2012-03

How Virtual Technology Can Impact Total Ownership Costs on a USN Vessel

Boothe, BJ S.

Monterey, California. Naval Postgraduate School

http://hdl.handle.net/10945/6771

NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

HOW VIRTUAL TECHNOLOGY CAN IMPACT TOTAL OWNERSHIP COSTS ON A USN VESSEL

by

BJ S. Boothe

March 2012

Thesis Advisor:	Glenn R. Cook
Second Reader:	Alberto Barreto III

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE

Form Approved OMB No. 0704–0188

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE March 2012

3. REPORT TYPE AND DATES COVERED Master's Thesis

4. TITLE AND SUBTITLE How Virtual Technology Can Impact Total Ownership Costs on a USN Vessel

5. FUNDING NUMBERS

6. AUTHOR(S) BJ S. Boothe

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 ______N/A______.

12a. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT (maximum 200 words)

This thesis investigates the development of virtual technology and how the Consolidated Afloat Network Enterprise Services (CANES) program can reduce Total Cost of Ownership when a ship's local area networks are upgraded. With the recent development of cloud computing technologies, everyone from Fortune 500 companies to defense organizations believes that there are tangible benefits to moving operations to the cloud. This idea is particularly resonant with Naval Surface operations, consistent with the downsizing of personnel on surface ships, and with the Chief of Naval Operation's vision for information systems to be agile, relevant, and cost effective. By building a scalable private cloud model that utilizes a centralized server for computer processing, thin client workstations were compared to current thick client architectures onboard surface vessels. With multicore server processors developed to handle several tasks simultaneously, the ability to consolidate and virtualize multiple servers and workstations aboard naval vessels is now possible from a blade server chassis. By consolidating the computer processing into a central location, total ship energy consumption could be reduced by31 kilowatts during peak usage. The reduced shipboard energy consumption cut shore power costs by $3.75 per hour and reduced fuel consumption by 2,400 gallons each operating quarter for a ship using Ships Service Gas Turbine Generators (SSGTG). Even with increased research and developments costs associated with the virtualization software, a ship's network becomes agile and elastic while reducing overall energy consumption.

14. SUBJECT TERMS Virtual Technology, Cloud Computing, Thin Client, Consolidated Afloat Network Enterprise Services, Total Cost of Ownership, Energy Consumption, Ships Service Diesel Generator, Ships Service Gas Turbine, Generator

15. NUMBER OF PAGES 85

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	**18. SECURITY CLASSIFICATION OF THIS PAGE**	**19. SECURITY CLASSIFICATION OF ABSTRACT**	**20. LIMITATION OF ABSTRACT**
Unclassified	Unclassified	Unclassified	UU

NSN 7540–01–280–5500

Standard Form 298 (Rev. 2–89)
Prescribed by ANSI Std. 239–18

HOW VIRTUAL TECHNOLOGY CAN IMPACT TOTAL OWNERSHIP COSTS ON A USN VESSEL

BJ S. Boothe
Lieutenant, United States Navy
B.S., University of Idaho, 2006

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN INFORMATION SYSTEMS AND OPERATIONS

from the

NAVAL POSTGRADUATE SCHOOL
March 2012

Author: BJ S. Boothe

Approved by: Glenn R. Cook
Thesis Advisor

Alberto Barreto III
Second Reader

Dr. Dan Boger
Chair, Department of Information Sciences

ABSTRACT

This thesis investigates the development of virtual technology and how the Consolidated Afloat Network Enterprise Services (CANES) program can reduce Total Cost of Ownership when a ship's local area networks are upgraded. With the recent development of cloud computing technologies, everyone from Fortune 500 companies to defense organizations believes that there are tangible benefits to moving operations to the cloud. This idea is particularly resonant with Naval Surface operations, consistent with the downsizing of personnel on surface ships, and with the Chief of Naval Operation's vision for information systems to be agile, relevant, and cost effective. By building a scalable private cloud model that utilizes a centralized server for computer processing, thin client workstations were compared to current thick client architectures onboard surface vessels. With multicore server processors developed to handle several tasks simultaneously, the ability to consolidate and virtualize multiple servers and workstations aboard naval vessels is now possible from a blade server chassis. By consolidating the computer processing into a central location, total ship energy consumption could be reduced by 31 kilowatts during peak usage. The reduced shipboard energy consumption cut shore power costs by $3.75 per hour and reduced fuel consumption by 2,400 gallons each operating quarter for a ship using Ships Service Gas Turbine Generators (SSGTG). Even with increased research and developments costs associated with the virtualization software, a ship's network becomes agile and elastic while reducing overall energy consumption.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

