODELS.....	24
I.	VIRTUALIZATION.....	25
J.	VIRTUAL DESKTOP INFRASTRUCTURE.....	27
K.	BRING YOUR OWN DEVICE	29
L.	THIN CLIENT	30
M.	CONTAINERS.....	31
III.	LAB TECHNOLOGIES AT THE NAVAL POSTGRADUATE SCHOOL AND SELECT UNIVERSITIES	37
A.	LRC COMPUTING AT NPS.....	38
1.	Customers and Services Provided	40
2.	LRC Hardware and Software Cost, Manpower, and Availability.....	42
B.	VDI COMPUTING AT NPS.....	48
C.	VDI INTERVIEW QUESTIONS AT NPS.....	50

D.	LAB COMPUTING AT THE NWC, CSUMB, AND STANFORD.....	52
E.	INTERVIEW QUESTIONS AT NWC, CSUMB, AND STANFORD.....	55
IV.	ANALYSIS OF VDI TECHNOLOGIES AT NPS AND IT TECHNOLOGIES AT THE NWC, CSUMB, AND STANFORD UNIVERSITY	59
A.	VDI COMPUTING AT NPS—REASONS FOR MIGRATING	59
B.	DISTANCE LEARNING	60
C.	EXPANDING CURRICULA REQUIREMENTS.....	62
D.	PRACTICAL LIMITATIONS	63
E.	SERVICES, COST, MANPOWER, AND AVAILABILITY OF VDI	65
F.	LAB COMPUTING AT THE NWC, CSUMB, AND STANFORD.....	71
1.	NPS.....	86
2.	NPS Cloudlab	88
3.	NWC.....	90
4.	CSUMB	92
5.	Stanford	94
V.	CONCLUSIONS AND RECOMMENDATIONS.....	97
A.	ANALYSIS	97
B.	CONCLUSIONS	100
C.	RECOMMENDATIONS.....	104
	LIST OF REFERENCES.....	107
	INITIAL DISTRIBUTION LIST	111

LIST OF FIGURES

Figure 1.	The ENIAC Computer. Source: United States Army (n.d.).	8
Figure 2.	Basic Client Server Network. Source: Dean (2012, p. 6).	14
Figure 3.	The OSI Model. Source: Rivero (2015).	15
Figure 4.	Simple Network Diagram Depicting a Cloud as the Internet. Source: Lindsey (2007).	18
Figure 5.	Cloud Computing Service Models. Source: “Cloud Services” (2015).	21
Figure 6.	Graphic Depiction of Structure of Virtual Machine from Virtualization Essentials. Adapted from VMware (n.d., p. 4).	26
Figure 7.	Representation of VDI Model versus Client Server Model. Source: Oglesby (2006, What Is VDI? section, para. 1).	27
Figure 8.	Example of a Dell Wyse Thin Client Computer. Source: “Dell Wyse D90D7 Thin Client” (n.d.).	30
Figure 9.	Docker Container Representation. Source: “What Is Docker?” (n.d., Package Your Application into a Standardized Unit for Software Development section, para. 1).	32
Figure 10.	Architectural Approach for Virtual Machines. Source: (“What Is Docker,” (n.d., Virtual Machines section, para. 3).	33
Figure 11.	Architectural Approach for Docker Containers. Source: “What is docker?” (n.d., Comparing Containers and Virtual Machines section, para. 2).	34

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

Table 1. University IT Solutions Compared Against Operational Criteria and IT Cost (Hardware, Software, and Manpower).83

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

ARPANET	Advanced Research Projects Agency Network
AWS	Amazon Web Services
BYOD	bring your own device
CAC	common access card
CP/M	control program/ monitor
CPU	central processing unit
CRM	customer relationship management
CSUMB	California State University Monterey Bay
DaaS	Desktop as a Service
DDCIO (N)	Deputy Department of the Navy Chief Information Officer
DL	distance learning
DOD	Department of Defense
ENIAC	Electronic Numerical Integrator and Computer
ESS	Essential Stanford Software
Ghz	gigahertz
GPU	graphical processing unit
GSBPP	Graduate School of Business and Public Policy
GSEAS	Graduate School of Engineering and Applied Sciences
GSOIS	Graduate School of Operational and Information Sciences
http	hypertext transfer protocol
I/O	input/output
IaaS	Infrastructure as a Service
IC	integrated circuits
ICA	independent computing architecture
IMSAI	IMS Associates Inc.
IT	information technology
ITACS	Information Technology and Communication Services
LAN	local area network

LRC	learning and resource center
MITS	Micro Instrumentation Telemetry Systems
NIC	network interface card
NIST	National Institute of Standards and Technology
NOC	network operations center
NWC	Naval War College
OPM	Office of Personnel Management
OSI	open systems interconnection
PaaS	Platform as a Service
PC	personal computer
PCoIP	PC over IP
PE	Popular Electronics
QDOS	86-DOS
RC	research computing
RDP	remote desktop protocol
SaaS	Software as a Service
SCP	Seattle Computer Products
SIGS	School of International Graduate Studies
TAC	Technology Assistance Center
TCO	telecommunications office
TCP/IP	transmission control protocol/Internet protocol
URL	uniform resource locator
USB	universal serial bus
VDI	virtual desktop infrastructure
VM	virtual machine
VNC	virtual network computing
WAN	wide-area-network

ACKNOWLEDGMENTS

First and foremost, I would like to thank Almighty God. My prayers have been answered as confirmed by the diploma in the frame. To my parents, thank you for providing the foundation that I relied upon and continue to rely upon every day. Your love, support, and encouragement has made the years of late evening and weekend schoolwork endurable. To my sisters and their families, Sonja, Caesar, Caesar Jovan, Alexander, and Danica, Jeff, Ryan, and Julia: thank you for your love and support. You always understood the commitment involved and the demands I faced while pursuing this degree working full time at the Naval Postgraduate School. I hope I have made you all proud.

Thank you to the faculty and staff at the Naval Postgraduate School for participating in the interviews. Also, I would like to thank the IT professionals at the Naval War College, Henry Simpson at CSUMB, and Addis O'Connor at Stanford University. Your time and dedication is appreciated.

To my colleagues in the ITACS department, thank you for all the encouragement and support that you provided. You are true professionals in every way. Thank you, ITACS and MAE management, for providing me with the opportunity to further my education. A special thanks to Malcolm Mejia, Toan Tran, Hedwig Fernando, and Edgar Mendoza. You motivated me during intricate times.

I would also like to thank Professors Glenn Cook and Albert "Buddy" Barreto for their guidance and patience throughout this process. Your mentorship and expertise were invaluable. Thank you, Chloe Woida, for your instruction. You made it fun.

Thank you all very much for supporting me through my journey.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

A. LAB AND CLASSROOM COMPUTING

Computers used in labs and classrooms have been employed to enhance or supplement the delivery of a traditional educational program. As computers have become less expensive and more accessible, the demand for computers has grown, thus allowing for their widespread use within labs and classrooms. We as a society have become highly dependent on technology to aid us in our work, studies, and research.

The methods and technology used to provide computing services have changed over the years. Unless otherwise specified, information technology (IT) related information pertaining to the learning and resource centers (LRCs) at the Naval Postgraduate School (NPS), presented in the thesis, is based on the author's personal experience as an IT Specialist working in the Information Technology and Communication Services (ITACS) department for over a decade. Currently at NPS, thick client computers deliver computing and software applications to students so that they can work on their homework, lab projects, thesis assignments, and dissertations. The lab support model that delivers the latest hardware and software to students in the labs and classrooms has been in use for over two decades. Due to the end strength issue, budget reductions (sequestration), unsupported software, and the directives set forth by the Deputy Department of the Navy Chief Information Officer (DDCIO (N)), the existing LRC solution at NPS has become increasing more difficult to maintain in an efficient manner. These hurdles warrant a look at implementing newer technologies in the LRCs that may potentially be used in the near future. The focus of this research is to advance a more efficient, streamlined, cost effective computing model and methods that may enhance the support of the LRCs.

B. PROBLEM STATEMENT

NPS is currently under a hiring freeze, has budget restrictions, and is using software that does not include the latest drivers for newer network interface cards (NICs). The existing LRCs, also known as labs, IT solution at NPS has become increasingly

difficult to maintain in an efficient manner. Manpower in the LRC support group is spread thin, the equipment in the computer labs is being updated less frequently due to budget issues, and the latest device drivers that support newer hardware are not included in one of the imaging software used in the support process. These reasons, coupled with the directive set forth by the DDCIO (N) requiring Echelon II commands to consolidate data centers by 25% and increase server utilization by at least 40%, as well as the goal of exploring viable measures to increase IT efficiency, have created a path to research alternative methods for lab computing.

C. RESEARCH QUESTIONS

The following research questions are addressed by the research done in the thesis. The questions pertain to the technology currently used in NPS's LRCs and to viable options available, as well as a cost comparison between the options.

- What is the current technology used for providing computing resources in the labs and classrooms at NPS? Is this technology still a viable solution?
- What viable options are available to provide the same level of computing capabilities in the LCRs at NPS?
- For the viable options, does a cost comparison show which solution would be preferable for the current labs and classrooms at NPS?

D. BENEFITS

Several benefits could be gained from the results of this research. The IT department spends a great deal of effort, budget, and labor hours in maintaining the labs and classrooms at NPS. The services are provided with the purpose of delivering the latest hardware and software to students, faculty, and staff to aid them with their curricula, research, and support of NPS. This research could benefit the school and the Navy, as this research looks to reduce costs for both hardware purchases and tools associated with deploying images to the lab computers. Another potential benefit would be to find effective methods of IT implementation to cut down on labor hours that the IT staff spends in maintaining the lab and classroom environment. As with most computing environments, downtime is inherent. However, several enterprises are touting 98% availability if their solution is implemented at an organization. Although downtime may

not be eliminated through the implementation of these solutions, it may be reduced. Students, faculty, and staff may potentially be more productive in their course work and research, as a few of the possible computing solutions are available from any smart device, from any location and at any time.

E. RESEARCH DESIGN AND METHODOLOGY

The research design used in this thesis includes a qualitative method approach that utilizes questions asked of several IT professionals in an interview setting or by conference call to gain a perspective about the technologies in use and motives behind their use. A set of questions were presented to faculty and staff at NPS to find what services were being provided and to which contingent of customers. The questions also aim to find the cost, the manpower associated with the technology, as well as the performance of the technology measured as availability.

The qualitative method is used when presenting the various IT solutions implemented by different universities for providing computing services to students. A similar set of interview questions were asked of the IT professional working at other universities with the goal of presenting case studies for each university. A quantitative methodology aids in gathering data and calculating values so that the various IT solutions can be compared. The main focus is the services provided, cost, manpower, and availability associated with current thick-client based solution used at NPS compared to the services provided, cost, manpower, and availability of the IT solutions used at the Naval College, California State University Monterey Bay (CSUMB), and Stanford University. The technologies used at those schools are considered for implementation at NPS. Other research methods consider services provided to the students, faculty, and staff at the university. The cost is based on the average price of the hardware and software being used to provide the services. The cost of manpower is shown as labor hours multiplied by cost of labor per hour. The system availability is shown as a percentage, which correlates to the computing solution and the amount of time it is available for use by students.

F. THESIS ORGANIZATION

Chapter II reviews the background information and literature review for the research. Several sections review the history of the digital computer and their evolution over time. Within those sections, the reader is introduced to the thick client computer and the client server architecture. These concepts are currently used in labs at NPS. Several sections discuss cloud computing and its various components. Included in those sections are examples showing the implementation of cloud technologies across various organizations. Two sections also review a component of cloud computing in use at NPS for teaching purposes, as well as to provide computing services to distance learning (DL) students.

Chapter III uses the concepts and knowledge provided in Chapter II as a starting point for introducing the current lab technologies used in the labs at NPS. IT professionals from the ITACS department provided information pertaining to the cost, manpower, and availability of the lab solution used at NPS. The virtual desktop infrastructure (VDI) solution in use at NPS is discussed. Four VDI use cases are presented that provide a foundation as to why the technology was selected for use at NPS. A series of questions are presented in the chapter to be presented to NPS IT professionals whose answers will help to deliver an estimate for the cost, manpower, and availability associated with maintaining the VDI solutions. The chapter is supplemented with various IT computing technology examples used at the Naval War College (Newport, Rhode Island), CSUMB, and Stanford University to provide computing services to students. A second series of questions are presented to IT professionals at these universities and their replies can help to show services provided, an estimated cost of the technology, manpower involved, and an estimate about the technology's availability.

Chapter IV contains the answers to the questions listed in Chapter III. The information provided by the four VDI administrators and the three IT professionals at NWC, CSUMB, and Stanford University are used to calculate a set of values to be categorized for comparison purposes.

Chapter V includes the analyses of the research. Data was gathered by interviewing various IT professionals at NPS, the Naval War College, CSUMB, and Stanford University. The chapter also includes the conclusions made based on the analyses, as well as recommendations. The recommendations include viewpoints made, based on the data presented in the case study, to advance a more efficient, streamlined, cost effective computing model in support of the NPS LRCs, as well as other learning institutions.

THIS PAGE INTENTIONALLY LEFT BLANK

II. LAB AND CLASSROOM TECHNOLOGIES (BACKGROUND AND DEFINITIONS)

This chapter provides the background needed to analyze the current technologies used in lab and classrooms environments as a basis for assessing which technology would be a good fit for implementation at universities looking to upgrade their infrastructure to save money, use resources more efficiently, and provide scalability. The research reviews the first modern day digital computer. Next, this chapter investigates the next generation of computers known as mainframe computers. The thick/fat client computer will be defined and shown how it is used in the client-server architecture used since the 1980s. This research also defines cloud computing and its various components. The term “cloud” refers to a remote location where this hardware and software are located. Exactly where the software and hardware is located is not important to the end user, only that the services remain available on-demand. For the research, the remote location is important, as the research analyzes whether these solutions can fulfill the goals for improved processes and services coupled with cost savings.

During the analysis of the different options available for lab and classroom computing, two specific cloud computing architecture service models are discussed. The first service model analyzed is the VDI model. This model may prove to be invaluable in providing the means to deliver software to NPS customers over the Internet and at a lower cost than the current solution. The private cloud service model is also evaluated. Although a newer concept, it may be a good alternative for NPS since a private cloud model offers many advantages.

A. COMPUTING BACKGROUND

This section explores the beginnings of the modern day computer. The computer, known as ENIAC (Electronic Numerical Integrator and Computer), opened the door to digital computing, which eventually led to today’s computer industry (Yost, 2005). The military and universities mainly used the first digital computers to perform number crunching and to solve mathematical type problems. It took several operators to program

computers by flipping switches and setting wiring similar to manually switched telephone boards. It was a centralized computer, which took a team to run it. As computer technology advanced, and the usefulness of the computer became evident, more people became interested in using it to solve complex problems, which would take far too long or be impossible to solve without its use. The computers that evolved from the first modern day digital computers were mainframes. The research shows how the mainframe computer evolved from those early beginnings and proved to be invaluable in the roles they played in several concepts (such as virtual machines and dumb terminals also known as thin clients) used in modern day Cloud service models.

Mechanical and electrical computers have been in existence since the 19th century. However, the modern day computer was born during the late 1930s and early 1940s. The ENIAC (see Figure 1) is credited as being the first general purpose digital computer. Bartik (2013) stated, “When ENIAC was announced in 1946, it was heralded in the press as a giant brain” (p. 23).




Figure 1. The ENIAC Computer. Source: United States Army (n.d.).

It was known as the “giant brain” not only because of its processing power, but also because of its size. Jeffrey Yost (2005) described, “The machine which would consist of 18,000 vacuum tubes and have dimensions of thirty by sixty feet” (p. 29). Computers were expensive in those days and were mainly used and funded by the military and universities. According to Goldstine (1972), the ENIAC computer was first intended to perform rudimentary calculations for military purposes, but was later repurposed to perform tasks deemed more important by the powers that be, including the development of the hydrogen bomb.

As digital computers evolved from the ENIAC, a group of computers came to be known as mainframes. These computers used advancements made over the years, through trial and error, using computers, such as the ENIAC. According to Beach (2012), “Mainframes are large, powerful computers that handle the processing for many users simultaneously (up to several hundred users) which caused an enormous expansion of the computing community” (para. 19). Operators use what is known as a terminal to connect to the mainframe. These bulky devices had a screen and keyboard used to input data and were also used to view the output as well, but did not do their own processing. Beach (2012) continued, “The processing power of the mainframe is time-shared between all of the users. Mainframes typically cost several hundred thousand dollars. They are used in situations where a company wants the processing power and information storage in a centralized location” (para. 19). Although mainframes were similar to the early digital computers of the 1940s and 1950s, they provided the ability for multiple users to be able to run programs on the computer, which interested businesses rather than just government and universities (Anderson, 1994). The paradigm used in the mainframe was centralized computing, but now it was a multi-user system. Unlike the early computers, mainframes were not good at number crunching, which deals with the types of problems whose primary constraint is calculation speed. Instead, they are better suited for processing large amounts of data reliably that deals with problems constrained by input and output. Mainframes are suited for performing thousands upon thousands of concurrent transactions, such as large amounts of data from credit card or payroll transactions (“Mainframes vs. Supercomputers,” 2013).

B. THE PC REVOLUTION

Many books and articles have been written about the personal computer (PC) revolution, which occurred between the mid-1970s and the mid-1980s; an exciting time for electronics enthusiasts. A paradigm shift was taking place at a rapid rate and many computer scientists and engineers were eager to get involved. With the use of new technology, decentralized computing was quickly taking shape. Computing was moving from the mainframe (which used many terminals in a room with the processing done at a centralized location) to being used in offices and in the home. The early PCs were faster, smaller, and less expensive than computers that came before them and the processing was being done in a box sitting on a desk.

PC magazine defines a personal computer as “A single user computer. The term was very popular in the 1980s when individuals began to purchase their own computers for the first time in history. ‘Microcomputer’ was another widely used term. Today, the terms PC, desktop, laptop, and just plain ‘computer’ are synonymous with personal computer” (“Personal Computer,” n.d., para. 1). One major factor that brought the PC into the home was the sale price, which was considerably lower than the previous generation of computers. Another factor that played a part of the PC’s popularity was the size, as most were able to fit on a desktop. Dhir (2004) stated, “beginning in the late 1980s technology advances made it feasible to build a small computer that an individual could own and use.” PCs were intended to be utilized by a single user without an operator being present, which contrasts with the mainframe computers that used time-sharing models that allowed use by several users, usually all at once.

Much debate has resulted about which computer holds the honor of being called the first PC. Pundits seem to agree that the advent of integrated circuits (IC) has transformed modern day society. The invention of the IC led to the invention of the microprocessor. The integration of a complete central processing unit (CPU) onto a single chip greatly decreased the cost of processing capability. With that said, a microprocessor is defined as “an integrated circuit which performs the central processing function in a digital computer system” (Cluley, 1983, p. 1) Several companies were working on IC that could process data at high speeds, but one company released a chip in

the early 1970s that garnered much attention. The Intel 4004, which was released in 1971, is generally regarded as the first commercially available microprocessor (Mack, 2005).

Shortly after its release, Intel made the 4004 available to customers. “In early 1971, Intel delivered the 4004 microchip, or microprocessor, to Busicom (a Japanese calculator manufacturer seeking a logic chip)” (Yost, 2005, p. 168). Even in the company’s infancy, Intel looked to make improvements to its processors. “Intel followed the 4004 with the more powerful 8008 in 1972 and the 8080 in 1974” (Yost, 2005, p. 169). The 8080 was an 8-bit processor that ran a few hundred thousand instructions per second. “Earlier processors were used for calculators, cash registers, computer terminals, industrial robots, and other applications, the 8080 became one of the first widespread microprocessors” (“Intel 8080 Microprocessor,” 2016, para. 3). Its use also spread due to its versatility. For example, “its enhanced instruction set and its subsequent role as the original target CPU for CP/M, the first de facto standard personal computer operating system” (“Intel 8080 Microprocessor,” 2016, para. 3).

Large and small companies were not the only ones interested in using the new microprocessor. Edward Roberts, who ran a small radio and airplane firm called Micro Instrumentation Telemetry Systems (MITS), recognized this emerging market and moved into the personal computer field in 1974 (Yost, 2005). Although other computer hobbyists successfully designed and built personal computers, Roberts was the first to garner much attention. Roberts and his team designed the Altair 8800 computer kit, which was built around the new Intel 8080 microprocessor (Yost, 2005). Roberts successfully marketed the Altair 8800 to the editors at *Popular Electronics* magazine who agreed to feature the, as yet incomplete, Altair. “The story appeared in the January 1975 issue with the cover photograph” (Yost, 2005, p. 170). This personal computer was among the first group of computers small enough and cheap enough to be made available to ordinary consumers.

An avid reader of *Popular Electronics* magazine was none other than Bill Gates III, who had read the article about the Altair 8800. “Bill was intrigued with computer technology and learned how to program in BASIC (Beginners All-purpose Symbolic

Instruction Code) on his high school's minicomputer time-sharing system back in 1969" (Yost, 2005, p. 170). Gates, and his friend, Paul Allen, recognized that just like minicomputer time-sharing systems, the Altair would need a BASIC translator to be more readily programmable rather than flipping switches to code in machine language. Gates and Allen quickly formed a company called Microsoft and contacted Roberts and developed a BASIC translator that the firm delivered to MITS in February 1975.

MITS competitor, IMS Associates Inc. (IMSAI), was gaining attention, "IMSAI would be the first to take advantage of a microcomputer operating system, a system built by a firm called Digital Research. In the mid-1970s Digital Research's founder, Gary Kildall, had developed an operating system, CP/M (Control Program/ Monitor)" (Yost, 2005, pp .171-172).

Larger corporations were now taking note of the attention garnered by smaller companies that provided these ready to assemble microcomputers and wanted to get involved. "IBM approached Digital Research in 1980 at the suggestion of Bill Gates as IBM was looking for an operating system for their IBM PC" (Isaacson, 2014, p. 358). Digital Research, a company located in Pacific Grove, California, was called upon to provide a unique product to the company, but that opportunity quickly disappeared due to less than ideal circumstances. Gary Kildall was out delivering software when IBM representatives sought to meet with him (Wallace & Erickson, 1993, p. 155). Due to the missed meeting, the IBM representative, Jack Sams, went back to Bill Gates to let him know about the lack of results from the meeting and "ordered him to get an acceptable operating system. Weeks later, Gates proposed using 86-DOS (QDOS), an operating system similar to CP/M that Tim Patterson of Seattle Computer Products (SCP) had made for hardware similar to the PC" (Wallace & Erickson, 1993, pp. 157-158). Wallace and Erickson went on to explain, "Microsoft made a deal with SCP to become the exclusive licensing agent and later the full owner of 86-DOS. After adapting the operating system for the PC, Microsoft delivered it to IBM as DOS" (Wallace & Erickson, 1993, pp. 157-158).

When the IBM 5150 PC was introduced in August 1981, customers had the option of selecting which operating system they wanted to purchase for an additional fee. Users

could choose “DOS, which is the IBM PC’s ‘standard’ operating system, CP/M-86, and the p-System, both of which are alternative choices. PC-DOS was priced at \$40 while the p-System cost \$675 and CP/M cost between \$300–\$350” (Glatzer, 1982, p. 51). As could be imagined, DOS was the most popular selection, as it was the most affordable at the time. The business relationship between IBM and Microsoft changed the course of history and still affects the computer industry to this day.

The personal computer did not become mainstream until several other technologies were integrated with the PC. This integration occurred when technologies, such as hypertext transfer protocol (http), uniform resource locators (URLs), transmission control protocol/Internet protocol (TCP/IP) protocol, browsers, and Advanced Research Projects Agency Network (ARPANET)/Internet were used to transform the PC into a powerful communication, processing, and multimedia device.

C. OVERVIEW CLIENT-SERVER ARCHITECTURE

The client-server architecture became an efficient way to connect computers together so that they could communicate over a network with servers sharing resources or services with the clients. “Systems which are well suited for the client-server architecture model are financial, mathematical, statistical analysis, CAD, medical, and software development to name a few” (Kavanagh, 1995, p. 39). The advent of the PC, coupled with computer applications, such as email, network printing, and the World Wide Web (for example), laid the foundation for the client-server architecture’s popularity. This architecture, when used with the thick client, supports the decentralized computing model that would remain the dominant model for years to come.

A client server network is made up of two basic parts, a remote client and a server. The server is a central computer, which is used “to facilitate communication and resource sharing between other computers on the network, which are known as clients” (Dean, 2012, p. 5). Usually, servers have better components that make them perform more robustly than most clients. A server consists of more physical memory, processors, and more storage space as compared to a personal computer. Servers may also be equipped with additional hardware that provides fault tolerance in case a hardware device

fails. For example, a server might contain an extra NIC, which is configured in such a way that if the primary NIC fails, the network traffic will fail-over to the secondary NIC to ensure network transmissions continue uninterrupted. Dean explains, “Clients take the form of personal computers, also known as workstations, or mobile devices, such as smartphones” (Dean, 2012, p. 5). (See Figure 2)




Figure 2. Basic Client Server Network. Source: Dean (2012, p. 6).

Dean (2012) states that every computer on a client/server network either acts as a client or as a server (p. 5). Dean goes on to state that it is possible, but not very common, for some computers to act as a client and as a server. “Clients on a network can still run applications from and save data to their local hard disk. But by connecting to a server, they also have the option of using shared applications, data, and devices. Clients on a client/server network do not share their resources directly with each other, but rather use the server as an intermediary” (p. 5). An example would be the accessing of a web page. The client (web browser) institutes a connection to the server over a network, usually a local area network (LAN) or a wide-area-network (WAN), such as the Internet, and requests services or resources. Specifically, the web browser requests services from a web server. The server processes the request and provides the services. In the example, the service provided by the web server is the delivery of a web page. Once the server has completed the client’s request, the connection is ended.

Communications among multiple computers on a network was accomplished through the open systems interconnection (OSI) model (see Figure 3). This model from

the 1980s was based on ARPANET fundamentals that allow computers to communicate based on the following seven layers: physical, data link, network, transport, session, presentation, and application (Dean, 2012, p. 42). This model is the foundation for the client-server architecture, which is currently in use in most networks including the labs and classrooms at NPS.




Figure 3. The OSI Model. Source: Rivero (2015).

D. OVERVIEW OF THICK CLIENT COMPUTING

Although the development of the PC was discussed in a preceding section, the PC has been utilized in different ways; thereby, gaining new names to differentiate itself from other classes of computers. In this section, the thick client is introduced. The ways in which it is used show how it differentiates itself from other computers, such as servers, workstations, and thin clients.

A thick client will generally perform the bulk of data processing on its end. “A thick client (also heavy, fat, or rich client) is a computer (client) used in the client-server architecture of networks that typically provides functionality independent of the central server” (Christensson, 2006). The thick client can benefit by communicating with a server, “The server may provide the thick client with programs and files that are not stored on the local machine’s hard drive” (Christensson, 2006). It is commonplace for “workplaces to provide thick clients to their employees where servers are used to provide some data and application support” (“Thick Client,” 2016, para. 1). Thick clients enable them to access files on a server or to use the computer as a standalone device. “Thick clients have their own operating system and software applications can be used offline (not connected to a network or server)” (“Thick Client,” 2016, para. 2). However, when part of a client-server architecture, the thick client (if used as a stand-alone) may not be able to use distributed computing applications and services that require communication with other computers or computer equipment to function. A good example would be a software program that obtains a license from a license server to provide the necessary information needed to run the program. Nor would a thick client be able to use a network printer if not connected to the network.

The Computer Hope website, <http://www.computerhope.com/jargon/t/thickcli.htm>, lists four advantages with using thick/fat clients. The first such advantage is the ability to work offline. Thick clients do not need to have a constant connection with the server as thin clients do. The second advantage is better multimedia performance. Thick clients have the advantage as far as multimedia-rich applications. These applications would be bandwidth intensive if fully served. An example is video gaming. Thick clients are more efficient, as the data is processed locally rather than transferred across the network and back. Thirdly, thick clients allow for increased flexibility. The article shows that “on some operating systems software products are designed for personal computers that have their own local resources” (“Thick Client,” 2016, para. 1). The Computer Hope website explains that running such software can be challenging on a thin client. The fourth advantage is greater server capacity. “The more work that is carried out by the client, the less the server needs to do, increasing the number of users each server can

support” (“Thick Client,” 2016, para. 2). Another benefit is the use of existing infrastructure. Since many computer users have fast computers, the infrastructure to run thick clients is already available without additional expenditures.

E. OVERVIEW OF CLOUD COMPUTING

The term “cloud computing” was first used in 1997 by IT and information systems Professor Ramnath Chellappa. He used the term to describe computers working in concert, which provide the new computing system (Chellappa, 2010). He said, “the cloud would be a new computing paradigm where the boundaries of computing will be determined by economic rationale rather than technical limits alone” (Chellappa, 2010, Cloud Computing section, para. 2). However, the “new cloud computing” paradigm has its roots entrenched in technologies that were used in the centralized mainframe computer paradigm. Cloud computing also has characteristics used in the decentralized computing paradigm that stretched from the 1980s until today. So a question that comes to mind is the following, is cloud computing centralized or decentralized? It can be said that cloud computing is centralized computing. For an example of cloud computing being used as centralized computing, look no further than the implementation of virtualization. Korzeniowski discusses the relationship between centralization and virtualization, and the benefits of virtualization:

Centralization consolidates a number of autonomous department servers onto a larger system. Replacing smaller, specialized systems reduces an organization’s technology footprint. Virtualization allows companies to collapse their processing tasks onto fewer systems, lowering storage, networks, power use and cooling costs. (Korzeniowski, 2013)

When people think of cloud computing, many different thoughts come to mind. Many people think of the old schematics depicted in books where a cloud is drawn in the schematic with clients, servers, and network equipment surrounding the cloud and are shown to connect to it with lines pointing to the cloud (see Figure 4).




Figure 4. Simple Network Diagram Depicting a Cloud as the Internet. Source: Lindsey (2007).

Computer science students may think of the cloud symbol used to show networks of computer equipment in the original ARPANET as early as 1977 when Cerf and Kahn successfully linked three networks in a dramatic round the world transmission in a cruising van (“Revolution,” n.d.).

To the author, cloud computing is the sharing of network resources, which leads to a reduction in operating costs while leveraging resources to increase profits and coverage. Cloud computing is based around sharing services supported on a larger installed infrastructure similar to the days when mainframes ruled the computing world. Some popular on-demand computing services utilized on a daily basis from the cloud are on-line data storage (including music, videos, photos), web-based email, database processing, software applications, and more.

Cloud computing is focused on maximizing the effectiveness of resources. “The Cloud” refers to a set of resources widely distributed while the underlying methods of delivery are unknown to the end user, much like the haziness of a cloud. One of the key features of a cloud environment is the ability of the resources granted to a customer to be scaled dynamically up or down easily by the cloud service provider (“Elastic Computing (EC),” n.d., para. 1). Thus, regardless of the distribution of queuing for resources,

customers can effectively grow their resource pool to meet high demands when necessary, and dynamically shrink their resources back to normal levels to reduce costs (“Elastic Computing (EC),” n.d., para. 3).

The National Institute of Standards and Technology (NIST) is a non-regulatory agency of the United States Department of Commerce. The institute’s mission is to “Promote U.S. innovation and industrial competitiveness by advancing measurement science, standards, and technology in ways that enhance economic security and improve our quality of life” (NIST, 2010). NIST’s definition on cloud computing as defined in Special Publication (SP) 800–145 titled, “The NIST Definition of Cloud Computing,” is as follows:

Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. (Mell & Grance, 2011)

F. ESSENTIAL CHARACTERISTICS OF CLOUD COMPUTING

The NIST publication by Mell and Grance is referred to throughout the next three sections. The NIST organization identifies five essential characteristics of cloud computing. These “five essential characteristics” are defined in SP 800–145, and NIST identifies them as follows.

- *On-demand self-service.* A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.
- *Broad network access.* Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations).
- *Resource pooling.* The provider’s computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. Generally, the customers feel a sense of location independence because they have no control or knowledge over the exact location of the provided resources but may be able to specify a location at

a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, and network bandwidth.

- *Rapid elasticity.* Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.
- *Measured service.* Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported providing transparency for both the provider and consumer of the utilized service. (Mell & Grance, 2011)

G. CLOUD COMPUTING SERVICE MODELS

Cloud computing is based upon three major service models discussed by NIST: Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS) (Mell & Grance, 2011). Figure 5 shows the component services as they are presented in a cloud computing architecture. Desktop as a Service (DaaS) is an emerging service model that merits discussion. The DaaS model was in its infancy at the time when NIST wrote the SP-800-145 document.

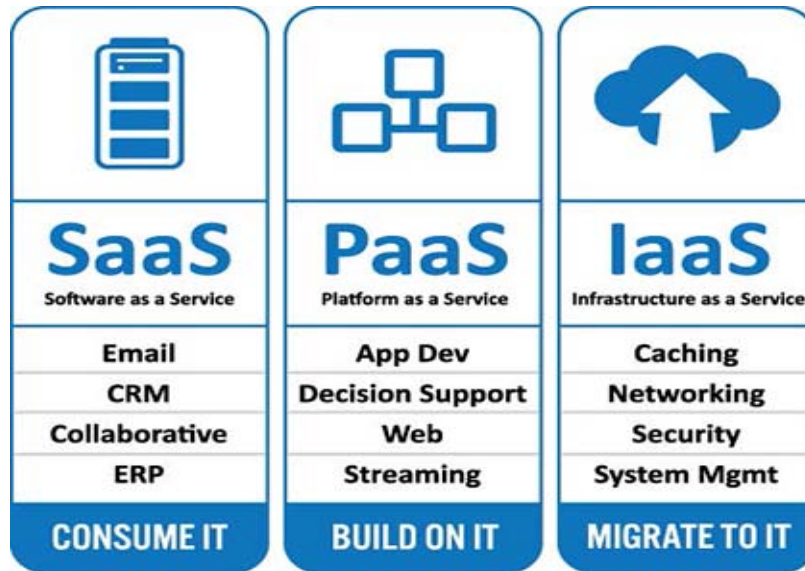


Figure 5. Cloud Computing Service Models. Source: “Cloud Services” (2015).

1. Software as a Service

The first service model is SaaS. SaaS is a software delivery model, which differs from the traditional method of installing software locally on a thick client computer by permitting applications to be deployed over the Internet. The delivery model allows software to be licensed on a subscription basis and is centrally hosted. These software applications are accessed via a web browser or a program interface (Mell & Grance, 2011). The SaaS service provider manages the cloud infrastructure including the network, servers, operating system, and storage. In some cases, the user may be able to manage limited user-specific application configuration settings (Mell & Grance, 2011).

An example is email hosting like Yahoo mail or Gmail. Another example is Salesforce.com, which provides customer relationship management (CRM) software, such as Sales Cloud, as well as the customizable support and helpdesk software called Service Cloud. Other popular SaaS providers include Google, which offers the Google Docs application among other SaaS offerings. Although many more SaaS providers exist, one familiar provider worth mentioning is Microsoft, which is a huge SaaS provider that provides software solutions to government and enterprise customers. Some popular

Microsoft SaaS offerings are Windows Live, Office Live, Exchange Online and SharePoint online to name a few (Tyson, 2014).

2. Platform as a Service

The second service model is Cloud Platform as a Service (PaaS). According to Mell and Grance (2011), this service model provides a platform that allows customers to develop or acquire software, such as programming languages, tools, and services, to deploy to the cloud infrastructure supported by the provider. In this case, the customer does not control or manage the Cloud infrastructure, as shown by Mell and Grance, but does have control over the deployed software and most likely the application hosting environment. In essence, a cloud provider provides the tools needed for application development to the customer. These tools are provided as a service.

There are many PaaS providers, which help developers build and deploy their applications to the cloud. One such provider is Amazon Web Services (AWS) with its Elastic Beanstalk service. Elastic Beanstalk changes how developers push their apps into Amazon's cloud. As stated by Amazon:

AWS Elastic Beanstalk is an easy to use service for deploying and scaling web applications and services developed with apps such as Java, PHP, Python, Ruby, Go, and Docker on familiar servers such as Apache, Nginx, Passenger, and IIS. End users can upload their code and Elastic Beanstalk handles the deployment, which includes, capacity provisioning, load balancing, auto-scaling to application health monitoring. ("Elastic Beanstalk," n.d., para. 1–2)

As stated on its website, "Windows Azure offers integrated tools, pre-built templates and managed services which make it easier to build and manage enterprise, mobile, web, and Internet of things apps" ("Microsoft Azure," n.d., para. 1).

3. Infrastructure as a Service

The third service model is Cloud IaaS. Mell and Grance stated that the IaaS model is defined by NIST as the ability delivered to the consumer:

to provision fundamental computing resources such as processing, storage, networks, and other resources where the consumer is able to deploy and

run subjective software, which can include operating systems and applications. The consumer does not manage or control the underlying Cloud infrastructure, but does have control over resources such as operating systems, storage, and deployed applications.” (Mell & Grance, 2011, Service Models section, 3)

Mell and Grance go on to explain that the consumer may gain partial control of some networking components, such as firewalls.

AWS EC2 has become one of the top IaaS providers. It is a web service that provides resizable compute capacity in the cloud. As shown on the Amazon website, “Its compute capacity set the standard for spinning up and taking down cloud capacity quickly and affordably with a pay-as-you-go model” (“Amazon EC2,” n.d., para. 2). Another example of an IaaS provider is AT&T Synaptic. It offers a services called SoftLayer. These products include a menu of corresponding services, such as security, system monitoring, storage, and application options included with virtual servers and standard systems. According to the Synaptic website, you can tailor your solution with a range of optional features such as software, networking and monitoring, security, and storage (“Softlayer Services,” n.d.).

4. Desktop as a Service

The final service model to be discussed is DaaS. DaaS is a cloud service in which a cloud service provider hosts the back-end of a VDI. DaaS has a multi-tenancy architecture and the service is purchased on a subscription basis. In the DaaS delivery model, the service provider manages the back-end responsibilities of data storage, backups, security, and upgrades (Rouse & Botelho, 2015). In a typical use case, the customer’s personal data is copied to and from the virtual desktop during the process of logging on and logging off the system. Access to the desktop is accomplished with a client device (e.g., desktop computer, laptop, thin client, zero-client, or a smart device). “Customers usually manage their own desktop images, applications, and security. The service is purchased on a subscription basis” (Rouse & Botelho, 2015).

VMware is a major provider for the DaaS platform. Their service, Horizon Air Cloud-Hosted Desktops and Apps, delivers Windows desktops and hosted apps as a

cloud service to any user, from any location, and on any device. Horizon Air Cloud-Hosted Desktops are fully customizable, just like a physical desktop. Air desktops also allow users to extend the life of their desktop hardware while using the latest OS and apps to the end users (“Simplify the Delivery,” n.d., para. 7). Another popular DaaS provider is Citrix. Citrix enables all businesses to embrace a mobile workstyle. “Citrix Service Provider-based DaaS solutions provide complete PC-style Desktops-as-a-service for applications and email securely delivered over the web. Desks-as-a-Service are simple to buy and easy to manage with no software for IT to maintain” (“Desktops-as-a-Service (DaaS) and Hosted Desktop Solutions,” n.d., para 1.).

H. CLOUD COMPUTING DEPLOYMENT MODELS

There are four primary computing models, which are known as deployment models. The private cloud, community cloud, public cloud, and hybrid cloud are defined by Mell and Grance (2011).

- **Private cloud.** The private cloud infrastructure is provisioned for the exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.
- **Community cloud.** The community cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.
- **Public cloud.** The public cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the Cloud provider.
- **Hybrid cloud.** The hybrid cloud infrastructure is made up of two or more distinct Cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds). (Mell & Grance, 2011)

As acceptance of Cloud computing has increased, corporations are validating the fact that Cloud consumption may possibly aid in IT cost reduction and can improve scalability of the environment.

I. VIRTUALIZATION

As a technology, virtualization is nothing new. The term dates back to the mid-1960s around the time that, “IBM developed the early incarnations of what would become the IBM System/360 Model 67, which was considered a virtual machine/virtual memory operating system” (McCafferty, 2014). More recently, companies, such as Citrix, VMware, and Microsoft, have emerged as leaders in the virtualized computing world as virtualization is beneficial in that, “virtualization is nothing more than an increasingly efficient use of existing resources that delivers huge cost savings in a brief amount of time” (Portnoy, 2012, p. xv).

At the center of virtualization is the virtual machine (VM) (see Figure 6):

A VM is a tightly isolated software container with an operating system and application inside. Because each virtual machine is completely separate and independent, many of them can run simultaneously on a single computer. The hypervisor (thin layer of software) decouples the virtual machine from the host and it dynamically allocates computing resources to each virtual machine as needed. (“VMware,” 2014, p. 4).




Figure 6. Graphic Depiction of Structure of Virtual Machine from Virtualization Essentials. Adapted from VMware (n.d., p. 4).

According to the VMware virtualization document, VMs offer several benefits, the first of which is partitioning (VMware, n.d., p. 3). With partitioning, a user can run more than one operating system on one physical computer. System resources can also be divided among computers. Another benefit of VMs is isolation. As noted in the VMware document, isolation provides fault and security isolation at the hardware level and preserves the performance with advanced resource controls (VMware, n.d.). Encapsulation is another benefit provided by VMs. With encapsulation, the entire state of the virtual machine can be saved to files and VMs can be moved and copied as easily as

files. The final benefit is hardware independence. “Any VM can be provisioned or migrated to any physical server” (“VMware Virtualization Essentials,” n.d., p. 4).

J. VIRTUAL DESKTOP INFRASTRUCTURE

According to Rouse and Barrett (2016), “Virtual Desktop Infrastructure (VDI) is the practice of hosting a desktop operating system within a virtual machine (VM) running on a centralized server. VDI is a variation on the client/server computing model, sometimes referred to as server-based computing” (p. 1).

As can be seen in Figure 7, a VDI model delivers the same functionality as a client/server model. This functionality provides a centralized desktop by using a protocol, such as PC over IP (PCoIP), remote desktop protocol (RDP), or independent computing architecture (ICA). Excluding the virtualization layer and the multiple operating systems, the two solutions look very similar.

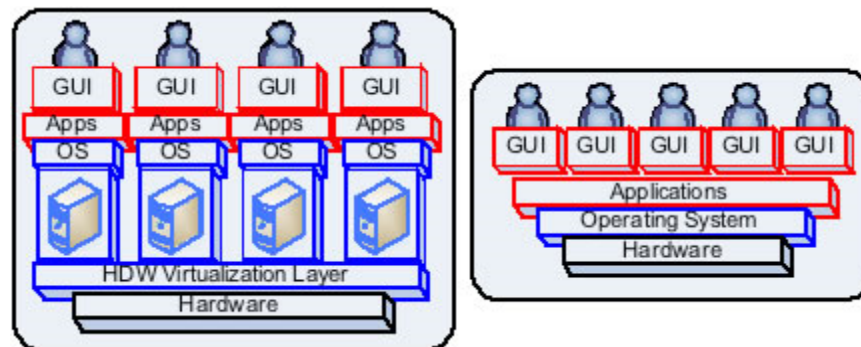


Figure 7. Representation of VDI Model versus Client Server Model. Source: Oglesby (2006, What Is VDI? section, para. 1).

The image on the right shows a traditional client-server model. “The users all have access to a desktop GUI via an individual session on the terminal server. This server has a single OS installed; an instance of Terminal Services to provide sessions and session management. The server also has software applications installed that can be used by all users” (Oglesby, 2013, para. 4). This type of use contrasts with the VDI model in several ways. “In the VDI model on the left, a single server is used again, but a hardware virtualization layer is added to this server in place of a more traditional OS like Windows

Server. The Virtualization layer provides numerous Virtual Machines that are each supplied with an operating system, applications and a unique GUI / desktop environment for each user” (Oglesby, 2013, para. 5).

However, as Oglesby (2013) states with VDI, the following benefits are gained:

- The ability to provide a unique environment for each user.
- Each environment can be customized with different apps and settings without impacting other users.
- Users can be granted control of their own virtual desktop to allow them to install and modify applications if needed.
- Applications that were not multi-user friendly can be run in this environment since each instance is like a new desktop. (para. 6)

Along with the aforementioned benefits, Madden (2014) noted two major technological advancements that came out in 2013 that are making VDI more attractive in many use cases. The first was in storage technologies that enable high-performance fully persistent disk images to be offered at a much lower cost than before. This advancement allows IT administrators to design and manage VDI desktops in the same way traditional desktops have been designed over the last 20 years. Simply stated, IT administrators can design VDI in a way that works in their environment (Madden, 2014). The second improvement to VDI was with graphics. PCs and laptops have had graphical processing units (GPUs) built into them for years. Most everything a user does on a computer requires a GPU these days. Traditionally, hypervisors have not virtualized GPUs in the past. In 2013, most companies that create hypervisors have begun to virtualize hypervisors. Led by Nvidia with its grid technology, full GPU support is now available in VDI desktops so that VDI users can now do everything in their remote VDI desktops that they could do on traditional PCs and laptops. As with any software, with each new version, hypervisors are more stable, have fewer bugs and more features added (Madden, 2014). All the benefits and improvements to VDI have without a doubt contributed to the analysts forecast that the VDI market in the United States is set to grow at 29.70% CAGR by 2019 (“Virtual Desktop Infrastructure (VDI),” 2015).