ARPA	Advanced Research Project Agency
ARPANET	Advanced Research Project Agency Network
BMC	Baseboard Management Controller
CENTRIX-M	Combined Enterprise Regional Information Exchange Maritime
CMC	Chassis Management Controller
DDCIO (N)	Department of the Navy Deputy Chief Information Officer
DoD	Department of Defense
DON	Department of the Navy
EDVAC	Electronic Discrete Variable Automatic Computer
ENIAC	Electronic Numerical Integrator and Computer
ET	Electronics Technician
FTP	File Transfer Protocol
GUI	Graphic User Interface
IBM	International Business Machines
IaaS	Infrastructure as a Service
iDRAC	Integrated Dell Remote Access Controller
iENCON	Incentivized Energy Conservation
IP	Internet Protocol
ISNS	Integrated Shipboard Network System
IT	Information Systems Technician
LAN	Local Area Network
MAC	Media Access Control Address
MIT	Massachusetts Institute of Technology
NASA	National Aeronautics and Space Administration
NCTSS	Naval Tactical Command Support System
NEC	Navy Enlisted Classification
NEURS	Navy Energy Usage Reporting System
NIAPS	Navy Information Application Product Suite
NIC	Network Interface Controller
NPV	Net Present Value

OMMS-NG	Organizational Maintenance Management System-Next Generation
OPTAR	Operating Target
PaaS	Platform as a Service
R-ADM	Relational Administrative Data Management
RAM	Random Access Memory
ROI	Return on Investment
RSupply	Relational Supply
SaaS	Software as a Service
SCI-LAN	Sensitive Compartmented Information Local Area Network
SSDG	Ships Service Diesel Generator
SSGTG	Ships Service Gas Turbine Generator
SUBLAN	Submarine Local Area Network
TCO	Total Cost of Ownership
UNIVAC	Universal Automatic Computer
WAN	Wide Area Network

ACKNOWLEDGMENTS

The author would like to extend his utmost gratitude to Mr. Glenn Cook and Mr. Albert “Buddy” Barreto for their mentorship, continued guidance, and infinite patience over the past year. Also, to a host of other faculty members such as Mr. Steve Iatrou and LCDR Dave Roberts, whose contributions to this writing have added immeasurable value. Lastly, the author would like to thank family, friends, and former shipmates who assisted with continuous support.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION AND BACKGROUND

This thesis investigates the development of virtual technology and how the Consolidated Afloat Network Enterprise Services (CANES) program can reduce Total Cost of Ownership when a ship’s local area networks are upgraded. With the development of cloud computing technologies, everyone from Fortune 500 companies to Defense organizations believe that there are tangible benefits to moving operations to the cloud. This idea is particularly resonant with Naval Surface operations, consistent with the downsizing of personnel on surface ships, and with the Chief of Naval Operation’s vision for information systems to be agile, relevant, and cost effective. A scalable private cloud model that utilizes a centralized server for computer processing and thin client workstations was compared to current thick client architectures onboard surface vessels. With multicore server processors developed to handle several tasks simultaneously, the ability to consolidate and virtualize multiple servers and workstations aboard naval vessels is now possible from a blade server chassis. By consolidating the computer processing into a central location, total ship energy consumption could be reduced in port and underway. Even with increased research and developments costs associated with the virtualization software, a ship’s network becomes agile and elastic while reducing overall energy consumption.

A. PURPOSE

The Navy currently has a vision of a 300-ship fleet. As each ship is commissioned, overhead costs may rise just to sustain this size of fleet at sea and in port due to increasing energy costs. One way to possibly reduce these costs is by following the cloud computing paradigm.

Cloud computing is based on virtual technology that allows resource sharing of a central server. Therefore, multiple users can run different operating systems simultaneously on a common computer. With the expanding capability of virtual technology, the Navy can upgrade legacy systems, shrink the footprint of local area networks, and have systems more interoperable as part of the CANES program.

By upgrading current shipboard computer networks with virtual technology, upgrades may become easier by updating either firmware or software in the fleet. With these potential savings in the fleet, this could make the initial investment of virtual technology beneficial by reducing manning and operating costs.