K. BRING YOUR OWN DEVICE

BYOD refers to the strategy of authorizing users to bring mobile devices that they own, such as tablets, smart phones, and laptops, to school or to work and to utilize such devices to access data and software.

The use of mobile devices in peoples' daily lives has increased over the years and has led to companies, such as Intel, to allow employees to bring their own devices to work. The light and compact device together with the fact that the devices have multiple uses has made it an attractive device to bring to work and school. The concept has gained popularity in recent years mainly due to the following reasons: employees are willing to spend money on their devices as they have ownership of the device, maintenance and protection of the device is a priority as the employee will be liable for losses, they allow employees to be more flexible and add more productive hours as they contribute more to the organizations growth from anywhere and at any time, reduces the burden of IT inventory maintenance, and hardware purchase costs are lowered (Deep, 2013, para. 1–7). Intel has not only been the leader with the microprocessors, but it has also led the way with IT implementation:

The trend started at Intel in 2009 when employees began using their own smartphones, tablets, and mobile storage devices on the job. Rather than reject the trend, as many organizations initially attempted, Intel's senior leaders were quick to embrace it as a means to cut costs and improve productivity. Since January 2010, the number of employee-owned mobile devices on the job has tripled from 10,000 to 30,000 and by 2014 Intel's CISO Malcolm Harkins expects that 70% of Intel's 80,000 employees will be using their own devices for at least part of their job. (Harkins, 2016)

When an organization implements a BYOD project within its IT infrastructure concerns need to be addressed. One concern is how are students going to access resources they need to complete projects and assignments, as well as collaborate. Many school networks are set up to give their devices access to the Internet only, and in some cases, to their email. "They had not anticipated how students would gain access to shared network drives, or to specialized applications that they could access without purchasing their own software license" (Raths, 2013, p. 9). One possible solution, which is evaluated, is to

provide students with a desktop by using VMware View (VDI) to permit students to access the stated resources from a variety of devices.

L. THIN CLIENT

Thin clients have their roots set in mainframe systems when multiple users would access the system using dumb terminals. As computers graphics became more advanced, these terminals transitioned from providing a command-line interface to a full GUI. In the 1990s, computer manufacturers began to shrink the desktop PC, which became a computer trend. As computer networking became more common, simpler to use computers became more common. In fact, these slimmed down computers came to be known as thin clients and are still currently in use (see Figure 8). “Thin client are computers that do not have internal hard drives, but rather let the server do all the work, but then display the information” (Velte, Velte, & Elsenpeter, 2010, p. 7). In essence, a thin client is a client computer that depends on the server to accomplish the workload.




Figure 8. Example of a Dell Wyse Thin Client Computer. Source: “Dell Wyse D90D7 Thin Client” (n.d.).

This thin client contrasts with the traditional PC/thick client, which mostly processes data on its own. Either a dedicated thin client terminal or a regular PC with thin client software (GUI, cloud access agents (RDP, ICA, PCoIP), a local web browser,


terminal emulators, and a basic set of local utilities) is used to send keystrokes and mouse input to the server and receive screen output in return (“Thin Client,” n.d., para. 1). The processing is done on the back end rather than at the terminal. “The thin client does not process any data; it processes only the user interface (UI)” (“Thin Client,” n.d., para. 1). The data processing occurs on the server or server farm. The server model makes use of cloud computing concepts, such as VDI, which brings up the reason why this research is interested in thin client technology. Thin clients are one proposed medium by which users can access VDI virtual machine desktops. This combination creates what is currently known as a cloud centric system where the datacenter provides desktop resources to be centralized. “Using products such as VMware Desktop Manager (VDM), the VDI component in Remote Desktop Services and Citrix XenDesktop, each user’s desktop (OS and applications) resides in a separate partition in the server known as the virtual machine (VM). Users are essentially presented with their own PC, except that it physically resides in a remote server in the datacenter” (“Thin Client,” n.d., para. 6).

Thin-clients typically do not have input/output (I/O) ports, hard disks, or other unnecessary features although newer models have ports for universal serial bus (USB) devices. A benefit of thin clients is that they are a low cost solution for businesses or organizations. “The benefits are improved maintenance and security due to centralized administration of the hardware and software in the datacenter” (“Thin Client,” n.d., para. 1). Other benefits include ease of use, low cost when compared with traditional PCs, and lower power consumption. Some disadvantages for thin clients are the lack of peripheral devices, such as CD and DVD drives, and limited performance with tasks, such as video rendering, graphics editing or gaming. With the advantages outweighing the disadvantages, it benefits this research to look into thin clients as a possible alternative for lab and classroom computing.

M. CONTAINERS

A relatively new technology is being implemented in the IT field. The technology is called Containers. “Docker containers wrap up a piece of software in a complete filesystem that contains everything it needs to run: code, runtime, system tools, system libraries,

anything you can install on a server. This guarantees that it will always run the same, regardless of the environment it is running in” (“What Is Docker,” n.d., Package Your Application into a Standardized Unit for Software Development section, para. 1). (See Figure 9).



Docker container representation of a piece of software in a complete filesystem that contains all software needed to run.

Figure 9. Docker Container Representation. Source: “What Is Docker?” (n.d., Package Your Application into a Standardized Unit for Software Development section, para. 1).

Containers are considered to be lightweight. “Containers running on a single machine all share the same operating system kernel so they start instantly and make more efficient use of RAM. Images are constructed from layered filesystems so they can share common files, making disk usage and image downloads much more efficient” (“What Is Docker,” n.d., Lightweight section, para. 2). They are also considered to be open source software. “Docker containers are based on open standards allowing containers to run on all major Linux distributions and Microsoft operating systems with support for every infrastructure” (“What Is Docker,” n.d., Open section, para. 2). The implementation of this software package allows for better security. “Containers isolate applications from each other and the underlying infrastructure while providing an added layer of protection for the application” (“What Is Docker,” n.d., Secure by Default section, para. 2).

Docker containers share many aspects with virtual machines. However, their differences are what make them unique. A VM (see Figure 10), “includes the application, necessary binaries, libraries, and an entire guest operating system.” Each software

component takes up space and when added together, the VM may be tens of GB in size (“What Is Docker,” n.d., virtual machines section, para. 3).

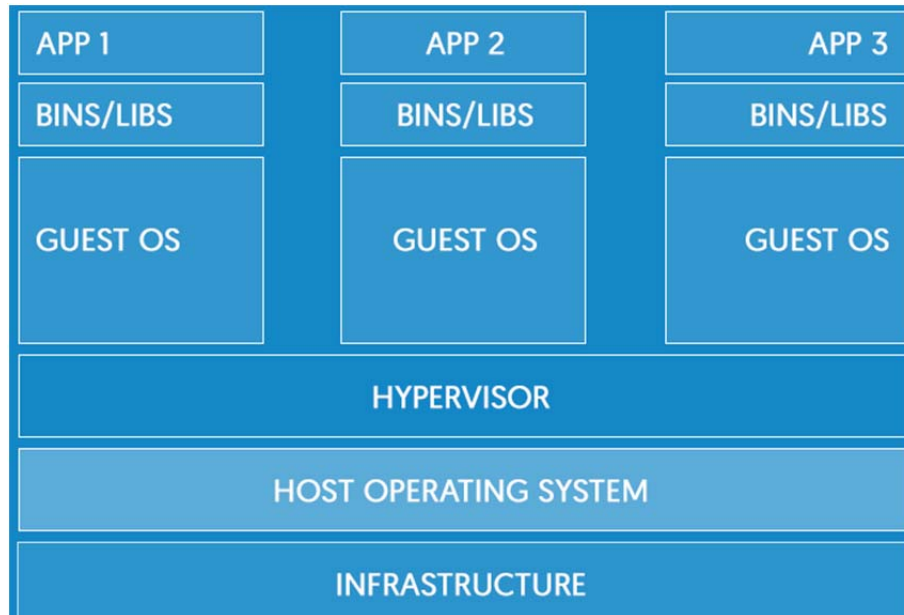


Figure 10. Architectural Approach for Virtual Machines. Source: (“What Is Docker,” (n.d., Virtual Machines section, para. 3).

Containers are all-inclusive yet allow for sharing of resources. “Containers include the application and all of its dependencies, but share the kernel with other containers. They run as a stand-alone process in user space on the host operating system. They are also not tied to any specific infrastructure—Docker containers run on any computer, on any infrastructure and in any cloud” (“What Is Docker,” n.d., Containers). (See Figure 11).

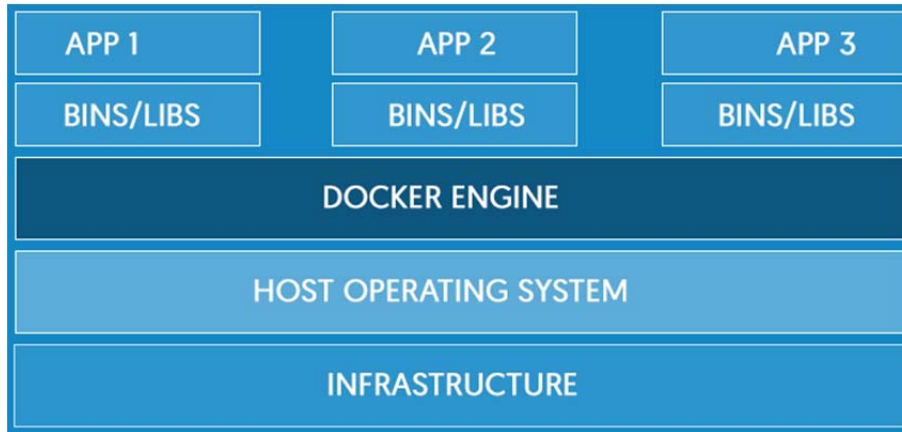


Figure 11. Architectural Approach for Docker Containers. Source: “What is docker?” (n.d., Comparing Containers and Virtual Machines section, para. 2).

As actions are performed to the base image of a Docker container, the union filesystems are formed and documented in such a way that the creation of an action is described in its entirety. “This strategy enables Docker’s lightweight images, as only layer updates needed to be propagated compared to full VMs” (“What Is Docker,” n.d., Lightweight section, para. 1).

In this chapter, the background information and concepts needed to analyze the technology used in the current labs and classrooms at NPS has been presented. This chapter looked at the beginnings of the modern day digital computer to show the size of the computer and the effort involved in running the computer (amount of resources, both human and physical), which was much less powerful than most smartphones in use today. The mainframe was shown to be a useful computer, which utilized the paradigm of centralized computing. Next, this research looked at the birth of the personal computer and the paradigm shift to decentralized computing. The PC computer, more specifically the thick client computer, performs all the processing and allows the user to run locally installed applications. This technology, combined with the client-server architecture, was shown to be an efficient way to connect computers together so that they could communicate over a network with servers sharing resources and services with the clients. This model is utilized in the LRC labs at NPS today. The research next took an in depth look at cloud computing along with the essential characteristics, service models, and

deployment models of cloud computing. VDI concepts were also discussed along with two technical devices, which can be used to connect to the VDI infrastructure followed by container application virtualization concepts.

In the next chapter, a few of these concepts and models are shown in a real world use case at NPS and select universities. The computing needs of the NPS student are described, as well as how those needs are being met with the current technology implemented in the LRC labs. The next chapter also presents the need for computing capabilities that has been created with the progression of computer technology in cloud computing. The ITACS department and faculty members at NPS have implemented the VDI technology to meet the evolving need for education in the new computing field. The administrators of VDI private cloud environments present the service driven requirements in this area along with how those needs are being met. The chapter also discusses various computing technology options at other universities. Information is gathered about the IT infrastructure used at the Naval College (Newport, Rhode Island), CSUMB, and Stanford University to provide computing capabilities to students. Those technologies are considered as to whether or not they are feasible in meeting the functional and cost (in terms of hardware, software, and manpower) and availability requirements for the LRC labs at NPS.

THIS PAGE INTENTIONALLY LEFT BLANK

III. LAB TECHNOLOGIES AT THE NAVAL POSTGRADUATE SCHOOL AND SELECT UNIVERSITIES

This chapter describes the current lab technologies used at NPS. These labs have been named Learning and Resource Centers (LRC) to distinguish them from other physical sciences labs at the school. This term is defined so that the reader can understand what services are provided to students in an attempt to outline goals that must be met for the possibility of replacing the existing technology with newer technologies. The existing LRC technology, known as the thick client computers, uses the client server architecture to provide computing services to students. The customers and services provided are shown, as well as an estimate of the cost, manpower, and availability associated with maintaining the LRCs. The information can be used to compare the other technologies shown in this research.

Over the course of the last decade, new technologies have been developed that have offered teaching faculty the ability to present coursework, which was not possible to accomplish with older/different technology, provide 24 hour access to computing resources, offer access to distance learning students, as well as provide access to specialized software not available to LRC users. The virtual LRCs at NPS have taken the form of a private cloud model offered using VDI technology. This chapter presents interview questions, which are asked of four NPS IT administrators to show services provided, the cost of the VDI solution, the manpower associated with maintaining the solution, and the availability related to each solution. The results are described in Chapter IV. The four use cases where VDI technologies are implemented at NPS to provide ubiquitous computing to graduate and PhD level students help the reader understand the reasons for the move to VDI.

The chapter is supplemented with an introduction to various IT computing technology examples used at three universities. Information is gathered about the IT infrastructure used at the Naval College (Newport, Rhode Island), CSUMB, and Stanford University to provide computing capabilities and software services to students. It is shown if they use traditional labs on campus or if they use a different model to provide

computing services. Next, a look at how their labs or equivalent services are configured to provide computer resources to students for completing homework, lab work, and thesis research are presented. A series of interview questions are presented in this chapter. The interviews either are conducted over the phone or in person, but are asked of IT professionals employed at those universities. Their responses are shown in Chapter IV. In addition, the questions can help to reveal their customer base, lab services provided, an estimated cost of the technology, the manpower involved in maintaining the solution, and an estimate of the availability of the various technologies to be provided. These resources may prove to be invaluable in providing options for NPS to deliver software to NPS LRC customers with new technologies that may offer a lower cost and higher efficiency than the current LRC solution.

A. LRC COMPUTING AT NPS

In this section, a description of the academic institution, NPS, is provided. The computer labs at NPS, LRCs, are defined and their functional goals explained. The information technology computing department, now known as ITACS, is discussed along with the groups contained within the department. This section then looks at the LRC customers and the services provided by the LRCs. Next, the thick client computers used in the LRCs are considered for understanding the hardware and hardware costs, manpower, and availability associated with the LRCs.

NPS is a fully accredited university offering masters and doctoral degrees. The United States Navy operates the university. Currently, NPS has approximately 1,600 resident and DL students enrolled at the school. NPS has four sub-schools, which further categorize the curricula taught at NPS. Those sub-schools are the Graduate School of Business and Public Policy (GSBPP), Graduate School of Engineering and Applied Sciences (GSEAS), Graduate School of Operational and Information Sciences (GSOIS), and the School of International Graduate Studies (SIGS). In support of that learning a variety of computer labs referred to as LRCs. According to the LRC Policy, an LRC is, “The LRCs are defined to be centrally funded and centrally maintained teaching

classrooms/laboratories that provide a computer station for each student” (“Learning Resource Centers (LRC) Policy,” 2013, para 1).

In 2004, the researcher accepted a position as an IT specialist working for ITACS at NPS. The position accepted by the researcher was IT Specialist in general, but the position was for a lab manager who was to provide support to the LRCs at NPS.

A set of standards must be met for the LRCs to function at peak performance. The standards are defined by the following:

The LRCs utilize common, centrally administered and funded lab resources to support their operation and maintenance. The standard hardware and software in the LRCs varies based upon the requirements of the particular teaching classroom. If the current hardware and/or software cannot meet specific class requirements, the equipment can be augmented to meet the needs of most departments. (“Learning Resource Centers (LRC) Policy,” 2013, para 2).

In essence, the labs are an extension of the classrooms. Faculty cover information and give assignments, which reinforces the knowledge presented to students or is included in the textbooks. Software programs are used to demonstrate the material, provide examples, and teach powerful software principles, which reinforce the knowledge learned in the classroom. Van Den Blink (2009) sums up the uses of computer labs, which also applies to the way NPS has designed and laid out its labs, by stating:

Traditionally, computer labs have been configured to support teaching and learning by providing rows of computers in a lecture-style classroom set-up. Lab computers and software allowed students to complete course assignments or learn new programs. The uses of technology for teaching have evolved, however, and so must the design and configuration of computer labs. (para. 2)

The IT department, now known as ITACS, has provided all communication services, telephone support, and network support and incorporated those services into the core computing functions at NPS since 1953. The ITACS mission statement is as follows: “Provide technology and communications support for the NPS core mission of teaching, research, and services to the Navy and Department of Defense, and to provide voice, and data infrastructure as mission-crucial enablers of innovation and experimentation within

the educational enterprise” (“ITACS Mission,” 2011). ITACS has several groups within the larger department all of which are dedicated to providing services and support to the NPS mission. The groups are comprised of the Technology Assistance Center (TAC)/Helpdesk (including the PC shop and LRC lab support), the network operations center (NOC), server management, web OPS, high performance computing, educational technologies, telecommunications office (TCO), and cyber security.

1. Customers and Services Provided

The student body at NPS consists of “officers from all branches of the U.S. uniformed services, civilian employees, employees of the federal, state, and local governments, as well as officers and civilians from 47 foreign countries. A limited number of contractors and enlisted personnel are also enrolled” (“NPS Academic Catalog,” 2016). The resident students are the general group of students who use the LRCs located across the campus since remote access to LRCs is not allowed, as it may interfere with the resident students. Also, the resident students are located on campus where the LRCs are physically located.

GSEAS, GSOIS, GSBPP, and SIGS all have access to and may utilize the LRCs. However, the students who use the LRCs use them in most cases based on hardware and/or software requirements mandated by their curricula. The lab computers provide the means for students to utilize software, store data, and have access to the network and the Internet. Furthermore, the computer hardware allows software programs and simulations to run, which helps students in their studies. The LRC computers are connected to the NPS network using Ethernet cables. Each computer station has (at a minimum) a thick client computer, monitor, keyboard, and a mouse. The software provided in the LRCs offers students the capability of working on their homework, lab assignments, preparing for exams, and working on their theses and dissertations.

In a nutshell, many students from all different curricula and certificate programs use the LRCs for the computing services provided, but to varying degrees, depending on with which school the students are associated. For example, GSEAS has many students who are engineering and applied sciences students. These students utilize software used

to create technical models, schematics, run simulations, calculate complex algorithms, and perform an array of complex calculations and number crunching among other requirements. The software, which is used to accomplish those goals, can be mildly memory and processor intensive or can rely heavily on processing and memory to generate output.

The LRC labs, which are located within GSEAS academic buildings, provide computer hardware and software used to accomplish these goals. Many of the software programs are costly and rely on a licensing model, which requires the software to communicate with a licensing server located on the school network. Also, many contracts with software vendors specify that the software can only be installed on government owned computers. That caveat gives validation as to why some of the software is currently provided via the LRC. For the aforementioned reasons, the GSEAS LRCs are used heavily by their students.

The situation is different in GSOIS. GSOIS is home to the computer science department, the information sciences department, the cyber academic group, as well as several other departments. Many various degrees are offered in each of the departments. GSOIS has three associated LRCs in which students can utilize the computers to complete their work. However, GSOIS in particular, has faculty members who provide options to using the traditional LRC computers. For instance, as discussed in a later section of this research, faculty members offer a VDI private cloud, which provides VMs to students for their coursework. Another faculty member in the computer science department offers a similar service; however, the VMs are set up to support courses taught in a particular area/field. Thirdly, another IT professional offers VMs via a VDI environment to faculty and staff in a different area of expertise. As noted by the researcher, who is a GSOIS student, some faculty members offer the Cloudlab private cloud (maintained by ITACS) as a means of providing software to students so they can complete their coursework. The students use BYOD (many of which are personal laptops) to access the VDI environments. The LRCs at GSOIS no longer exist in buildings where they existed at one time and the LRCs that do exist are not used as heavily due to the move to BYOD by GSOIS.

In the case of GSBPP, the use of the LRCs is light at times. LRCs are available for resident students enrolled in the school. Much of the coursework has to do with business and public policy. Many of the software applications are made available to students via the NPS software portal. The faculty also offers the Cloudlab VMs in support of the coursework. The students who use the LRCs use them as they are convenient to use between classes as was stated by Professor A. All the necessary software is installed in one place and a printer is available he added (Professor A, personal communication, August 11, 2016).

Concerning SIGS, the school has two LRCs associated with it. The school specializes in the study and teaching of international relations and policy-based studies. SIGS also offers non-degree courses and programs. Many of the applications used at the school are available via the NPS software portal. The students also have access to the Cloudlab VMs in some cases. The LRCs for this school are not heavily utilized, as many students use BYOD for their coursework.

2. LRC Hardware and Software Cost, Manpower, and Availability

In support of the learning described previously, a variety of computer labs/LRCs are at NPS. Currently, NPS has 10 LRC labs that provide computing resources to students. One Macintosh lab houses 25 iMac computers; however, another group within ITACS maintains this lab and is not discussed. In addition, approximately three dozen LRC classrooms are associated with those 10 LRC labs. The LRC classrooms are used for lectures and have anywhere from 18–80 seats for students with one thick client computer, one monitor, a mouse, and a keyboard located at the front of the classroom for the faculty member. For the scope of this research, the LRC labs are discussed from this point forward and the term LRC is used to refer to the computer labs.

The NPS LRC computing environment has not changed much over the years in terms of the architecture used to deliver software and IT services to the NPS student. The thick client model has been used throughout the 10 physical labs since the early part of the 1990s. Each LRC has anywhere from 18 to 36 thick client computers available for students use. Currently, all computers in the LRCs are Dell computers. The computer

models in half of the labs vary from the models in the other half of the labs due to the recapitalization cycle. Half of the lab computers are purchased with lab recap funds in the course of a fiscal year. The other half of the labs have their computers purchased the next fiscal year. The third fiscal year in the cycle is used to purchase audio/visual equipment along with computers for the LRC classrooms before the process starts all over in the next fiscal year.

An important factor to consider when evaluating hardware to be used in labs is the cost of such hardware. The computer hardware purchased for the labs are similar in terms of processing power, memory, hard disk space, and graphics cards. The computer model may vary based on recap year and availability of the model. For example, GL-203 has 24 Dell Optiplex 1040 computers. Those computers have an i7 3.0 gigahertz (Ghz) dual core processor with 4 gigabytes of RAM and a 250-gigabyte hard drive. The computers are equipped with Nvidia graphics cards, which have one gigabyte of video RAM onboard. The manager of the LRC lab group generally purchases hardware considered “middle of the road” in terms of the range of technology offered by Dell to its customers and regarding the cost of the computer (Administrator H, personal communication, March 18, 2016). At the time of purchase, the computers purchased during the lab recap cycle cost approximately \$2,000 each (Administrator H, personal communication, March 18, 2016). In that year, enough funding was available for five labs. The same model computer was purchased for all with the exception of the form factor (size).

The Dell computers are comprised of similar hardware configuration models with the exception of two LRCs, which have a need for more robust computers due to the processor and graphics intensive requirements set by the software programs used in those labs. These two LRCs are part of GSEAS whose service requirements were discussed previously. For example, the Watkins lab has 36 Dell Precision 7910 computers. These computers use two Xeon 4 core processors running at 2.4 Ghz. They come with 32 gigabytes of RAM and a NVIDIA graphics card that has two gigabytes of onboard video RAM. This model was ordered with 500-gigabyte SSD drives. Each station has dual 21-inch monitors for the students use. The deciding factor for the purchase of this computer model was the software applications utilized in the lab. Software, such as Ansys,

SolidWorks, SolidEdge, and Matlab, often require processor and graphics intensive computing. On the Ansys Corporations website, the Precision computer used in the Watkins lab was highly recommended for running large-scale models efficiently and effectively. These computers cost approximately \$3,000 each. The average price of a LRC lab computer is assumed to be \$2,500 based on the information provided previously.

Another important criterion to consider is that of software costs. Note that the software programs used by students at the universities are not included in the calculations for software cost. When applicable, the operating system, Office products, anti-virus software, as well as the software and core supporting software are taken into account. In some cases, the Landesk software used to distribute the LRC image to the lab computers is considered as part of the cost.

Many software applications are commonly used in all the LRC labs at NPS. However, certain software application specific to the department or class are offered by that department. The software applications common to all LRCs are installed on the LRC lab image, which is referred to as the base lab image.

The software cost per seat per year was partially obtained by communication with Administrator I who is the IT specialist in charge of software licensing (Administrator I, personal communication, September 16, 2016). Administrator I provided the cost for the Windows operating system (including Microsoft Office software), which was shown to be \$26.39 per unit. He also provided the cost of Landesk for a year, which is \$77,000 for a 3,000 user license. The Landesk cost per unit was \$25.67. Administrator J, a NPS IT specialist who works on supporting anti-virus software, provided information regarding the cost for the anti-virus software. Administrator J provided information that revealed the cost per unit for Symantec anti-virus was \$1.00 per seat (Administrator J, personal communication, September 16, 2016).

When the desktop operating system price was added to the cost for Landesk and the cost for anti-virus software, the total was \$53.06 per seat. In essence, an operating system, such as Windows 7, is installed on a computer. The Windows updates are then

run and installed on the computer. Next, the anti-virus software and updates are installed. A host of other supporting applications (provided free of charge) are installed to support the applications to be installed. Lastly, the most frequently used applications among all the various departments at the school are installed, configured, and tested. However, as previously mentioned, the academic software applications are not considered part of the software cost. The software is then activated as required by the product. Finally, the latest updates are installed and then all programs within the software suite are tested. All these software components make up what is known as the LRC base image.

A factor that should be considered when evaluating IT systems for use in a production environment is manpower. Over the last several years, many changes have occurred due to the Department of Defense (DOD) budget cuts, end strength cuts, as well as DOD IT server consolidation directives. These changes have impacted the ITACS department in general and the lab support group specifically. Currently, the team has one dedicated full time lead lab manager. This lab manager is responsible for managing the 10 LRCs and approximately three dozen classrooms. These tasks include overseeing the building of the LRC base image, the building of one of the LRC lab images, and overseeing the building, configuring, and testing of the other four LRC images by other ITACS technicians. The 10 LRCs that are currently in use have five Windows 7 LRC images.

The first step in the process is the creation of the LRC base image. The software to be included was discussed in a previous section. The task of building, configuring, and testing the base lab image is performed by one IT technician and takes approximately 80 hours to complete (Administrator K, personal communication, March 30, 2016). The base image is then used as a foundation to build the five lab images, which contain specific software for various departments.

Fundamentally, the educational software programs required by a group of departments are installed onto a lab image. The lab image is then copied or deployed onto the computers, which are housed in LRCs located in the corresponding building. This task is accomplished to provide the software programs, which support the curricula being taught in the building, to faculty and students who use the programs in support of their

coursework. Students working on their theses, as well as dissertations, may also use these programs if their work requires. The task of creating five LRC images is complex and requires attention to detail. The five IT technicians use the base lab image, which takes two weeks to build. The five technicians each spend approximately 80 hours installing, configuring, and testing the software necessary for the five LRC lab images every six months (Administrator H, personal communication, March 18, 2016).

In addition to the lead lab manager, two PC shop members assist when the lab image creation and imaging functions must be performed. An IT specialist working for ITACS, who is assigned to an academic department as a dedicated resource, builds and maintains the WA lab image as part of his responsibilities (Administrator H, personal communication, March 18, 2016). Similarly, another ITACS IT specialist who is a dedicated resource to a department at the school is responsible for building and maintaining of the RO lab image (Administrator H, personal communication, March 18, 2016).

Once the LRC images have been completed, the next step in the process is to copy the images onto computers in the corresponding LRCs. To accomplish this task, the IT lab support group relied upon various available methods. These methods ranged from cloning using specialized software to using a disk-cloning machine, to using PXE boot server, and most recently, to using Landesk to deploy the images over the network to the target computers. For this research, the manpower involved using the latest method is considered.

To deploy the LRC images to the LRCs, an IT technician uses a software tool called Landesk. Once the LRC image is ready, the IT administrator is then able to deploy the software to many computers, usually one lab at a time, during off business hours. The task is usually scheduled to be performed overnight so that labor hours are not used for the task. The next day, the IT administrator must go to the lab to complete some post-processing tasks, which are necessary on the LRC computers. The LRC technician is required to log onto each computer in the lab to perform the remaining tasks. The tasks of scheduling the image for deployment, as well as completing post-processing tasks, requires approximately four hours of manpower per lab (Administrator K, personal

communication, March 30, 2016). The LRCs are re-imaged in groups of five LRCs biannually. Five LRCs are imaged every six months, as the number of LRCs is 10 (Administrator L, personal communication, April 5, 2016).

NPS has a helpdesk, as do most other universities. The NPS Helpdesk supports many diverse technologies and software, such as student, faculty, and staff office computers, laptops (both government owned and personal), MAC computers (government owned and personal), smart devices (both government owned and personal), operating system (Windows, MAC, and Linux) related issues, account related issues, common access card (CAC) and DOD certificate related issues among other tasks. The labor hour related maintenance in the LRCs is difficult to gauge, as the lab group is contacted through many different avenues. Also, as noted, the LRC labs require low maintenance once the quarter has begun. Most manpower hours for the NPS Helpdesk are spent on the other items listed previously; therefore, the LRC tickets are inconsequential. The NPS Helpdesk is not considered in the calculations for LRC manpower.

Another important criterion to consider when discussing computing environments is their availability. The computers housed in the LRCs are located in secure rooms made available to students at 0800. The doors to the labs are closed and locked at approximately 1700. Access to the labs can be obtained for after hour use if required. Resident students have the ability to walk into any LRC on campus and utilize the computers. However, software specific to a department is generally found on the LRC computers in the building where the department is located. The LRCs are closed during the re-imaging process that takes approximately four hours per year per lab. The labs are also unavailable when a power outage occurs. In the recent past, these outages have been scheduled events and have happened once per year. The labs are unavailable for the eight-hour duration.

This section described NPS. The computer labs at NPS, known as LRCs, were defined and their functional goals explained. The IT computing department, now known as ITACS, was discussed along with the groups contained within the department. The customers of the LRCs were shown to be resident students. The services provided by the LRCs are computing and software services. The lab computers provide the means for

students to utilize software, store data, and have access to the network and the Internet. The thick client computers used in the LRCs were considered for understanding the hardware costs. Although they range in process, the thick client computers cost approximately \$2,500 each.

The manpower associated with the maintenance of the equipment was discussed. The thick client solution used at NPS is labor intensive and time consuming. The task of building, configuring, and testing the LRC images is performed by various IT technicians and takes approximately many hours to complete. The tasks of scheduling the image for deployment, as well as completing post-processing tasks, requires approximately four hours of labor per lab.

Lastly, the availability of the computing environment was considered. The criteria take into account the up-time (also known as availability) of the computers during the year. Availability differs from accessibility, as the latter refers to a barrier in accessing the computer labs, such as hours of operation.

B. VDI COMPUTING AT NPS

In recent years, a new technology has become available at NPS and has been beneficial to students from all curricula. While supporting the LRC labs and classrooms at NPS, the researcher had received a request from an IS military faculty member for the Root Hall image that the researcher had worked to create. The faculty member went on to explain how he and another IS faculty member were creating a VDI infrastructure to support his department and classes. His goal was to convert the lab image into a virtual machine, which could then be cloned for use by students who would access those VMs with the use of client computers. This technology had not been implemented at NPS in any department, as it was still in its early stages as a technology. The pilot project has matured and due to the proof of concept, VDI is now being used at NPS to provide services to students. The researcher was made aware, by taking courses in the ITM curriculum, and in one instance by supporting the VDI system in an IT administrative capacity, of other VDI environments being used at the school to provide computing services to students. These use cases prompted the researcher to seek out the

administrators of these VDI private clouds to develop an understanding behind the migration to this technology versus using the existing LRCs.

Four IT administrators were interviewed at NPS with the goal of gaining insight into why the VDI technology was selected for use rather than using the traditional LRC. A set of 15 questions were asked that helped to gain information about the motivations for using the technology, as well as gaining an understanding of the cost of the systems, manpower involved, and the availability of the system. These metrics may be compared against the metrics pertaining to the existing LRCs, which were provided in the previous section.

Based on the replies from the administrators, three major reasons were the driving factors contributed to the migration to the VDI infrastructure as opposed to use of the LRCs. The first reason is computing services, which included software applications, were mostly limited to resident students through the use of LRCs. NPS had a growing DL component who needed access to computing services that were limited in reaching students. The second reason is the cutting edge curricula being taught at NPS demanded computing resources for classes that taught about topics, such as malware, cyber-attack and defense, and ethical hacking to name a few. These courses would require a separate network, as they would not be allowed onto the university network. Thirdly, LRCs had practical limitations, such as the man hours spent to maintain the LRCs and research labs, the physical space limitations at the school, the cost of maintaining the LRCs, accessibility, and the need for additional capabilities (such as virtual network computing (VNC)).

At NPS, of the four IT administrators interviewed, two of the four administrators teach courses at NPS. These administrators work in various departments and provide services to particular groups of students, faculty, and staff at the school. Administrator A works for the IS department. His customers are comprised of students, faculty, and staff from GSOIS, the defense analysis department, and select customers from the computer science department. Services provided are hands on development of actual VDI environments. Students work on racks, install servers and operating systems, and configure networks for the VDI environment. Services also include providing

applications and virtual machines for coursework. In some cases, the VMs provide software, which is different from software found in the LRCs.

Administrator B works for the cyber academic group. The students, faculty, and staff from the cyber systems and operation group and students from the computer science department use the VDI system, which he supports to meet the challenges of the curriculum. For the most part, the students are located at the NPS campus, but on occasion, he provides services to DL students. Services provided include virtual machines, which contain software for scanning and penetration testing, malware analysis, and attack and defend scenarios.

Administrator C works for the ITACS department at NPS. His customers include students, faculty, and staff, however, a large proportion of students are comprised of DL students located all over the world. Services provided are virtualized applications and virtual machines that contain a host of software required from various departments at NPS.

Administrator D works for the computer science department. His customers are comprised of resident and DL students. He teaches courses all over the world and his students use the VDI system that he administers at NPS to complete their homework, lab work, and projects. Services provided include virtual machines and a complete virtual environment where students can test the skills acquired through the ethical hacking course.

C. VDI INTERVIEW QUESTIONS AT NPS

The four interviews were conducted face-to-face and at the faculty/staff members' offices. They were conducted to gain perspectives on the reasons behind the migration to the VDI technology. They were also conducted to learn specific information about the cost of the solution, manpower needed to administer the environment, and lastly, the performance of the system measured availability. The following 15 questions were asked of VDI administrators:

- Who are your customers and what services do you provide to your customers?

- Why did you choose to implement a VDI solution rather than use the NPS LRC labs provided by ITACS or use ITACS Cloudlab?
- Do you experience outages where students cannot access the VDI environment? If yes, how frequently and for how long is the average downtime?
- How much time do you spend providing technical support to your customers? How much time do you spend on maintenance (software updates such as Microsoft updates, Java updates, etc.)?
- What if a hardware, software, or network outage occurs on Administrator D's VDI environment?
- Given the success of the VDI environments (ITACS Cloudlab, Administrator A's vLab, Administrator B's vLab, Administrator D's vLab) can we come up with a framework which supports a common architecture which supports resources (either dynamic or static) required for each lab/class?
- How long does it take to create a Windows image which includes software for classes in the VDI environment? How long does it take to duplicate that image to let us say 36 VMs?
- How many sessions (users concurrently logged into VMs) can you support?
- Are there any safeguards and mitigations in place for students using BYOD (such as Safeconnect on the ERN domain) to address vulnerabilities?
- How much does VDI hardware cost (servers, storage, switches, etc.)? How long before hardware needs to be replaced? Do you have a warranty period for hardware?
- How much does/did the VMWare ESXi software cost (ESXi operating system, licenses, etc.) initially and how much does it cost to maintain the VDI environment (ongoing fees)?
- What is the availability of LRCs? What is the availability of VDI? Can users VPN/RDP to LRCs?
- Are we still tied to 32 computers in a classroom? Or does VDI, which works well in Cloudlab and vLab environments, have the ability to extend into the LRC environment?

- How can LRCs leverage Zero clients/BYOD and VDI in order to provide the current computing LRC functionality (software, web access, storage space) to students?
- Do you have your own Active Directory environment? Is it required? Do you have a DHCP server? Do you use a DNS server and a web server in the VDI environment?

D. LAB COMPUTING AT THE NWC, CSUMB, AND STANFORD

In this section, information was gathered from IT professionals who work at the Naval War College, CSUMB, and Stanford University. The universities were selected as the researcher had contacts at those universities who either agreed to be interviewed or passed his request for interview to IT professionals who had knowledge of the lab/computing process. The IT professionals were interviewed via telephone conference, in one case, and in person for the other two cases. The objective of the interviews was to gain an understanding about the technologies used to provide computing resources/services to students for completing homework assignments, lab work, and research at the selected universities and why they selected those technologies. Other points of interest were to gauge how much the system cost, determine how much manpower is involved in maintenance, and to gauge the performance of the system measured in availability. The final goal was to determine if any of the technologies used at the three universities could be used as a possible alternative solution to the current thick client solution being implemented at NPS.

The first interview conducted was with Administrator E who is affiliated with the IT department at the Naval War College (NWC) in Newport, Rhode Island. The school's missions today are "developing strategic and operational leaders, helping the Chief of Naval Operations define the future Navy, strengthening maritime security cooperation and supporting combat readiness" as stated on the NWC website. When asked, "Who are your customers that use the technology (for example computer labs) and what services are provided"? Administrator E stated that the computing services, which they provide, are different from the computing services offered at NPS. Most of the services they provide are desktop services, which run Microsoft Office for word processing, spreadsheets, and presentation creation software among other such services

(Administrator E, personal communication, July 25, 2016). Administrator E informed the researcher that their customers are students, faculty, and staff at the NWC. In regards to the students, Administrator E conveyed, they are comprised of approximately 600 U.S. students enrolled full time, with about 200 international students, and approximately 600 faculty and 200 staff personnel (Administrator E, personal communication, July 25, 2016). The services provided are access to software applications, such as Microsoft Office and other Navy applications with a Zero client connecting to a VDI platform (Administrator E, personal communication, July 25, 2016).

The second interview was with Administrator F who is the Director of Technology Support Services at CSUMB. Educators and community leaders founded CSUMB in 1994. The campus is located on the former site of the Fort Ord military base. CSUMB provides education to students mostly from the entire state of California; however, one third of the undergraduate population comes from Monterey County and surrounding counties. CSUMB offers undergraduate and graduate degrees to students who attend. When asked about who are the customers that use the technology (for example, computer labs) and what services are provided, Administrator F replied that his customers were mostly resident students from all different programs at the school (personal communication, July 26, 2016). Approximately 7,000 students are enrolled at CSUMB. The services provided are software and computing/processing power (Administrator F, personal communication, July 26, 2016). Administrator F stated that the services are provided with thick clients located in labs across campus (personal communication, July 26, 2016).

The third interview was with Administrator G who is a High Performance Computing Systems Administrator at the Stanford Research Computing Center. The university is located at the heart of the Silicon Valley. According to its website, it is one of the world's leading teaching and research universities. Stanford University opened in 1891 and offers both undergraduate and graduate degrees to students. Administrator G was asked the following, "Who are the customers that use the technology (for example, computer labs) and what services are provided"? He stated that the customers were all students, faculty, and staff and Emeritus faculty at Stanford University. Anyone with a

sponsored student I.D. can log into its systems. He stated that the university has over 6,000 undergraduate students and approximately 9,000 graduate students in attendance. The services provided by the IT organization are basically software applications and computing/processing power.

The information collected in the interviews made it possible to gain perspective on which technologies are being used at the universities for providing computing resources to students and why those technologies were selected. The first interview, which was performed via telephone conference, showed the NWC is using the VDI infrastructure to provide computing to students, faculty, and staff at the university. They chose the VDI technology as it was less labor intensive.

The second interview, which was conducted face-to-face, showed that CSUMB was using the thick client architecture to provide services to resident students, faculty, and staff. This technology was selected, as it was the technology used at the time the labs were setup over 20 years ago. It is costly and time consuming to make changes to that model. Students, whether resident or DL (of which they have over 1,000), have the option of downloading software from the CSUMB portal. The software is either offered free of charge or at a discounted price.

The third interview, which was also conducted face-to-face, showed that Stanford University uses a combination of technologies to provide computing services to its user base. One such technology is BYOD, combined with the ability to connect to a Stanford web portal (via WiFi). This option provides students with software, which can be downloaded and installed onto student laptops or smart devices (Administrator G, personal communication, August 4, 2016). If the software is not available through the web portal, Stanford University provides a shared computing environment called FarmShare. FarmShare is a private cloud located in the datacenter at Stanford University, which provides hardware, and software services that run on servers (Administrator G, personal communication, August 4, 2016). The main reason why the university chose the private cloud model was that students were not working in physical labs anymore (Administrator G, personal communication, August 4, 2016). They wanted to be able to sit anywhere on campus and have access to computing services. Another reason why this

model was chosen was because “space is at a premium so if spaces are being underutilized and can be transformed into something more useful to our faculty, staff, and students then that will play a role in fulfilling our goals of providing the highest level of service to our students” explained Administrator G.

Another computing option made available to students is the Dell all in one computers, which are located at the Tech Center. The researcher located the lab while at Stanford University. The Dell computers are running Windows with numerous applications installed. The Dell computers have access to the Internet and the Stanford.edu network. The computers located at the Tech Center are not considered in the calculations in Table 1, shown in Chapter IV, as they are inconsequential.

E. INTERVIEW QUESTIONS AT NWC, CSUMB, AND STANFORD

The interview with the NWC was conducted via conference call from the researcher’s office while the interview with CSUMB was conducted face-to-face from the office of Administrator F. The interview with Stanford University was conducted face-to-face from a conference room located the library. The following 19 questions were asked of university IT professionals:

- How many students attend the University?
- Do you have computer labs on campus?
- Do you have redundancy in the computing environment which offers computing/software services to students?
- Does the University have distance learning students? If so, how many? How do they access/use the software programs, which supports their coursework?
- Who are the customers that use the technology (for example, computer labs) and what services are provided?
- Which technology do you use to provide students with the computing resources (hardware and software) to accomplish their homework, lab projects, thesis, and dissertations?
- Why did the IT department/University choose to use the current model/technology to provide students with computing/software resources?

- Are there any outages where students cannot access the computing environment? If yes, how frequently and for how long is the average downtime?
- How much time is spent providing technical support to customers on the technology?
- How much time/manpower is spent on maintenance (software updates, patches, installations) on the technology used to provide students with computing/software resources?
- How do you implement changes made to the services (hardware/software)? For example, at NPS, we use Landesk to deploy images to lab/LRC computers. How much time/manpower is spent on the task?
- How long does it take to create a Windows image, which includes software for classes using the technology?
- How long does it take to duplicate an image to 36 users/seats if applicable?
- How many users (users concurrently logged in to the system) can the system support?
- Are there any safeguards or mitigations in place for students using BYOD to address vulnerabilities?
- How much did the hardware cost? (Can the cost be broken down to 36 seats?) How long before the hardware needs to be replaced? Do you have a warranty period for hardware?
- How much did the software cost initially and how much does it cost to maintain on a yearly basis?
- What is the availability of labs and/or the technology used to provide computing services?
- Do you have your own Active Directory environment? Is it required? Do you have a DHCP server? Do you use a DNS server and a web server in the VDI environment?

The goal of determining which of the three technologies implemented at the universities can be a possibly implemented at NPS is discussed. CSUMB utilizes the thick client solution where “gold” Windows images are created and deployed to the thick client computers in the lab using Landesk. This solution is currently being implemented at NPS and is analyzed for comparison purposes. The thin client/VDI solution

implemented at the NWC should be considered. Currently, NPS has a VDI presence at the university. Three faculty members have designed and implemented three VDI infrastructures at the school for use by students, faculty, and staff in various departments. However, as these VDI infrastructures are supported by faculty, and are used by departments within NPS, the solutions are not included in the analysis found in Chapter V. ITACS has implemented an Enterprise VDI cloud made available to students (both resident and DL), faculty, and staff located in all departments at the university. The ITACS Enterprise VDI infrastructure is used as a case for analysis in Chapter V. The possibility as to whether the LRC computers can be combined with the ITACS private cloud infrastructure to provide computing services to students is investigated.

Chapter III described the current lab technology used at NPS. The existing LRC technology, known as the thick client computers, uses the client server architecture to provide computing services to students. The services provided to students were highlighted to outline goals that must be met for the possibility of replacing the existing technology with newer technologies. An estimate of the cost, manpower, and availability associated with maintaining the LRCs was provided. The information is used for comparison with the other technologies shown in this research. New technologies have been developed over the last 10–20 years, which have offered faculty the ability to present coursework not possible to accomplish with older technology. The technology also offered 24-hour access to computing resources, allowed DL students to access computing services, as well as provide access to specialized software not available to LRC users. These functions are made possible with the use of the VDI technology. A list of interview questions was provided, as they were asked to four NPS IT administrators to show services provided, the cost of the VDI solution, the manpower associated with maintaining the solution, and the availability related to each solution.

The chapter was complemented with an introduction to various IT computing technology examples used at three universities. Information was gathered about the IT infrastructure used at the Naval College (Newport, Rhode Island), CSUMB, and Stanford University to provide computing capabilities and software services to students. It shown if they use traditional labs on campus or if they use a different model to provide

computing services. A series of interview questions to be asked of the IT professionals at the aforementioned universities was presented in this chapter. Their responses are shown in Chapter IV. In addition, the questions help to reveal their customers, the lab services provided, an estimated cost of the technology, the manpower involved in maintaining the solution, and an estimate about the availability of the various technologies. The information may prove to be invaluable in providing options for NPS to deliver computing services using advanced technologies that may offer a lower cost and higher efficiency than the current LRC solution. An analysis of the extended capabilities offered by NPS, as well as the technologies used at NPS, the NWC, CSUMB, and Stanford University are presented in the next chapter.

IV. ANALYSIS OF VDI TECHNOLOGIES AT NPS AND IT TECHNOLOGIES AT THE NWC, CSUMB, AND STANFORD UNIVERSITY

VDI technologies used at NPS are presented by using data obtained by asking four NPS VDI administrators, who each manage a different VDI private cloud, 15 questions about the infrastructure that they administer. The answers to the questions helps to show what services are provided by the technology. An estimate of the hardware cost, the software cost, the manpower, and the availability associated with each VDI system is shown as well. The purpose is to provide criteria by which to compare the LRC services, the cost, the manpower, and the availability to those of the VDI systems.

The chapter also discusses various IT computing technologies used at the Naval College (Newport, Rhode Island), CSUMB, and Stanford University to provide computing capabilities and software services to students. It shows how their labs or equivalent services are configured and used to provide computing resources to students. The estimated hardware and software cost of the technology followed by the manpower involved in maintaining the solution and an estimate about the technologies availability is provided. This research includes interviews, either conducted over the phone or in person, with IT professionals employed at the various schools. The IT professionals were asked 19 questions and their replies are presented throughout the chapter followed by a table that calculates the cost per seat, as well as the availability of each IT solution implemented at the three universities. These resources may prove to be invaluable in providing options for NPS to deliver software to NPS customers with new technologies, which offer a lower cost and higher efficiency than the current solution.

A. VDI COMPUTING AT NPS—REASONS FOR MIGRATING

Based on the replies from the four VDI administrators, valuable insights were gained to the question of why they migrated to the VDI technology. The reasons included a growing DL component at the school, unique demands of the curricula being offered, and the practical limitations, imposed by the traditional LRCs. Their answers also revealed various aspects when compared to the NPS LRC model, such as the cost of

funding versus the cost for LRCs, the manpower involved in maintaining VDI systems as opposed to LRCs, and the availability of VDI systems versus LRCs.

B. DISTANCE LEARNING

The first topic to be discussed was the growing contingent of DL students. NPS was embracing an influx of students located all over the United States, and in some cases, in various parts of the world. These students were enrolling in various programs, one of which was the EMBA program. Also, the systems engineering department was catering to students located across the country, many of whom were located at military installations. Other departments, such as the electrical and computer engineering department, as well as the computer science department, were seeing DL students enrolling in their programs.

In all cases, it was apparent that the DL students needed a way to access computing resources/services so that they could complete their coursework, most of which utilizes costly software applications. NPS is able to receive educational discounts and volume licensing discounts for use on systems located at NPS, which is not available to students in most cases. Also, in some cases, a number of license agreements only allow certain software to be installed on government computers as opposed to student-owned laptops. This limitation would affect DL students, as most use their own laptops or computers to complete course assignments. Furthermore, with regards to residents students, not all had BYOD (particularly eight years ago and years prior) so the school was obligated to provide an option for computing resources (both hardware and software) to students.

To reinforce the point of growing DL enrollment, the VDI administrators were asked questions regarding their customers. When the interviewees were asked, “Who are your customers and what services are provided,” Administrators B, C, and D all indicated that they supported a mixture of resident and DL students. Administrator A informed the researcher that his customers consist of information sciences, defense analysis, and computer science students going through masters programs in various curricula (Administrator A, personal communication, May 26, 2016). Administrator B replied by

stating that his customers were all instructors in the computer systems and operations group, the cyber academic group, and the computer science department who want to use the virtual infrastructure (Administrator B, personal communication, June 8, 2016). Administrator C was quoted as saying, “NPS students, faculty, and staff.” The services provided through Cloudlab are virtual desktop so that students can do their class work and get to the NPS network from anywhere around the world 24 hours a day seven days a week. With the VDI technology, there is no such excuse as not enough time to get the assignment done” (personal communication, June 16, 2016). Administrator D teaches courses, as well as administers a VDI environment. He stated that his customers are local and international students (Administrator D, personal communication, June 20, 2016). He teaches a CS course that teaches the concepts of hacking and uses the VDI environment for those classes. The class is also taught from various locations around the globe. The students access the lab environment, located at NPS, from wherever they are and have access to it 24/7 no matter the time zone. Each class has roughly 30 resident or DL students and eight classes are taught per year.

Adding to the case for an alternative computing model for DL students was evident when the administrators were asked “Why did you choose to implement a VDI solution rather than use the NPS LRCs provided by ITACS or use ITACS Cloudlab?” Two of the VDI administrators made some interesting points about the topic. Administrator C provided information stating that the main reason for the use of the VDI technology was that DL students needed to access software programs to complete their schoolwork. These students are enrolled in various curricula and need access to different types of software. He stated that students are able to go on business or personal travel and can continue to work on their schoolwork. Furthermore, students can go home after school and continue to work on project from home by accessing the Cloudlab enterprise VDI environment. Administrator D mentioned a combination of reasons as to why he chose the VDI environment to compliment his off-site and international teaching objectives. Several key reasons are mentioned later in this section, but the point that fits the topic of discussion were DL students. He stated that first and foremost international students did not have access to LRCs. As stated, remote access to LRC computers is not

allowed. Due to that fact, when traveling to foreign countries, he would have to bring 35 laptops with him, which were preconfigured and ready for use. Or, he would have to rely on having the students use their laptops. However, an added hurdle was that it was unknown which operating systems were running on students' laptops, as some students would run Vista while others ran Ubuntu, etc. He needed a controlled environment with all the hacking software installed that all students could access.

C. EXPANDING CURRICULA REQUIREMENTS

When asked the question about why VDI was selected as an option to LRCs, a second yet compelling reason was due to the curricula being taught in several departments at NPS. Some curricula require software unavailable in the LRCs. Administrator A revealed that the courses he teaches and/or supports consist of labs where students get a hands-on experience on the development of the VDI environment (personal communication, May 26, 2016). Students work on racks, install servers and operating systems, configure networks and perform tasks to create a VDI environment (Administrator A, personal communication, May 26, 2016). He noted that some of these courses require applications and virtual machines configured differently or have slightly different software from that found in the LRC. For example, he teaches a big data analytics course that utilizes a Hadoop cluster and Cloudera software (Administrator A, personal communication, May 26, 2016). The virtual machines used in that course are Linux based with software and software agents that run on the Linux operating system (Administrator A, personal communication, May 26, 2016). The LRCs at NPS are Windows-based so Linux is not an option.

Other curricula require software that must be run in a secure environment, as it would pose a threat on the standard NPS network. Administrator B was quoted as saying that the courses he supports have labs that run cyber-attack and defend classes (Administrator B, personal communication, June 8, 2016). Due to the intrusive nature of those labs, faculty and staff wanted to host these labs on isolated networks and VDI was the easiest way to achieve that goal (Administrator B, personal communication, June 8, 2016). As with the previous case, they did not want to risk infecting the schools network.

Administrator B reiterated that the classes he supports consists of labs, which practice attack and defend exercises, malware exercises, scanning, network penetration testing among other exercises (Administrator B, personal communication, June 8, 2016). They wanted the ability to create real world networks and the best way to do so was virtually. In the VDI environment, which he supports, hundreds of VMs with dozens of independent subnets are all connected together (Administrator B, personal communication, June 8, 2016). This curricula requirement would not be allowed on the schools network.

Still another curriculum in the computer science department required a stand-alone network to support the course labs and final project. Administrator D was excited while answering the question why VDI was selected as the technology for use in his department, as he teaches courses, as well as being the administrator of the technology. He stated that he teaches an ethical hacking course and all the labs are done using the VDI environment (Administrator D, personal communication, June 20, 2016). Students have access to virtual machines based on both Windows and Linux operating systems and most of the tools required for the course are installed on those VMs. Students also have access to an entire virtual hacking environment on which they can work on their assignments and projects. This virtual environment is separate from the school's network, and therefore, poses no threat of infecting the campus computers (Administrator D, personal communication, June 20, 2016). The system will be used for another cyber related course, which will be taught to resident and DL students in the near future (Administrator D, personal communication, June 20, 2016).

D. PRACTICAL LIMITATIONS

A third category behind the reasons why the move to VDI was made was due to the practical limitations set by LRCs. At times, a use case was encountered that caused faculty to allocate funds that support a course taught once per year. In some cases, space constraints were the cause of contention. Administrator A stated that LRCs were not cost effective (personal communication, May 26, 2016). Instances occurred when faculty members wanted to use a piece of software that cost \$30,000 to license for a class of

students for a one-year period (Administrator A, personal communication, May 26, 2016). The professor would only use the software for one quarter out of the year and was unused for the remainder of the year. Part of Administrator A's funds would go to support cases as the one described (personal communication, May 26, 2016).

Falling under the practical limitations category was that of space constraints in the department. Since the LRCs in his building doubled as classrooms, conflicts resulted for users who wanted to use the lab versus the classes being taught throughout the day (Administrator A, personal communication, May 26, 2016). By using the VDI environment, the computers were removed from the LRC lab room and tables (which had power and Ethernet ports built in) were installed (Administrator A, personal communication, May 26, 2016). Students started to use the BYOD model and were thus able to plug in their laptops and get access to the software and tools they needed for the course they were taking via the VDI system without having to find a computer (Administrator A, personal communication, May 26, 2016). When the LRCs were occupied, students who wanted to use the system could go to a number of students' study areas or to teaching and research lab spaces to use the VDI system.

Administrator D explained that since LRCs were not allowed to host software for malware and for ethical hacking labs, he had to install hardware and software in teaching and research labs for use by resident students (Administrator D, personal communication, June 20, 2016). Those labs had to be re-imaged at least every quarter and many times more often due to the nature of the courses/labs (Administrator D, personal communication, June 20, 2016) That task alone was very time consuming and cumbersome. Also, he had to use funds to refresh the research hardware and pay into the LRC fund as well (Administrator D, personal communication, June 20, 2016).

Adding to the case for practical limitations of LRCs, Administrator D wanted additional features not present in LRCs. One such feature is being able to connect remotely to students' computer while they are still logged in (VNC onto a student's desktop) and see what the students were doing when they encountered problems (Administrator D, personal communication, June 20, 2016).

When asked about the accessibility of LRCs versus VDI, three of the four VDI administrators commented that LRCs being open from 0800–1700 (unless a keycard was obtained to access the lab after hours) was a factor in the decision of going to VDI. Another factor was that VDI was accessible 24/7/365 and gave students excellent access options from anywhere, which also showed the limitations of LRCs.

E. SERVICES, COST, MANPOWER, AND AVAILABILITY OF VDI

In addition to finding the reasons behind the migration to VDI, the questions asked of the VDI administrators revealed the services provided by the VDI environment, as well as the approximate hardware costs, the software costs, the manpower involved, and the availability of the VDI systems. The questions were asked so that the criteria and subsequent values could be compared with the LRC services offered, LRC hardware and software costs, LRC manpower, and LRC availability. They help to determine whether the VDI technology should be considered as a viable option for the LRCs at NPS, which is investigated in the next chapter.

As specified in an earlier section of the chapter, the interviewees were asked “Who are your customers and what services are provided”? The second part of the question addressed in this section are based on the interviewees’ replies and the researcher’s experience with administering two VDI solutions and attending the courses that utilized three of the four VDI infrastructure solutions discussed in the research. Computing services are provided to resident and DL students with the use of the VDI infrastructure. The services include VMs that provide Windows and Linux desktops. The VM desktops include software programs that can be used so that students can perform their class work, lab assignments, and thesis research. The VMs provide processing capability, as well as software functionality to accomplish coursework. Additional services provided are Internet access and storage. The VDI systems are accessible from anywhere in the world and at any time 24/7.

The VDI administrators were asked about cost of the VDI solution, “How much does VDI hardware cost (servers, storage, switches, etc...), how long before the hardware needs to be replaced and “How much did the VDI software cost (operating system,

licenses, etc.) initially and how much does it cost to maintain the VDI environment (ongoing fees)”? When asked about the hardware cost of the system, which Administrator A supports, he pointed out that a different business model and different undertaking exists in regards to LRC versus VDI. The LRC model is required to support the enterprise, while the model that he supports works well in a dynamic environment (Administrator A, personal communication, May 26, 2016). The system, which he supports, is able to change configurations often to support the next set of courses. He went on to say that the hardware costs \$80,000 in storage, \$30,000 in network equipment, and \$100,000 for servers (Administrator A, personal communication, May 26, 2016). When asked about software costs, Administrator A replied that \$30,000 was spent to purchase a license, which supports 100 concurrent users per cluster, and the room has two clusters (personal communication, May 26, 2016). An option exists to expand, as the current VDI system can support 300–400 users with the existing hardware (Administrator A, personal communication, May 26, 2016).

When asked about hardware and software costs, Administrator B was quoted as saying that he was not involved in the purchase of the system four years ago, but understands the hardware and software costs together were in the ballpark of \$500,000 four years ago (Administrator B, personal communication, June 8, 2016). He estimates that the hardware costs were approximately \$350,000, while the software costs were roughly \$150,000 (Administrator B, personal communication, June 8, 2016). Currently, all hardware and software are running without issue. Servers in the rack would cost \$50,000 to replace, but currently no plans are in the works to replace any equipment in the foreseeable future (Administrator B, personal communication, June 8, 2016).

Administrator C acknowledged that the hardware cost for the server and storage was approximately \$100,000 (personal communication, June 16, 2016). They use an all in one solution, which had benefits as opposed to engineering the system from scratch. The software cost was \$170,000 for a 400-user license (Administrator C, personal communication, June 16, 2016).

Administrator D stated that another group was excessing 40 of the 50 Dell 1950 servers, so those servers did not cost a penny (Administrator D, personal communication,

June 20, 2016). The higher end servers in the system cost \$10,000 per server and he is able to purchase 1–3 servers per year (Administrator D, personal communication, June 20, 2016). With the purchase of one high-end server, he is able to remove (life cycle them out) four Dell 1950 servers. The network equipment cost roughly \$20,000 (Administrator D, personal communication, June 20, 2016). The storage system cost approximately \$15,000 (Administrator D, personal communication, June 20, 2016). When asked about the software costs, Administrator D stated that ongoing fees are \$1,500–\$2,000 (Administrator D, personal communication, June 20, 2016). The total cost for VDI software is \$10,000 explained Administrator D (personal communication, June 20, 2016).

The topic of manpower needed to maintain the system was a point of interest, as it has been shown that the current economic conditions at NPS do not allow for an increase in manpower in the IT organization. A newer solution for LRCs is being researched that should require less manpower and/or a more efficient use of the existing manpower. Questions four and seven were asked to the VDI administrators as they pertained to manpower. The questions asked are how much time is spent providing technical support to your customers and how much time is spent on maintenance (software updates, such as Microsoft updates, Java updates, etc.)? Questions seven asked, “How long does it take to create a Windows image which includes software for classes in the VDI environment”? How long does it take to duplicate the image to 36 VMs? When asked about the labor involved in running and maintaining the VDI system, it should be noted that this task could be looked as having two parts. One part is patching and maintaining the servers, which run the VDI environment while the other part is the maintenance of the virtual machines that contain the software used in the courses.

Administrator A stated that the process of connecting to the system remains constant so they provide a static set of instructions to students. However, the accounts roll with students so he spends 3–4 hours at the beginning of the quarter refreshing VMs and creating new accounts (Administrator A, personal communication, May 26, 2016). He supports 50–100 customers per quarter (Administrator A, personal communication, May 26, 2016). The total time spent for customer service is 3–4 hours per quarter (Administrator A, personal communication, May 26, 2016). When asked about time spent

on maintenance, Administrator A replied that he has a Nessus server, as well as a Wuss server to perform updates automatically in addition to Java updates. The servers used for VDI are maintained more closely than VMs. The VMs are rebuilt every quarter so students use fresh ones (Administrator A, personal communication, May 26, 2016). They have templates, make changes to the templates, and redeploy from that newer template. He manually patches the VMware servers every week. It takes less than one hour to complete that task (Administrator A, personal communication, May 26, 2016). His reply for question seven was that it takes 3–4 hours to create a Windows VM that includes all software used in the course (i.e., MS Office, Visio, Project, wire-shark, WinSCP, puTTY, and network mapping tools) (Administrator A, personal communication, May 26, 2016). To duplicate the VMs to 36 VMs will take approximately 30 minutes (Administrator A, personal communication, May 26, 2016).

When Administrator B was asked about time spent on technical support and time spent on maintenance, he replied that it took several days at the end of every quarter (Administrator B, personal communication, June 8, 2016). This number includes time spent on creating new VMs and accounts for new students, as well as patching the servers (Administrator B, personal communication, June 8, 2016). When asked about the time it takes to create a new Windows image from scratch, he stated anywhere from 6–8 hours (Administrator B, personal communication, June 8, 2016). He also noted that the task of creating images is not something he does often. He informed me that the instructor starts with the .ISOs and creates the VM and then installs the software for the class. The instructors may also have the students create the class image rather than the instructor. To duplicate 36 VMs from one class VM, it takes anywhere from 20–40 minutes (Administrator B, personal communication, June 8, 2016).

Administrator C was quoted as saying that his department has a separation of duties; different engineers performing different functions support the enterprise Cloudlab (Administrator C, personal communication, June 16, 2016). A group of engineers provides the backend support and frontend customer support is handled by another group (Administrator C, personal communication, June 16, 2016). On average, the time spent on providing customer support is between 1–3 hours per week (Administrator C, personal

communication, June 16, 2016). When asked how much time is spent on server maintenance, he stated that an average of four hours are spent once a month while a few hours is spent on desktop VM maintenance at the end of each quarter (Administrator C, personal communication, June 16, 2016). On question seven, Administrator C replied that it takes several days to create a Windows image from scratch, as the standards for the images are more stringent for the enterprise (personal communication, June 16, 2016). Also, Cloudlab supports departments from all over the school, and as a result, many software packages must be installed, configured, and tested. Upon completion, it can take 30–45 minutes to create 36 VMs (Administrator C, personal communication, June 16, 2016).

When Administrator D was asked the questions contained in question 4, he responded that 10 hours per week are spent on technical support and on maintenance (Administrator D, personal communication, June 20, 2016). When asked question seven, he stated that it takes 2–3 hours to install the operating system and software required for the classes (Administrator D, personal communication, June 20, 2016). It only takes a few hours, as there is less software than the software installed on a LRC image (Administrator D, personal communication, June 20, 2016). Once the VM is setup, he can create 36 clones in five minutes (Administrator D, personal communication, June 20, 2016).

The final point made apparent in the interviews with the VDI administrators was that of performance measured as availability. The four administrators were asked, “Are there any outages where students cannot access the VDI environment? If yes, how frequently and for how long is the average downtime?”

Administrator A had been administrating the VDI system for eight years. In that timeframe, four outages occurred that ranged anywhere from four hours to two years (Administrator A, personal communication, May 26, 2016). These outages included power outages, WiFi upgrades, and network infrastructure upgrades.

Administrator B stated that at the end of every quarter, he performs a flush or clean out as he likes to call it (Administrator B, personal communication, June 8, 2016).

He recovers resources, changes passwords, and performs other IT related maintenance. The system is brought down for a few days for maintenance (Administrator B, personal communication, June 8, 2016). He also pointed out that with the View infrastructure, they have run into issues where a view client could not connect for whatever reason, which resulted in merely a 1–2 hour outage (Administrator B, personal communication, June 8, 2016). He further explained that the VDI system is a very robust product. He recalled once instance where it was necessary to relocate the server racks, which took a longer stretch, but it was planned and users were given notice.

Administrator C was direct in revealing that yes, sometimes there are hardware outages and sometimes there are power outages occur, and during those times, users cannot access Cloudfab (Administrator C, personal communication, June 16, 2016). Failure on hardware has happened. He went on to say downtimes of 1–2 hours for maintenance occurred roughly once a month (Administrator C, personal communication, June 16, 2016). On one occasion, the system was down between 2–3 days due to a technical issue (Administrator C, personal communication, June 16, 2016). However, for the most part, the system is up consistently.

Administrator D jokingly said that yes, system outages occur because the school keeps taking the power down. Interestingly enough, a power outage was scheduled the day of the interview. He stated that the only outages are power outages that have happened once or twice over the last two years (Administrator D, personal communication, June 20, 2016). It takes one day to power the system down gracefully and four days to bring it back up due to configuration and hardware errors (Administrator D, personal communication, June 20, 2016).

In this section, the VDI systems used at NPS for lab purposes prompted the researcher to seek out and interview four VDI administrators with the goal of discovering why the technology was implemented at the university. Three main reasons can be gathered, the first of which is that LRCs lacked support for DL students. DL students are not allowed to connect to LRC computer; hence, did not have complete access to the computing and software services that the LRCs offered. The second reason was the expanding curricula requirements. Many courses being taught at NPS included lab

exercises that pertained to malware use, cyber-attack and defense, ethical hacking principles, and other related coursework. The software to support these courses was not allowed on the LRC lab computers or the university network for that matter. Thirdly, LRCs had practical limitations that helped the departmental decision makers to seek out alternative technologies. Criteria, such as man-hours spent imaging research labs, physical space limitations in some departments, accessibility, as well as the LRC's lack of certain capabilities, were limiting factors. When considering a computing solution with which to provide computing and software services to customers, several criteria should be investigated. A set of questions were asked of the VDI administrators to quantify the criteria. The services required for meeting the mission should be considered. Other criteria, such as the system cost, the manpower involved in the maintenance of the system, and the availability of the system should be considered. The replies from the administrators provided information to gauge such values so that they may be compared to the values obtained for LRCs.

F. LAB COMPUTING AT THE NWC, CSUMB, AND STANFORD

Based on the answers to the 19 questions asked of the three IT professionals, valuable insight was gained about the technologies used at the three universities to provide computing and software services to students. Their answers also revealed various aspects when compared to the LRC model, such as the cost of funding for LRCs, the manpower involved in maintaining LRCs, and the availability of the various IT systems.

When discussing technologies for use at universities, the topic of cost is one of the driving factors when making a selection. Questions were asked to the interviewees regarding the cost of the computing solutions. The first set of questions that IT administrators were asked, question 16, "How much did the hardware cost, how long before the hardware needs to be replaced and do you have a warranty period for hardware? Another set of questions pertaining to cost were question 17, "How much did the software cost initially and how much does it cost to maintain on a yearly basis (ongoing fees)?"

When asked about the hardware cost of the system that the NWC uses, Administrator E stated that the VDI system is integrated as part of its infrastructure. It is difficult to break out a cost estimate for VDI hardware by itself. He approximated the hardware to cost between \$250,000–\$350,000 (Administrator E, personal communication, July 25, 2016). Administrator E also stated that the zero clients cost \$210 per client. He next addressed the question about how much it costs to maintain the VDI environment (ongoing fees). Administrator E explained that they have 900 user licenses with a yearly fee. It is necessary to buy a license and pay a maintenance fee. He referred the researcher to go to Navy ESL to get an estimate on how much 900 licenses would cost as he did not have the information available (Administrator E, personal communication, July 25, 2016). Fortunately, an ITACS VDI administrator at NPS had current cost estimates for the server and storage hardware used at NPS. For the analysis, a value of \$300,000 for the servers, storage, and switches used in the NWC VDI environment are used as it falls in the range provided by the Administrator E. The 900 VDI software licenses cost approximately \$500,000 (Administrator C, personal communication, August 1, 2016). The maintenance fee is approximately \$5,000 per year.

When asked about hardware and software costs, Administrator F from CSUMB first addressed the hardware costs. He stated they buy high-end thick client computers, as they are looking for longer life. With that said, they pay \$2,000 on average per computer, which come with a three-year warranty (Administrator F, personal communication, July 26, 2016). Those computers have i7 microprocessors and have 16 gigabytes of RAM with 24-inch monitors (Administrator F, personal communication, July 26, 2016). They do not buy extended warranties, as they purchase 20% more computers than they need so that they can keep them on hand. If a computer breaks and is out of warranty, he swaps it out with a new computer. Administrator F also stated that the school does have between 300–400 workstation computers that have higher requirements, so they cost \$3,000 per computer (Administrator F, personal communication, July 26, 2016). These computers are used in the science department as they have greater data needs. They use huge datasets so the computers are more robust. In regards to the software costs, they use

Landesk to push the images to the lab computers (Administrator F, personal communication, July 26, 2016).

When Administrator G from Stanford University was asked about how much does lab hardware cost, and in his groups case, it is called Farmshare. Administrator G replied, within Farmshare are three separate environments (Administrator G, personal communication, August 4, 2016). He went on to explain the “cardinal” machines are VMs intended for processes that run for days yet are not resource intensive (personal communication, August 4, 2016). The interactive machines that can be logged into are called “corn” computers (Administrator G, personal communication, August 4, 2016). Administrator G stated that Stanford University has 40 “corn” machines (personal communication, August 4, 2016). A subset of the “corn” machines is called “Rye” machines (Administrator G, personal communication, August 4, 2016). These computers are general purpose Ubuntu machines like the “corn” machines, only these have Nvidia GPUs (Administrator G, personal communication, August 4, 2016). Then, 20 MPI machines are called “barley” machines that cannot be logged into directly, but a “corn” machine can submit jobs to them (Administrator G, personal communication, August 4, 2016). The current hardware equipment has an estimated cost of \$250,000–\$300,000 including storage, servers, racks, and switches (Administrator G, personal communication, August 4, 2016). The equipment is purchased with a standard three-year warranty (Administrator G, personal communication, August 4, 2016). Most hardware is replaced when it reaches the three-year mark. The question regarding software costs did not apply to this computing environment.

As in the case for VDI administration, the topic of manpower needed to maintain the system was a point of interest, as the current economic conditions at NPS do not allow for an increase in manpower in the IT organization. A more robust solution for LRCs is being researched that should require less manpower and a more efficient use of the existing manpower. A series of questions were asked to attempt to gauge the manpower associated with using the technology at each university. Questions 9–13 were asked to the IT professionals as they pertained to manpower. The following sets of questions were asked:

- How much time is spent providing technical support to customers
- How much time/manpower is spent on maintenance (software updates, patches, installations) on the technology used to provide students with computing/software resources
- How do you implement changes made to the services (hardware/software)? For example, at NPS, we use Ghost or Landesk to deploy images and how much time/manpower is spent on the task
- How long does it take to create a Windows image which includes software for classes using the technology
- How long does it take to duplicate an image to 36 users/seats if applicable?

When asked about the labor involved in running and maintaining the various IT systems, Administrator E from the NWC responded with the following information when asked question 9. Basically, the NWC Helpdesk is dedicated to providing technical support to customers all day long (Administrator E, personal communication, July 25, 2016). Administrator E added that they have 14 technicians who include asset management and audio/visual functions. Ten technicians provide support on the VDI technology other computing related issues at the Helpdesk (Administrator E, personal communication, July 25, 2016). Approximately half their day is spent on VDI related support issues (Administrator E, personal communication, July 25, 2016). When asked question 10, Administrator E stated that four technicians perform maintenance twice a month, and at each maintenance window, they spend approximately four hours (personal communication, July 25, 2016). By maintenance, he clarified that he meant they perform application updates, Microsoft updates, and other requirements. The maintenance is done in parallel to the systems running. All students have an entitlement. During the maintenance windows, the student entitlements are removed if they are logged on (Administrator E, personal communication, July 25, 2016). The IT specialist creates a new pool and gives the students entitlements to the new pool. Therefore, the next time the students log on, they will get the new VM with the latest updates. For question 11, Administrator E informed the researcher that the changes are implemented with the use of a VMware Horizon administrator. This person uses the tool to recompose the pool,

which takes approximately four hours, but once the pool is composed, it is available to all users immediately (Administrator E, personal communication, July 25, 2016).

Concerning question 12, Administrator E explained that it takes 6–8 weeks to get the image working right, as it has to be compliant and everything has to be fully tested (personal communication, July 25, 2016). For this task, it takes two IT specialists working in the datacenter for an extended period of time before the image is ready for release (Administrator E, personal communication, July 25, 2016). He made a distinction about the fact that they have specialty software that customers use. Those pieces of software are relatively expensive so they roll those programs out as thin app'd applications and allow access to the users who need the applications as opposed to making it part of the base image (Administrator E, personal communication, July 25, 2016). The application is rolled out separately from the image, but it shows up as an icon on the desktop as if it were installed. When asked questions 13, Administrator E simply stated it would take minutes to spin up 36 VMs (Administrator E, personal communication, July 25, 2016).

Administrator F of CSUMB was asked question 9, how much time is spent providing technical support to customers? His reply was that 90% of their time is spent on planning and setup (Administrator F, personal communication, July 26, 2016). They meet with faculty members to see what software is needed. They then build and test the load sets or images as they are also referred to as. For those 1,000 lab computers, he has four full-time staff working on that task (Administrator F, personal communication, July 26, 2016). Once the semester starts, it takes minimal support to keep the labs up and running. When asked question 10, Administrator F explained that once the labs are setup and running, approximately 20 hours per week are spent on maintenance (software updates, patches, and installations) (personal communication, July 26, 2016). In regards to question 11, Administrator F quickly replied that they use Landesk for their imaging (Administrator F, personal communication, July 26, 2016). They push the images out to the computers using that software. He said that there is always a straggler or two when they push the images out to a lab, so the IT technicians have to be present to complete the task (personal communication, July 26, 2016). Also, they use a software called

Deepfreeze that freezes a machine in the current state. During the term, they thaw the machines out, which allows them to receive updates and then they freeze them again. Pushing out the images can take approximately two weeks (10 workdays) at which equates to (80 hours x four IT Techs) (Administrator F, personal communication, July 26, 2016).

When asked question 12, which pertains to the amount of time it takes to create a Windows image, the interviewee answered that it takes 16 hours to create and test a base image (Administrator F, personal communication, July 26, 2016). However, adding the software takes time. For example, Microsoft Office 2016 has dozens of decisions that need to be reviewed and made during the installation process (Administrator F, personal communication, July 26, 2016). Depending on the image, it can take weeks. Administrator F went on to say that it takes two months to build and test the images, as they have 10 different load sets/images (personal communication, July 26, 2016). They also install specialty software. These images are deployed to those 1,000 machines and then the post-configuration process ensures they can print (Administrator F, personal communication, July 26, 2016). Administrator F's reply to question 13 was that it can take 2–4 hours to deploy an image to 36 computers including post-configuration (Administrator F, personal communication, July 26, 2016).

Administrator G of Stanford University was asked the five questions pertaining to manpower to which he replied that it is a team effort. When asked question nine about how much time it takes to provide technical support to customers, Administrator G replied they have a Helpdesk that fields questions pertaining to the different computing architectures (personal communication, August 4, 2016). The general helpdesk department contain tier1 and tier2 support specialists (Administrator G, personal communication, August 4, 2016). Administrator G added that his team is considered tier3 due to the server systems they maintain (personal communication, August 4, 2016). Stanford has an IT support phone line, and a HelpSue web-ticketing system (Administrator G, personal communication, August 4, 2016). Since the topic of discussion is the necessary manpower in providing computing resources so that students can complete their school work, and since an IT administrator from the research

computing group was interviewed, the manpower considered is provided for the technology used to provide services through the Farmshare private cloud technology. Most tickets for the research computing team come in through the web-ticketing system and are routed to his department if they pertain to any of the systems that his team administers. All team members provide support, as members have their own areas of responsibility. Administrator G is responsible for Farmshare (Administrator G, personal communication, August 4, 2016). Time spent providing support on Farmshare is approximately 20% of his time (Administrator G, personal communication, August 4, 2016). Question 10 allowed Administrator G to show that roughly 10% of his time is spent performing maintenance on servers (personal communication, August 4, 2016). They have quarterly patching and are also responsible for security requirements patches in addition to standard patching (Administrator G, personal communication, August 4, 2016). When asked question 11, Administrator G paused before replying to the question. He does install software requested by faculty and makes it available to students. When a new release of the software is made available, the research computing (RC) specialist removes the previous version and installs, configures, and tests the latest version. Roughly 10% of his time is spent performing these tasks (Administrator G, personal communication, August 4, 2016). When asked questions 12 and 13, Administrator G replied that these questions are not applicable to the Cloud technology that he manages (personal communication, August 4, 2016). They do not use Windows images to provide computing services to students.

An important criterion that should be considered when providing computing services to customers is the performance metric known as availability. A series of questions asked to the three interviewees helped provide an idea regarding the reliability of the services provided to students with the use of the technology implemented at the various universities. The three IT Professionals were asked question 3, Do you have redundancy in the computing environment which offers computing/software services to students? Question 8, “Are there any outages where students cannot access the computing environment? If yes, how frequently and for how long is the average downtime”? Question 14, How many users (users concurrently logged into the system) can the system

support? Question 18, What is the availability of the labs and/or technology used to provide computing services?

Administrator E from the NWC answered with an emphatic “definitely” when asked if they have redundancy in the computing environment (personal communication, July 25, 2016). He went on to say that the infrastructure is built in a high availability model. If they have to swap a blade or replace storage, they can do that live and the users will not notice (Administrator E, personal communication, July 25, 2016). The system will be available. When asked about outages, Administrator E explained that no major outages have occurred. Whenever work needs to be done, such as a switch replacement and those types of things, they are done during off hours. They have enough redundancy built into the system that users will still have access during those times. He went on to say that if an outage were to occur, they would announce it to users well in advance. Administrator E stated, “For the most part students will always have access to the system while they are on campus”! (personal communication, July 25, 2016). The researcher asked how many sessions (users concurrently logged onto the system) can the system support? Without taking any time at all, Administrator E said that they have a 900-user license so they could easily support 900 users logged into the system. Administrator E added that the hardware could support many more users than that. In regards to the availability of the labs and/or technology, it was conveyed that if the labs are not booked for use by faculty or staff then they are available 24/7 as the campus is open 24/7 (Administrator E, personal communication, July 25, 2016).

Administrator F from CSUMB was asked about the up-time regarding the computer labs at the university. He was quoted as saying that the university has no additional capacity, so every single computer lab they have is booked (Administrator F, personal communication, July 26, 2016). The university does not have spare space or labs. If a lab were flooded, Administrator F went on to say, then we would not have the capacity to move it to another room (personal communication, July 26, 2016). They work at a state institution for which every square inch is accounted. When asked if any outages occur, and if yes, how frequently do they occur and for how long, Administrator F stated he could not recall an instance when the campus was open and a computer lab was not

(personal communication, July 26, 2016). He gave an example by stating that currently, part of the campus is without power, but other labs are open. If the university does not have power, then the campus closes. He pointed out that scenario is rare. The infrastructure is extremely solid as Administrator F stated, and no issues have been detected with the hardware. His reply to question 14 was that they have 1,000 computers in the labs so they can support 1,000 users concurrently logged in (Administrator F, personal communication, July 26, 2016). When asked about the availability of the labs, Administrator F explained that the availability of labs is from 8:00 a.m. to 10:00 p.m. (Administrator F, personal communication, July 26, 2016). The library has 60 computers and those are available 18 hours per day. The Café has 20 machines that are ready for use 24/7. If students want to meet for working on a project, they have options.

“Yes, we have redundancy on the equipment which we use to provide users with computing services” was the reply provided by Administrator G from Stanford University (personal communication, August 4, 2016). When question 8 regarding outages was asked, Administrator G replied that yes, they had two outages over the last two years. He informed the researcher that one outage might have impacted users and lasted a few hours while the other outage lasted less than an hour (personal communication, August 4, 2016). The longer outage was due to a server going down because of a hardware issue (Administrator G, personal communication, August 4, 2016). The server in question had to be replaced and rebuilt. At that time, some services were interrupted, but if students had the client installed, they were able to use that application to access the system (Administrator G, personal communication, August 4, 2016). In theory, they should have been able to access everything even though a server was unavailable explained Administrator G (personal communication, August 4, 2016). Since then actions were taken to provide redundancy.

When asked question 14 about concurrent users logged onto their service, his reply conveyed that at any given time, over the past year, there has been 5,500 unique user logins (Administrator G, personal communication, August 4, 2016). Therefore, at any given time, about 800–1,000 are users logged into the Farmshare system

(Administrator G, personal communication, August 4, 2016). He went on to say the hardware can support many more users.

Lastly, Administrator G commented that the computing services that he and his team manage are available 24/7/365 (personal communication, August 4, 2016). Quarterly maintenance is scheduled and conveyed to users ahead of time (Administrator G, personal communication, August 4, 2016). The system has enough redundancy so that users should not be impacted. If one of the Korn machines goes down in the middle of the night, then users are expected to use one of the other 64 available until it can be brought back up the next day (Administrator G, personal communication, August 4, 2016).

In this section, three IT professionals who work at the Naval War College, CSUMB, and Stanford University were interviewed. The goal of the interviews was to gain an understanding about the technologies used to provide computing resources/services to students for the purpose of completing homework assignments, lab work, and research at the selected universities and why those technologies were selected.

From the interviewees statements, it was conveyed that the NWC uses thin clients and VDI infrastructure to provide computing services. CSUMB utilizes the tried and true thick client technology to provide students, faculty, and staff with computing services in labs. At Stanford University, various technologies are offered to students, faculty, and staff. Software is offered through the ESS (Essential Stanford Software) portal free of charge or at a discounted rate. Thick client computers are available at the tech center. Those computers are thick client computers running Windows. A plethora of software is available on those computers. Also, a private cloud service known as Farmshare is hosted at the Stanford datacenter. The private cloud service replaced the thick client computers that were installed in the traditional labs in years prior. Many licensed software applications are offered through the Farmshare service.

An import criterion to consider when evaluating technologies with which to provide computing services is that of cost. Several interview questions were aimed at gauging how much the IT system cost at each university. Both hardware and software pertaining to the lab technology in use at the four universities were considered. At the

NWC, the hardware cost was difficult to estimate. However, an ITACS VDI administrator at NPS had current cost estimates for the computing and storage equipment, as well as the software license costs used at in the VDI environment at NPS. Those estimates were used to determine pricing for the NWC. At CSUMB, a value of \$2,500 is used for the average cost of a thick client computer. They have 1,000 computers throughout the various labs on campus. In the computing environment at Stanford University, the private cloud infrastructure was estimated to cost \$300,000. The equipment has a lifespan of three years. The cost of the small group of thick clients were not included in the calculations.

When considering a computing solution with which to provide computing services to customers, the manpower associated with the maintenance of the equipment must be taken into account. At the NWC, it was shown that the predominant model used is thin clients connected to a VDI infrastructure. Ten IT technicians provide computing support at the Helpdesk. However, the support provided covers various technologies and not just the VDI environment. Approximately 10 technicians provide four hours/day each on thin client/VDI issues totaling 40 hours per week. Maintenance performed twice a month totals 32 hours per month. For image creation, it takes (two techs x eight weeks) 640 hours per year. The manpower involved in spinning up 36 VMs is one technician taking 10 minutes. At CSUMB, the manpower involved in supporting the thick client lab environment for 1,000 computers is intensive in preparation and deployment time. Technical support on the labs is minimal. Ten percent of the time for four IT technicians is equivalent to 16 hours per week for tech support. The time spent on maintenance is 20 hours per week. Approximately 320 hours are spent deploying lab images using Landesk to 1,000 computers per year. It would require four hours to deploy an image to 36 thick client computers. In the case of Stanford University, the IT technician spends eight hours per week providing technical support. The task of performing maintenance requires approximately four hours per week. Installing, testing, and configuring software requires roughly four hours dedicated to this task.

The final topic of interest is system availability. At the NWC, it was discovered that they have high availability in their computing model. No system outages have been

reported. A total of 900 users can be logged into the system concurrently although the hardware is able to support many more users. The computing environment is available 24 hours per day seven days per week. At CSUMB, redundancy is not built in, as space on campus is limited. No outages have occurred that resulted in labs being unavailable. The IT solution at CSUMB is able to support 1,000 concurrent users logged onto the system. The lab computers are available 24 hours a day, but access is restricted to 8:00 a.m.–10:00 p.m. In the case of Stanford University, the computing environment does have redundancy built in. Over the last two years, two outages have occurred. One outage lasted four hours while the other lasted one hour. They system can support 1,000 users concurrently logged in. Farmshare is available 24 hours per day seven days per week 365 days per year.

Table 1 provides calculations that show the cost per seat per year that each university pays to provide computing lab services to students so that they can complete their homework, labs, presentations, and analysis to meet their educational requirements. The platform used at each university is shown in the first row. The table provides a description of the computing services provided by the IT lab solution at each university in the second row. The next row shows how the availability of the system is calculated, as well as the total value for each university. A list of lab hardware cost is shown in the fourth row followed by software cost for the IT lab solution. The manpower associated with each solution is calculated based on the time it takes in hours to maintain the system multiplied by the cost per labor hour of an IT Specialist. Lastly, the total cost per seat per year is shown. A detailed description of the values used to determine the cost per seat per year is included in the following sections. The values obtained are discussed in the next section of this chapter.

Table 1. University IT Solutions Compared Against Operational Criteria and IT Cost (Hardware, Software, and Manpower).

	NPS	NWC	CSUMB	Stanford	NPS Cloudlab (VDI Enterprise Solution)
<i>IT Platforms</i>	Thick Client	Zero Client VDI Infrastructure	Thick Client	Private Cloud	VDI Infrastructure
<i>Computing Services</i>	LRC computers provide software applications, store data, network and Internet access; hardware allows software and simulation programs to run.	Desktop computing service, which run Microsoft office for word processing, spreadsheets, presentations, as well as other such services.	Lab computers located in various labs across the CSUMB. The services provided are software and computing power.	Private cloud offering that provides software applications and computing/processing power.	Desktop computing service runs MS office & dozens of software programs; supports resident & DL curricula
<i>Availability</i>	Lab unavailable = 4 hrs. (Image deployment) + 8 hrs. (power outage) = 12 hrs. TOTAL = 12 hrs. per year; 365 x 24 hrs./day = 8,760 hours in a year – 12 hours downtime; 8760 – 12 = 8748; 8748/8760 x 100 = 99.86%	High availability 24/7/365 days per year. Greater than 99.9% availability	Computers are available 24 hrs. x 365 days per week = 8760 hours per year; labs unavailable 4 hrs. per year 8760 - 4 = 8756; 8756/8760 x 100 = 99.95% availability	Greater than 99.9% availability	Unavailable maint. = 24 hrs. /year Recompose – 6 months = 4 hrs. = 8 hrs. /year. Power Outage = 24 hrs. /year; TOTAL = 56 hrs. /year – 8760 hrs. in a year – 56 hrs. = 8704; 8704/8760 x 100 = 99.3%

	NPS	NWC	CSUMB	Stanford	NPS Cloudlab (VDI Enterprise Solution)
Hardware Cost	Client computer = \$2500.00/3 year warranty = \$833.33	Zero client = \$210 Hardware (servers, storage, and switches) = \$300,000 / replaced ever 3yrs = \$100,000 per year - 900 seats = \$111.00 /seat per year - Total = \$210 + \$111 = \$321.00 /seat per year	Client computer = \$2500.00/3 year warranty = \$833.33	Servers & storage = \$300K; \$300,000/3 = \$100,000 per year/ 5500 = \$18.18 TOTAL = \$18.18 /seat per year	Servers & Storage = \$100,000; \$100,000 /3 = \$33,333 per year / 400 = \$83.33 TOTAL = \$83.33 /seat per year
Software Cost	Desktop OS + MS Office = \$26.39 /seat Landesk = \$77,000/3,000= \$25.67 /seat Symantec A/V = \$1.00 /seat TOTAL=\$26.39 + \$25.67 + \$1.00 = \$53.06 /seat per year	900 user licenses = \$500K; \$500K/3 yrs. = \$166,667 /yr. Maintenance fee = \$5K /yr. - \$166,667 + \$5,000 = \$171,667 / 900 = \$190.74 /seat -Desktop OS & MS Office = \$26.39; Anti-virus = \$1 Total = \$190.74 + \$26.39 + \$1 = \$217.39 /seat per year	Desktop OS + MS Office = \$26.39 /seat Landesk = \$77,000/3,000= \$25.67 /seat Symantec A/V = \$1.00 /seat TOTAL=\$26.39 + \$25.67 + \$1.00 = \$53.06 /seat per year	No additional cost to access Farmshare; users utilize BYOD software or software downloaded from software portal TOTAL=\$0	400 user license = \$170K/3 yrs. = \$56,667/400 = \$141.67 Desktop OS & MS Office = \$26.39 A/V = \$1.00 /seat; Desktop Total=\$53.06 TOTAL = \$141.67 + \$26.39 + \$1 = \$169.06 /seat per year
Manpower involved	Base Image Creation = 80 hours; Each lab image = 80 hours TOTAL=Base+Lab1+Lab2+Lab3+Lab4+Lab5 = 480 hours deployment = 4 hours x 5 labs = 20 hours TOTAL=480 hours + 20 hours = 500 labor hours every 6 months 500 labor x 2 (twice a year) = 1000 x \$30/hr. = \$30,000 / 225 = \$133.34 /seat per year	Maintenance backend & front end – 32 hours per month x 12 mo. = 384 hr. x \$30 = \$11520/900 = \$12.80 – Image – 2 techs x 4 weeks = 640 hours per year x \$30/hr. = \$19,200 /900 users = \$21.33 /seat per year – HelpDesk – 10	Base image = 16 hrs. – Lab image per Tech = 2 months = 320 labor hrs. – 320 hrs. x 4 techs = 1280 hrs. – (16 hrs. + 1280 hrs. =1296) x \$25 per	Technical support = 20% (160 hours per month x .20 = 32 hours per month) Maintenance = 10% (160 x .10 = 16 hours per month) Software installation = 10% (160 x .10 = 16 hours per	Backend = 4 hrs. /month; front end = 8 hours / month(customer support) = 144 / year; Desktop maintenance = 4 hrs. per quarter (4 x 4 =16 hrs.

	NPS	NWC	CSUMB	Stanford	NPS Cloudlab (VDI Enterprise Solution)
		<p>techs x 4 hrs./day =200 hrs. per week; 200 hrs. x 4 weeks = 800/month x 12 months= 9600 hrs. x \$30 per labor hr. = \$28800/ 900 seats = \$320.00 /seat per year - Total = \$12.80+\$21.33+\$320.0=\$354.13 /seat per year</p>	<p>hr. / 1000 = \$32.04 /seat per yr. Deployment = 80 hrs. x 4 IT techs = 320 hours /year TOTAL= 320 labor x \$25 = \$8,000 / 1000 = \$8.00 /computer – Maintenance = 80 hrs. /month – 960 /year x \$25 / hr. /1000 = \$24.00 Total = \$32.04 + \$8.00 + \$24.00 = \$64.04 /seat per year</p>	<p>month) --- (32 +16+16 =64 hrs. / month) x 12 months = 768 hrs. x \$60 per hr. = \$46,080 /yr. in labor / 5500 users /yr. = \$8.38 /seat per year</p>	<p>per year); Base Image = 80 hours & 40 hrs. for lab image = 120 hours every 6 months = 240 per year. Total labor spent is 144 + 16 + 240 = 400 hours x \$30 per hour = \$12,000 / 400 = \$30 /seat per year</p>
Total Cost Per Seat Per Year	$\$833.33 + \$53.06 + \$133.34 = \1019.73	$\$321 + \$217.39 + \$354.13 = \892.52	$\$833.33 + \$53.06 + \$64.04 = \950.43	$\$18.18 + \$8.38 = \$26.56$	$\$83.33 + \$169.06 + \$30 = \282.39

1. NPS

At NPS, the lab computing services are provided by using thick client computers. The services provided give students the ability to utilize software, store data, access the network and Internet, and run software and simulation programs. Thick clients are standalone computers; however, they must be connected to the NPS network, as many applications require a license from the NPS license server to run.

a. Hardware

The hardware cost for a thick client computer ranges from \$2,000 to \$3,000 (Administrator H, personal communication, July 11, 2016). Administrator H is the helpdesk manager at NPS. The LRC lab function falls under him so he is responsible for ordering the equipment needed for the LRCs. The value of \$2500 was used, as it is the average price of the thick client computer. The thick client computers are replaced every three years (Administrator H, personal communication, July 11, 2016). Therefore, the \$2,500 is divided by three to obtain the hardware cost per year (see Table 1).

b. Software

The software cost per seat per year was partially obtained by communication with Administrator I who is the IT specialist in charge of software licensing (Administrator I, personal communication, September 16, 2016). He provided the cost for the Windows operating system (including Office products), which was shown to be \$26.39 per unit. Administrator I also provided the cost of Landesk for a year was shown to be \$77,000 for a 3,000-user license (personal communication, September 16, 2016). The Landesk cost per unit was \$25.67. Administrator J, a NPS IT specialist who works on supporting anti-virus software, provided information regarding the cost for the anti-virus software (personal communication, September 16, 2016). Administrator J provided information that revealed the cost per unit for Symantec anti-virus was \$1.00 per seat (personal communication, September 16, 2016). When the desktop operating system price was added to the cost for Landesk and the cost for anti-virus software, the total was \$53.06 per seat (see Table 1).

c. Manpower

Concerning manpower, several sources were asked to provide input. Personal communications with Administrator K showed that it takes 80 hours to create the base image (personal communication, July 19, 2016). Administrator L, a NPS IT specialist, stated that creating the lab image that he is responsible for takes 80 hours (Administrator L, personal communication, July 19, 2016). On July 20, 2016, Administrator M, a NPS IT specialist, conveyed that it takes 80 hours to create the lab image that he was responsible for creating (personal communication, July 20, 2016). Administrator N, NPS IT specialist, stated that creating the lab image he maintains takes 80 hours (Administrator N, personal communication, July 22, 2016). Based on experience, the researcher, a NPS IT specialist, has found the image he maintains takes 80 hours to build. Personal communications on July 22, 2016 with Administrator H, manager of the helpdesk and lab support group, informed the researcher that it takes 80 hours to create the fifth lab image (Administrator H, personal communication, July 22, 2016). The manpower is calculated by taking the labor hours for the base lab image and then adding the value to 80 hours per lab image multiplied by five for a total of 480 labor hours (see Table 1). Personal communications with Administrator K, a NPS IT specialist, showed that it takes four hours to deploy an image using Landesk (personal communication, July 19, 2016). The four-hour timeframe includes post-processing of all the computers in the lab. The maintenance performed once the quarter has started is minimal and is not considered, as it is inconsequential. Five labs are imaged every six months, so the four hours that it takes to deploy an image is multiplied by five labs totaling 20 labor hours. The 480 labor hours that it takes to build the images is added to 20 labor hours (time spent deploying the images and post processing) totaling 500 labor hours every six months. That value is multiplied by two, as this process is performed twice a year. The total labor hours spent for the lab process is 1,000 labor hours (see Table 1). Per discussions with Administrator H, the lab support technicians fall in the GS-09 pay scale. The GS-09 pay scale range for 2016 is \$58,132 to \$75,567 (per www.opm.gov/policy-data-oversight/pay-leave/salaries-wages/salary-tables/16Tables/html/SF.aspx). The dollar per hour rate was selected to be \$30.00 per hour, as this rate falls in the pay scale for GS-

09 IT technicians in Monterey (San Francisco scale). The 1,000 labor hours is multiplied by \$30 and divided by the number of labs seats (225) to obtain the value of \$133.34 per seat per year.

d. Availability

Next, the availability of the thick client lab solution was calculated. The normal business hours for the LRC labs are 0800–1700. However, any NPS resident student can obtain after-hours access by obtaining permission from the chair's department; thereby, making the lab computers available 24 hours a day. The lab is unavailable for four hours during lab imaging that occurs once per year per lab (Administrator L, personal communication, July 19, 2016). A scheduled power outage (lasting eight hours) occurs approximately once per year (Administrator H, personal communication, July 11, 2016). The labs are unavailable 12 hours per year. The year has 8,760 hours (365 x 24). Subtract 12 hours that the labs are unavailable from 8,760, which equal 8,748, and divide that value by 8,760 multiplied by 100. The labs are available 99.86% (see Table 1).

2. NPS Cloudlab

The computing services provided via Cloudlab are desktop computing services that provide an operating system that runs Microsoft Office and dozens of software programs that support resident and DL curricula. Faculty offer Cloudlab VMs as an option to resident students, as special software may be installed on the VMs and not in the LRCs. The VDI infrastructure allows students access to VMs 24 hours per day seven days a week from any location.

a. Hardware

Students use their own computing devices to access Cloudlab. The servers and storage used to provide the Cloudlab service costs approximately \$100,000 (Administrator C, personal communication, August 1, 2016). Per the interview with Administrator C, he stated that the hardware is replaced every three years (personal communication, August 1, 2016). The cost per year equated to \$33,000 when divided by 400 seats per year, and the cost equates to \$83.33 per seat per year (see Table 1).

b. Software

The software cost is determined by adding the various aspects of software used for the VDI solution. Cloudlab has a 400-user VDI license that cost approximately \$170,000 (Administrator C, personal communication, August 1, 2016). The \$170,000 cost was determined based on special pricing for government/educational pricing. Administrator C is involved with pricing for the NPS VDI licenses and was able to provide an estimate on August 1, 2016. The licenses are refreshed every three years so \$170,000 is divided by three to equal \$56,667 (Administrator C, personal communication, August 1, 2016). The cost per year is then divided by 400 to get \$141.67 per seat per year. The desktop operating system (including Microsoft Office) is estimated to cost \$26.39 per seat based on pricing obtained from NPS (Administrator J, personal communication, September 16, 2016). The cost of \$1.00 per seat for the anti-virus software is based on the cost that NPS pays per seat (Administrator J, personal communication, September 16, 2016). The total cost for software is obtained by adding $\$141.67 + \$26.39 + \$1.00$, which totals \$169.06 per seat per user (see Table 1).

c. Manpower

The manpower involved with the front end and back end support is based on the 12 hours per month obtained during the interview (Administrator C, personal communication, August 1, 2016). That value is multiplied by 12 months providing a value of 144 hours per year for backend and frontend support. In addition, four hours per quarter are spent on desktop maintenance, which equates to 16 hours per year (Administrator C, personal communication, August 1, 2016). The labor hours spent working on the Windows image is 240 hours per year (see Table 1). The IT specialists are on GS-09 pay scale, which has a salary range of approximately \$50,000–\$70,000 (per www.opm.gov/policy-data-oversight/pay-leave/salaries-wages/salary-tables/16Tables/html/SF.aspx). A labor rate of \$30.00 an hour value is used in the calculations, as the rate falls within the range. The labor hours when totaled equal 400 hours. The total labor (400 hours) is multiplied by \$30 per hour, which equals \$12,000. This value is divided by the 400-seat license, which yields \$30.00 per seat per year (see Table 1).

d. Availability

The system is unavailable 24 hours per month for maintenance (Administrator C, personal communication, August 1, 2016). The pools are recomposed every six months, which makes the system unavailable for eight hours per year (Administrator C, personal communication, August 1, 2016). A power outage occurs once a year, which causes the system to be unavailable 24 hours per year (Administrator C, personal communication, August 1, 2016). The total time the system is unavailable is 56 hour per year. The 56 hour value is subtracted from the total number of hours per year is 8,760, which equals to 8,704. The value is divided by 8,760 and multiplied by 100 to show an availability of 99.3% per year (see Table 1).

3. NWC

The NWC provides desktop computing services that run Windows VMs, which offer Microsoft Office for word processing and other such programs. They also offer other Navy applications, but by far, the biggest suite of applications is Office on the desktop platform (Administrator E, personal communication, July 25, 2016).

a. Hardware

The NWC uses zero clients to connect to the VDI infrastructure, which cost \$210 per seat (Administrator E, personal communication, July 25, 2016). The backend hardware (servers, storage, and network equipment) costs approximately \$300,000 (Administrator E, personal communication, July 25, 2016). The equipment is replaced every three years, so \$300,000 is divided by to get a cost of \$100,000 per year. The cost per year is divided by 900 seats, which yields \$111.00 per seat. That number is added to the cost per zero clients, which comes out to \$321.00 cost per seat per year for hardware (see Table 1).

b. Software

The software cost is determined by adding the various aspects of software, which is used for the VDI solution. The NWC has a 900-user VDI license, which costs approximately \$500,000 (Administrator C, personal communication, August 1, 2016).

The \$500,000 cost was determined based on pricing for the government/educational sector. Administrator C is involved with pricing for the NPS VDI licenses and was able to provide an estimate. The licenses are refreshed every three years, so \$500,000 is divided by three to equal \$166,667. The yearly maintenance fee of \$5,000 (Administrator C, personal communication, August 1, 2016) is added to the cost \$166,667 to yield \$171,667, which is then divided by 900 to get \$190.74 per seat per year (see Table 1). The desktop operating system (including Microsoft Office) is estimated to cost \$26.39 per seat based on pricing obtained from NPS (Administrator J, personal communication, September 16, 2016). The cost of \$1.00 per seat for the anti-virus software is based on the cost, which NPS pays per seat (Administrator J, personal communication, September 16, 2016). The total cost for software is obtained by adding $\$190.74 + \$26.39 + \$1.00$, which totals \$217.39 (see Table 1).

c. Manpower

The manpower involved is based on the 32 hours per month obtained during the interview (Administrator E, personal communication, July 25, 2016). That value is multiplied by 12 months providing a value of 384 hours per year. The IT specialists are in the GS-09 pay scale, which has a salary range of approximately \$50,000–\$70,000 (per www.opm.gov/policy-data-oversight/pay-leave/salaries-wages/salary-tables/pdf/2016/BOS.pdf). A labor rate of \$30.00 an hour is used in the calculations, as the rate falls within the range. The labor hours of 384 (see Table 1) multiplied by \$30 per hour equals \$11,520. This value is divided by 900, which yields \$12.80 per seat per year (see Table 1). Labor for the creation of lab images is 640 labor hours (see Table 1) per year while the IT Specialists are GS-09 pay scale, so \$30.00 per hour is used (per www.opm.gov/policy-data-oversight/pay-leave/salaries-wages/salary-tables/16Tables/html/SF.aspx). The total is \$19,200 divided by 900 users, which equals \$21.33. The 10 Helpdesk technicians spend approximately four hours per day (Administrator E, personal communication, July 25, 2016) on support of the thin client solution, which when calculated equals 800 hours per month. That value multiplied by 12 months is 9,600 hours per week while the average rate of a GS-09 technician is assumed to be \$30 per hour (per www.opm.gov/policy-data-oversight/pay-leave/salaries-wages/salary-tables/16

Tables/html/BOS.aspx) labor cost for the Helpdesk is \$288,000 divided by 900 users equals \$320 per seat per year. The total cost of the VDI solution at the NWC is \$321.00 + \$217.39 + \$354.13, which equals \$892.52 (see Table 1).

d. Availability

Administrator E stated that the VDI system has high availability and high redundancy. No outages have been reported. The VDI system availability is greater than 99.9% (Administrator E, personal communication, July 25, 2016). The system is architected so that it may continue to provide services if hardware, such as a hard drive, power supply, or server goes down.

4. CSUMB

At CSUMB, thick client computers are used to provide students with software services and computing power. The thick clients are installed in labs throughout the campus. The computers have applications installed so that students can complete their coursework, as well as access the Internet and CSUMB network (Administrator F, personal communication, July 26, 2016).

a. Hardware

The interview with Administrator F revealed that the thick client computers used at CSUMB range in price from approximately \$2,000 to \$3,000 each (Administrator F, personal communication, July 26, 2016). The average value for the cost per computer will be assumed to be \$2,500. The computers are replaced every three years, so the cost per year for hardware is \$833.00.

b. Software

Since CSUMB is a state university and received educational discounts for hardware and software, it is assumed that the cost for the Desktop (including Office products) operating system, Landesk, and anti-virus software is similarly priced to that of NPS. It is assumed that the cost for desktop operating system is \$26.39 per seat. The

Landesk price is \$25.67 per seat and anti-virus software costs \$1.00 per seat. The total for software cost is \$53.06 per seat per year.

c. Manpower

The values for manpower involved with creating and deploying images was obtained during the interview with Administrator F. The base image requires 16 hours to complete (Administrator F, personal communication, July 26, 2016). All technicians utilize the same base image so efforts do not need to be duplicated. The time required to build lab images 320 labor hours per technician (see Table 1). The group has four technicians working on this task, which totals 1,280 labor hours (Administrator F, personal communication, July 26, 2016). That 1,280 value is added to 16 hours, which it takes to create the base image, and totals 1,296. An hourly rate of \$25 per hour is used in the calculations (Administrator F, personal communication, October 11, 2016). That value is then multiplied by \$25 per hour for labor and then divided by 1,000 computers. The cost per seat per year is \$32.04 (see Table 1). The deployment of the images takes 80 hours x four technicians, which totals 320 labor hours (Administrator F, personal communication, July 26, 2016). The time it takes is then multiplied by \$25 per hour, which equals \$8,000 in labor costs for deployment of the images. That value is divided by 1,000, which yields \$8.00 per seat per year. Administrator F stated that approximately 20 hours per week are spent on maintenance (Administrator F, personal communication, July 26, 2016). Maintenance equates to approximately 960 hours per year, which is then multiplied by the \$25 rate and then divided by 1,000, which equals \$24.00. The total cost for manpower is obtained by adding \$32.04, \$8.00, and \$24.00, which equals \$64.04 (see Table 1).

d. Availability

The values for availability of the lab were derived from the fact that the lab computers were available (up-time) 24 hours per day x 365 per year. They were not available during imaging, which was four hours per year. Administrator F's comments about the hours of operation, labs are available from 8:00 a.m.–10:00 p.m. which equates to 14 hours per day, were not included in the calculations, as the computers are available

24 hours per day even though they are not accessible due to restrictive lab hours. However, the library has 60 computers, which are available 18 hours per day. The Café has 20 machines, which are ready for use 24/7. The calculations are based upon the primary labs, which are up 24 hours per day x 365, which equates to 8,760 hours per year. The lab computers are unavailable four hours per year due to imaging. The lab computers are available a total of 8,756 hours per year (Administrator F, personal communication, July 26, 2016). The image deployment takes place at times when the labs are closed. However, the time it takes to deploy images does count against the availability, as the lab computers are unavailable for four hours per year (Administrator F, personal communication, October 11, 2016). When asked if any outages occur, and if yes, how frequently do they occur and for how long, Administrator F stated he could not recall an instance when the campus was open and a computer lab was not. He pointed out that scenario is rare. The infrastructure is extremely solid as Administrator F stated with no issues with the hardware. The labs solution at CSUMB is available 99.95% of the time (see Table 1). Due to restrictive lab hours, the lab computers are accessible 14 hours per day.

5. Stanford

At Stanford University, a private cloud offering provides software applications and computing/processing power (Administrator G, personal communication, August 4, 2016). Students are able to connect to the private cloud, known as Farmshare, to run applications so that they may complete their coursework. The Farmshare resources are also made available to students, faculty, and staff with Stanford IT accounts to facilitate research at the university (Administrator G, personal communication, August 4, 2016).

a. Hardware

Students use their own computing devices to access Farmshare (Administrator G, personal communication, August 4, 2016). The servers and storage used to provide Farmshare cost approximately \$300,000 (Administrator G, personal communication, August 4, 2016). Per the interview with Administrator G, he stated that the hardware is replaced roughly every three years. The cost per year equated to \$100,000 and when

divided by 5,500 users per year, the cost equals \$18.18 per seat per year (Administrator G, personal communication, August 4, 2016) (see Table 1).

b. Software

No additional cost is incurred for accessing Farmshare, as students use BYOD software or the software is downloaded from the software portal (Administrator G, personal communication, August 4, 2016). BYOD devices come with operating systems installed. The operating system in some cases comes with software, which can be used to access Farmshare. If the OS does not include access software, Stanford has a portal that provides the software free of charge to students (Administrator G, personal communication, August 4, 2016). Other software programs can also be downloaded from the Internet free of charge.

c. Manpower

The manpower involved with supporting Farmshare is divided into different fields (see Table 1). Time spent providing technical support is 20% of 160 hours per month, which equates to 32 hours per month (Administrator G, personal communication, August 4, 2016). The time spent providing maintenance is 10% of 160 hours, which totals 16 hours per month (Administrator G, personal communication, August 4, 2016). The task of software installation for Farmshare requires 10% of 160 work hours in a month, which totals 16 hours per month (Administrator G, personal communication, August 4, 2016). When these tasks are added together, 768 labor hours are spent on supporting Farmshare. The labor hours are multiplied by \$60 per hour, which equals \$46,080 per year (see Table 1). That cost is divided by 5,500 users, which nets a cost of \$8.38 per seat per year (see Table 1).

d. Availability

Concerning availability, Administrator G stated two outages had occurred over the last two years. He informed the researcher that one outage occurred, which may have impacted users and lasted a few hours, while the other outage lasted less than an hour and had no impact to users. The longer outage was due to a server going down because of a

hardware issue. The server in question had to be replaced and rebuilt. At that time, some services were interrupted, but if students had the client installed, they were able to use that application to access the system. In theory, they should have been able to access everything even though a server was unavailable explained Administrator G (personal communication, August 4, 2016). With that said, the availability of Farmshare is shown to be greater than 99.9%.

Chapter IV covered four use cases pertaining to VDI technologies used at NPS. Data obtained during the interview of the four NPS VDI administrators was gathered by asking 15 questions about the VDI infrastructure, which they administer. The answers to the questions showed the services provided by the technology, as well as a cost estimate for hardware, software, and manpower. The availability associated with each VDI system was shown as well.

The chapter also discussed various IT computing technologies used at the Naval College (Newport, Rhode Island), CSUMB, and Stanford University to provide computing capabilities and software services to students. The IT professionals were asked 19 questions and their answers were presented throughout the chapter. It was shown how their labs or equivalent services are configured/used to provide computing resources to students. The estimated hardware and software cost of the technology followed by manpower involved in maintaining the solution was shown. The availability of the solutions was also provided.

The data was presented throughout the chapter and was then followed by a table that calculated the cost per seat per year for the hardware and software cost and manpower associated with each IT solution used at each university. The availability for each solution was shown as well. The NPS case included values for LRCs, as well as values for the enterprise private cloud, Cloudlab. These resources may prove to be invaluable in providing options for NPS to deliver software to NPS customers with new technologies that offer a lower cost and higher efficiency than the current solution. The analysis for the hardware costs, software costs, manpower costs, and availability are discussed in the following chapter.

V. CONCLUSIONS AND RECOMMENDATIONS

The values for hardware and software cost, manpower, and availability were calculated for the NPS LRCs, NPS Cloudlab, NWC, CSUMB, and Stanford University. The goal was to show the cost per seat per year for each category at each university and the availability of the solution was shown as percent available. The NPS use case showed values for both the NPS LRCs and NPS Cloudlab (which is considered to be an extended LRC for the enterprise). The process of how each value was determined was explained in the previous chapter and a table was provided for simplicity. The following sections analyze which university has the highest cost per seat in each of the four categories: hardware cost, software cost, manpower, and availability followed by the next highest cost, etc. The most cost effective solutions are determined followed by the conclusion of the research. The final section includes recommendations made by the researcher that may provide for an efficient, cost effective, labor effective solution that can provide availability equivalent to or better than the availability provided by the current LRC technology.

A. ANALYSIS

The hardware cost of an IT solution is an important aspect to consider when making decisions on which solution should be implemented at a university. The hardware cost was determined on a cost per seat per year basis for NPS, NWC, CSUMB, Stanford, and the NPS enterprise cloud known as Cloudlab.

Based on the calculations, and by ranking them from highest to lowest, the highest cost solution was that of NPS and CSUMB. The average cost per computer at both universities was \$833.33. The next highest cost was that of the NWC. The cost for the NWC solution was \$321.00. Once again, the zero client computers cost \$210 each and have increased the overall cost of the solution as the computers are used to access the VDI system. Next on the list was the NPS Cloudlab solution. The hardware cost for the Cloudlab solution was \$83.33 per seat per year. Lastly, the cost of the Farmshare solution was least expensive as the hardware cost \$18.18 per user per year when taking into

account the 5,500 unique user accounts that accessed the system in a year. Note that the Stanford Farmshare private cloud is accessed by students using BYOD systems.

Another important factor to consider is that of software costs. Note that the software applications used by students at the universities were not included in the calculations for software cost. Where applicable (depending on the IT solution), the operating system and core supporting software (such as Microsoft Office, Landesk, and anti-virus software) are taken into account. In some cases, the software used to access the system was considered as part of the cost, such as software license costs for the VDI solutions. The software cost was determined on a cost per seat per year basis for NPS, NWC, CSUMB, Stanford, and NPS Cloudlab. Based on the calculations, and by ranking them from highest to lowest, the highest cost software solution was that of the NWC. The software cost was \$217.39 per seat per user. The NWC has a 900-seat user license. The second highest cost solution was that of NPS Cloudlab. The software cost for the 400-user VDI solution was shown to be \$169.06. The third most costly solution was shared among NPS and CSUMB. The software used to support the IT lab solutions implemented at the two universities was \$53.06, as both shared a common pricing model for universities and also utilized similar software to create and deploy images. The most cost effective solution was that of Stanford University and its Farmshare private cloud. The software used to access Farmshare is provided by Stanford University through its portal. In some cases, software used to access Farmshare is included with some operating systems used on the BYOD systems, which the students provide, or is available for free on the Internet. The cost to access Farmshare is included with the BYOD device or available for download via the portal.

The next factor considered when evaluating IT systems for use in production environments is manpower. The manpower associated with the maintenance of the equipment was taken into account as labor hours. The labor hours were then multiplied by the hourly rate of the IT specialist based on labor charts obtained from the Office of Personnel Management (OPM) or based on information provided by the interviewee. The manpower cost was determined on a cost per seat per year basis for NPS, NWC, CSUMB, Stanford, and NPS Cloudlab. Based on the calculations, and by ranking them

from highest to lowest, the highest cost manpower solution was that of the NWC. The solution used at NWC is labor intensive. As seen in Table 1, the manpower was determined by three factors: backend/frontend maintenance, image creation and helpdesk support. The cost for helpdesk support was the most expensive aspect of the manpower cost. Since the solution provided to customers at NWC is the VDI solution, the helpdesk spends approximately half their time working on VDI related issues (Administrator E, personal communication, July 25, 2016). The second most costly solution was the thick client model at NPS. The task of creating one base image and five lab images in addition to deploying to five LRCs is performed twice a year. The NPS thick client solution costs \$133.34 per seat per year. The third most costly in terms of labor was CSUMB as it cost \$64.04 per seat per year. The task takes of building, configuring, and testing the base image and 10 lab images is performed once per year. The fourth most costly solution based on manpower was that of NPS Cloudlab. The cost for labor hours is \$30.00 per seat per year and the VDI cloud solution supports 400 concurrent users. The least costly solution in terms of manpower is the Stanford Farmshare cloud. The total cost for manpower is \$8.38 per seat per year. The manpower considered is for technical support, maintenance, and software installation and when 5,500 user are considered, the costs are considerably lower (see Table 1). Note that Farmshare does not create a desktop for the client side, and therefore, does not use manpower on creating lab images whereas the other universities do.

The final criteria to be considered when discussing lab computing environments is their availability. The availability in the NPS and CSUMB cases evaluated had to do with the amount of time the computers were physically accessible to students. The lab hours were a factor in the amount of time the computers were available. Another factor to be considered was the amount of downtime (whether power outages or maintenance performed), which caused the IT solution to be unavailable. These factors played a greater role in the NWC, Stanford, and NPS Cloudlab solutions. The availability was determined as a percentage of time that the IT solution was available at NPS, NWC, CSUMB, Stanford, and for the NPS Cloudlab. Based on the calculations, and by ranking them from highest to lowest, the greatest system availability was at the NWC, CSUMB

and Stanford University. All three solutions were available 99.9% of the time. The NWC uses a VDI private cloud while the Stanford solution is a private cloud offering. CSUMB uses the thick client technology throughout their labs. Both private cloud solutions have high redundancy built into the systems while the thick client solution does not have the option for redundancy other than the fact that a lab has multiple clients. If one thick client is not available, others are that can be utilized. However, if the lab is not available, then a secondary lab is not available as a backup. A close second was the thick client solution used at NPS. The computers are available 99.86% of the time and the labs are open 24/7/365 with a minimal amount of downtime. The third most available system is NPS Cloudlab, which is available 99.3% of the time. The technology is a private cloud technology that is highly available other than during a planned power outage or during maintenance. All IT solutions used at the four universities have an availability greater than 99%.

The total cost per seat per year was calculated by adding the hardware costs plus the software costs plus the manpower costs for each university. Based on the calculations, and by ranking them from highest to lowest, the highest cost per seat per year was at NPS followed by CSUMB, the NWC, NPS Cloudlab, and lastly, by Stanford University. The costs were calculated to be \$1019.73, \$950.43, \$892.52, \$282.39, and \$26.56, respectively. Based on the results, the Stanford University solution is the most cost effective followed by the Cloudlab solution.

B. CONCLUSIONS

NPS relies on thick client computers to deliver software to students so they can work on their homework, lab projects, thesis assignments, and dissertations. The lab support process used to deliver the latest hardware and software to students in the labs and classrooms has been in use for over two decades. Due to the end strength issue, budget reductions (sequestration), unsupported software, and the directives set forth by the DDCIO (N) the existing LRC solution at NPS has become increasing more difficult to maintain in an efficient manner. Several benefits could be gained from the results of this research. The NPS IT department spends a lot of effort, budget and manpower/labor

hours maintaining the labs and classrooms at NPS. This effort is done to provide students, faculty and staff with the latest thick client hardware and software to aid them with their curricula, research, and support of the NPS mission. The following research questions were asked to guide the researcher in hopes of meeting the aforementioned goals:

- Research question #1: What is the current technology used for providing computing resources in the labs and classrooms at NPS? Is this technology still a viable solution?

Based on the information provided through the research, the current technology used for delivering computing resources in the LRC labs and classrooms at NPS is the thick client technology. This technology is still a viable solution when considering the resident contingency, as thick clients have proven to be effective in providing students with the computing services necessary to perform their schoolwork while attaining a master's degree or a doctoral degree while at NPS. The research also showed that CSUMB has implemented a thick client computing solution to provide lab services to their students, which further strengthens the argument that the thick client solution is still viable. However, the thick client solution is not considered a viable solution when the DL component of the university is taken into consideration. DL students are not allowed to connect to LRC computers. Since some software licenses restrict use to government owned computers, DL students do not have access to this software and may be forced to purchase the software or rely on NPS to provide the software, which is being done using NPS Cloudlab. The other option is to do without the software or find alternative software that meets the needs of the curriculum.

- Research question #2: What viable options are available to provide the same level of computing capabilities in the LRCs at NPS?

To search for viable options with which to provide the same level of computing capabilities in the LRCs at NPS, the researcher interviewed IT professionals at the Naval War College, CSUMB, and Stanford University to gather data about each IT solution. The researcher also interviewed an IT specialist who supports the NPS enterprise private cloud offering known as NPS Cloudlab to assess if the technology can be extended into the LRC labs. Based on the information gathered, CSUMB uses the same thick client technology and deployment methods utilized at NPS. For that reason, the CSUMB lab

solution is not considered for further analysis. The availability of all IT solutions are greater than 99% and are no longer considered when making a determination. The VDI solution used at the NWC is of interest, as the total cost of the solution is less than that of the NPS solution. The manpower cost is greater than that of NPS. However, when reviewing Table 1, it is shown that the helpdesk incurs the greatest portion of the cost. If NPS were to adopt this technology, its helpdesk cost would rise, as more labor hours would be spent supporting the VDI solution. With that in mind, the VDI solution at NWC is evaluated further. The private cloud offering used at Stanford University has been shown to be a viable option due to the low hardware and software costs and low requirement for manpower spent maintaining the system. The private cloud solution is considered although the thick client aspects of their solution matches those used at NPS and are not considered.

- Research question #3: For the viable options, does a cost comparison show which solution would be preferable for the current labs and classrooms at NPS?

The NPS lab solution (including the enterprise Cloudlab solution) along with the NWC IT solution, CSUMB lab solution, and the Stanford private cloud solution were evaluated to gain a better understanding about which technologies were used at those universities. The hardware and software costs were assessed by gathering information from interviewees at NWC, CSUMB, Stanford, and NPS Cloudlab administrators. The manpower and availability were also based on the information provided in the interviews. A series of calculations were performed (see Table 1) to determine a cost per seat per year for each criteria so that the various IT solutions could be compared. Based on the information provided, a cost comparison shows which IT solution may be preferred for use in the LRCs at NPS.

In terms of hardware costs, Stanford University has the most cost effective solution costing \$18.18 per seat per year. This cost is possible because the cost of client hardware used to access Farmshare is passed down to the student as the school uses a BYOD model. The second most cost effective hardware solution is that of NPS Cloudlab. Once again, no costs are incurred regarding client hardware. Students are expected to access the VDI solution via BYOD or from their personal computers.

With regards to software costs, Stanford University provides the most cost effective method for using software applications through the Farmshare cloud. The software used to access Farmshare is provided to students through a web portal and is free of charge. Another option for students would be to download the software from the Internet free of charge if the BYOD device that they use does not include the software. The second most cost effective software solution is provided through the NPS and CSUMB thick client model. The software cost includes the operating system, Microsoft Office products, Landesk software used for deployment, and anti-virus software. The fact that universities obtain special pricing due to the educational services they provide, as well as purchasing more licenses helps to drive down the software costs.

When considering manpower, Stanford University provides the most cost effective solution with its Farmshare private cloud. The labor hours are divided between tech support, maintenance, and installing software to be used by students. The system is used by 5,500 users per year, which helps bring down the cost of the solution. The second most cost effective solution is that of NPS Cloudlab. The manpower spent on the NPS VDI cloud is divided between backend support, frontend support, and the building of the images. The system is used by 400 users per year, which lowers the cost.

The final point to consider is the overall cost of the solutions. Included in the overall costs are the hardware cost, software cost, and manpower. The most cost effective solution is the Stanford private cloud technology. The cost is \$26.56 per seat per year. The large number of users was a factor in the low cost of the solution, as was the fact that the students use BYOD devices for client computing services. No software costs were incurred either as the user provides software used to access Farmshare or Stanford does. The second most cost effective solution was the NPS VDI solution. The cost for Cloudlab was \$282.39 per seat per year. The BYOD model required to access Cloudlab along with the educational pricing for the 400-user license contributed to the low cost per user.

This research could benefit the university and the Navy, as this research looks to reduce costs in both hardware purchases and tools associated with deploying images to the lab computers. Another potential benefit would be to find effective methods of IT implementation to cut down on labor hours that the IT staff spends in sustaining the LRC

and classroom environment all while providing the highest percent availability. Students, faculty, and staff may potentially be more productive in their work, coursework and/or research, as some of the possible computing solutions will be available from any smart device, from any location, and at any time.

C. RECOMMENDATIONS

Recommendation made based on the research provided are as follows:

- NPS should implement a private cloud solution similar to the technology being used at Stanford University. The software applications used in the LRCs, those of which that can be provided via the cloud technology, should be made available to students with the use of the private cloud technology. The software used to access the private cloud should be made available to students through the NPS portal, which currently exists. The implementation of the private cloud technology will provide DL students an alternate access method to computing and software services, which was one factor that prompted the employment of the VDI private cloud solutions at NPS.
- A BYOD policy should be implemented at NPS in general so that the devices can be incorporated into the LRCs. Students will be able to access the private cloud with BYOD to run the applications, which have been made available with the use of the technology. The applications, which could not be made available with the use of the private cloud solution, will be accessed using the VDI technology. All students should be required to own or have access to a BYOD system. The ITACS department should provide minimum specifications for the BYOD systems to be used at NPS.
- The NPS Cloudlab hardware should be upgraded so that it is capable of supporting 1,000 users. The VDI license should be increased from a 400-user license to a 1,000-user license. This increase will replace the 225 thick client computers currently located in the LRCs, as well as the thick client computers located in the classrooms, yet support the 225 plus user capacity that the LRCs and classrooms offer. The 1,000-user license will also provide for an increase in resident and DL student use.
- Approximately half of the thick client computers used in the LRCs should be replaced with zero client computers once the computers are out of warranty. Until then, a software client can be installed that will redirect the thick client computer to the newly upgraded Cloudlab. This action will greatly reduce the hardware cost for LRCs, as it is currently the greatest expense. The other half of the thick client computers located in the LRCs should be removed and tables should be installed that will have power and Ethernet ports built into the tables. Students can then plug their BYOD

technology into those power and Ethernet ports to access the newly added private cloud technology or may access the Cloudfab virtual machines, as well as have access to Internet, storage, and the NPS network.

- Concerning the VDI technology, the use of a base lab image should be considered together with VMware's App Volume technology. Using App Volume technology, the application will be "virtualized" and thereby allow it to be assigned to students based on their username. Once assigned, the student logs into the virtual machine and the application appears as if it were installed on the virtual machine's operating system so that the student can use the application. Thus, manpower associated with creating and maintaining LRC and lab images is reduced, as it is the most costly at the NWC and NPS.
- A review of most commonly used software applications in LRCs at NPS should be performed. The costs per seat should be determined for each software application. The vendor for the software should be contacted to ascertain if the software is offered as a SaaS solution. If so, the cost per seat should be obtained from the vendor to determine if moving to the SaaS model would be beneficial. The availability of the SaaS software should be compared to the availability of the locally installed software in the LRCs.

The availability of the recommended systems will be greater than 99%, which is equivalent to the current LRC solution used at NPS.

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

- Amazon EC2. (n.d.). Retrieved from https://aws.amazon.com/ec2/?nc2=h_13_c
- Anderson, A. (1994). *Foundations of computer technology*. Boca Raton, FL: CRC Press.
- Beach, T. E. (2012, January 24). Computer concepts and terminology. Retrieved from <http://www.unm.edu/~tbeach/terms/types.html>
- Chellappa, R. K. (2010). Retrieved from <http://www.bus.emory.edu/ram/>
- Christensson, P. (2006). Thick client. Retrieved from <http://techterms.com/definition/thickclient>
- Cloud Computing Forum & Workshop. (2016). Retrieved from <https://www.nist.gov/news-events/events/2010/05/cloud-computing-forum-workshop>
- Cloud services. (2015). Retrieved from <http://aws.idglabs.net/?p=2442>
- Cluley, J. C. (1983). Basic microprocessor architecture. In *Interfacing to microprocessors* (pp. 1–12). London: Macmillan Education UK.
- Dean, T. (2012). *Network+ guide to networks* (6th ed.). Boston, MA: Course Technology.
- Deep, M. (2013). Getting your BYOD policy right. Retrieved from <http://niiconsulting.com/checkmate/2013/07/getting-your-byod-policy-right/>
- Dell Wyse D90D7 thin client. (n.d.). Retrieved from <http://www.optiodata.com/dell-wyse-d90d7-thin-client>
- Desktops-as-a-Service (DaaS) and hosted desktop solutions. (n.d.). Retrieved from <https://www.citrix.com/products/daas/overview.html>
- Dhir, A. (2004). *The digital consumer technology handbook: A comprehensive guide to devices, standards, future directions, and programmable logic solutions*. Burlington, MA: Elsevier Science.
- Elastic beanstalk. (n.d.). Retrieved from https://aws.amazon.com/elasticbeanstalk/?nc2=h_13_dm
- Elastic computing (EC). (n.d.). Retrieved from <https://www.techopedia.com/definition/26598/elastic-computing-ec>
- Glatzer, H. (1982, Feb–Mar). Operational choice. *PC*, 1(51), 50–53.

- Goldstine, H. H. (1972). *The computer: From Pascal to Von Neumann*. Princeton, NJ: Princeton University Press.
- Harkins, M. (2016). Mobile: Learn from Intel's CISO on securing employee-owned devices. Retrieved from <http://www.govinfosecurity.com/webinars/mobile-learn-from-intels-ciso-on-securing-employee-owned-devices-w-264>
- Intel 8080 microprocessor. (2016). Retrieved from <http://techxway.com/2016/02/11/intel-8080-microprocessor/>
- Isaacson, W. (2014). *The innovators: How a group of inventors, hackers, geniuses, and geeks created the digital revolution*. New York: Simon & Schuster.
- ITACS mission. (2011). Retrieved from <https://wiki.nps.edu/display/ITACS/ITACS+Mission>
- Jennings, Bartik, J., Rickman J., & Todd, K. (Eds.). (2013). *Pioneer programmer: Jean Jennings Bartik and the computer that changed the world*. Kirksville, MO: Truman State University Press.
- Kavanagh, P. (1995). *Downsizing for client/server applications*. Cambridge, MA: Academic Press Inc.
- Korzeniawski, P. (2013). Following both sides of the decentralized vs. centralized IT debate. Retrieved from <http://searchdatacenter.techtarget.com/opinion/Following-both-sides-of-the-decentralized-vs-centralized-IT-debate>
- Learning Resource Centers (LRC) policy. (2013). Retrieved from <https://wiki.nps.edu/display/IT/Learning+Resource+Centers+%28LRC%29+Policy>
- Lindsey, Mark R. (2007, December 10). On network designs. Retrieved from <https://200ok.info/2007/12/10/on-network-diagrams/>
- Mack, P. (2005). The microcomputer revolution. Retrieved from <http://pammack.sites.clemson.edu/lec122/micro.htm>
- Madden, B. (2014). The state of the VDI industry in 2014. Retrieved from <http://www.computerweekly.com/opinion/The-state-of-the-VDI-industry-in-2014>
- Mainframes vs. supercomputers. (2013). Retrieved from <http://aspg.com/mainframes-vs-supercomputers/#.VyU5vU3SlhA>
- McCafferty, D. (2014). Useful virtualization stats, trends and practices. Retrieved from <http://www.cioinsight.com/it-strategy/cloud-virtualization/slideshows/useful-virtualization-stats-trends-and-practices.html>

- Mell, P., & Grance, T. (2011). *The NIST definition of cloud computing, recommendations of the national institute of standards and technology*. (Special Publication No. SP-800-45). Gaithersburg, MD: Computer Security Division, Information Technology Laboratory, National Institute of Standards and Technology. doi: <http://dx.doi.org/10.6028/NIST.SP.800-145>
- Microsoft azure. (n.d.). Retrieved from <https://azure.microsoft.com/en-us/overview/what-is-azure/>
- NPS academic catalog. (2016). Retrieved from <http://my.nps.edu/web/registrar/academic-catalog>
- Oglesby, R. (2006, July 20). Virtual desktop infrastructures (VDI): What's real today, what's not, and what's needed. Retrieved from <http://www.brianmadden.com/blogs/ronoglesby/archive/2006/07/20/virtual-desktop-infrastructures-vdi-what-s-real-today-what-s-not-and-what-s-needed.aspx>
- Personal computer. (n.d.). *PC magazine*. Retrieved from <http://www.pcmag.com/encyclopedia/term/49133/personal-computer>
- Portnoy, M. (2012). *Virtualization essentials*. Indianapolis, IN: John Wiley & Sons, Inc.
- Raths, D. (2013). Crossing the device divide. *The Journal*, 40(5), 9–13.
- Revolution. (n.d.). Retrieved from <http://www.computerhistory.org/revolution>
- Rivero, Gonzalo. (2015). The OSI model. Retrieved from <https://commons.wikimedia.org/wiki/File:Osi-model-jb.svg>
- Rouse, M., & Barrett, A. (2016). Virtual desktop infrastructure (VDI). Retrieved from <http://searchservvirtualization.techtarget.com/definition/virtual-desktop-infrastructure-VDI>
- Rouse, M., & Botelho, B. (2015). Desktop as a service (DaaS). Retrieved from www.searchvirtualdesktop.techtarget.com/definition/desktop-as-a-service-DaaS
- Simplify the delivery of windows desktops and applications with VMware horizon air. (n.d.). Retrieved from <http://www.vmwhorizonair.com/overview/>
- Softlayer services. (n.d.). Retrieved from https://www.synaptic.att.com/clouduser/html/productdetail/Softlayer_services.htm
- Thick client. (2016). Retrieved from <http://www.computerhope.com/jargon/t/thickcli.htm>
- Thin client. (n.d.). Retrieved from <http://www.pcmag.com/encyclopedia/term/52832/thin-client>

- Tyson, B. (2014). Top 10 SaaS providers. Retrieved from <http://www.brighthub.com/environment/green-computing/articles/105724.aspx>
- United States Army. (n.d.). Historical monograph: Electronic computers within the Ordnance Corps. <http://ftp.arl.mil/ftp/historic-computers/>
- van den Blink Clare. (2009). Uses of labs and learning spaces. Retrieved from <http://er.educause.edu/articles/2009/3/uses-of-labs-and-learning-spaces>
- Velte, A. T., Velte, T. J., Elsenpeter, R. C. (2010). *Cloud computing: a practical approach*. New York: McGraw-Hill.
- Virtual desktop infrastructure (VDI) market in U.S. to grow at 29.70% CAGR by 2019. (2015). Retrieved from <http://libproxy.nps.edu/login?url=http://search.proquest.com/docview/1652052179?accountid=12702>
- VMware. (2014). Virtualization essentials. Retrieved from <http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/ebook/gated-vmw-ebook-virtualization-essentials.pdf>
- Wallace, J., & Erickson, J. (1993). *Hard drive: Bill Gates and the making of the Microsoft empire* (Reprint ed.). New York City: HarperCollins.
- What is docker? (n.d.). Retrieved from <https://www.docker.com/what-docker>
- Yost, J. R. (2005). *The computer industry*. Westport, CT: Greenwood Press.

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California