B. RESEARCH QUESTIONS

This research will focus primary on virtual technology for medium-sized surface naval vessels that run either Ships Service Gas Turbine Generators (SSGTG) or Ships Service Diesel Generators (SSDG) for producing electricity. To determine how efficient cloud computing is, the operating parameters for the Intel Xeon processor since its introduction 17 years ago will be researched. By comparing computing power measured in Floating Operating Points per Second (FLOPS) against processor speed, number of cores, and heat dissipation. This will determine how processors became efficient and powerful in today’s computing environment for virtual technology.

With future networks shifting toward a thin client computing environment, current shipboard systems energy consumption in kilowatts (kW) will be compared to a newer computing system with similar computing capacity. The difference in energy consumption will be calculated for a surface ship on shore power and underway utilizing either a Ships Service Gas Turbine Generator or Ships Service Diesel Generator.

The Total Cost of Ownership will be calculated for a five-year period based on the results of the private cloud model. By comparing operating cost savings with initial hardware investments, this will determine if virtual technology has a return on investment.

C. BENEFITS OF THE STUDY

The Secretary of the Navy is constantly seeking ways to improve energy conservation for commands ashore and afloat. With the U.S. Navy seeking cost-saving measures while upgrading current platforms to meet technology demands, virtual technology can be the answer. With the ability to consolidate legacy systems into a common computing environment such as CANES, legacy systems can have extended cycle lives and operate on new hardware vice on older systems that are no longer supported. This research will help identify how efficient virtual technology is and how it can meet the Secretary of the Navy's goals of reducing energy consumption afloat while having an agile network capable of meeting future requirements.

D. METHODOLOGY

The following methodology will be used to conduct research and analysis. A literature review will be conducted to understand the background of computer and network fundamentals. Next, the research will study the advantages and disadvantages of thin clients and cloud computing coupled with features of the CANES program will be applied to a surface vessel. The infrastructure and process of building a private cloud at Naval Postgraduate School will be researched and analyzed to discuss energy efficient features of both the hardware and software that could be utilized on a surface vessel. For the analysis, a series of measurements will be taken on the private cloud model to compare existing and possible differences in energy consumption of a shipboard computer network. Lastly, the results will be applied to a surface ship, and the Total Cost of Ownership (TCO) and Net Present Value (NPV) will be calculated for a five-year period if virtual technology was implemented for surface ships based on the current and potential prices of energy.

E. ORGANIZATION OF THE THESIS

This thesis will be organized according to the following chapters:

1. Literature Review

Chapter II will inform the reader with the history of computers and client server computing fundamentals. An introduction to cloud computing fundamentals and an overview of the CANES program will be analyzed.

2. Research Method-Shipboard and Systems Applications

Chapter III will go into detail about the Secretary of Navy's vision for energy conservation and the Navy's Incentivized Energy Conservation Program and how they apply toward surface vessels. The basics of SSGTGs and SSDGs will be reviewed and how a change in electrical load affects the hourly fuel burn rate. Next, the review of the hardware and software of a private cloud built at Naval Postgraduate School will be reviewed to give the reader and in-depth understanding how virtual software turns an ordinary server into an efficient computing environment.

3. Analysis

Chapter IV details the comparison of operating parameters of the Intel Xeon processor family. Additionally, the performance of the Dell M1000e chassis is measured while running 50 virtual desktops to simulate a potential shipboard private cloud. The results of the private cloud model are then analyzed for an Arleigh Burke Destroyer and Amphibious Landing Dock Ship while on shore power and underway. The TCO and NPV are calculated for a five-year period to determine potential operation savings.

4. Conclusion and Recommendations

Chapter V presents a summary of previous chapters and analysis. The findings of Chapter IV are summarized for the reader discussing how cloud computing can reduce Total Cost of Ownership on a naval vessel in five years. Finally, future research and areas of study will be recommended.

II. LITERATURE REVIEW

A. CLIENT SERVER

1. Computer Background

The Army developed the first computer known as the Electronic Numerical Integrator and Computer (ENIAC) during World War II. The ENIAC was designed so engineers can reduce calculating time to develop firing tables for ballistic trajectories from 30 minutes to 100 seconds. Once completed, the ENIAC consisted of 18,000 vacuum tubes, 1,500 relays, weighed 30 tons, and consumed 130 kW of power (Bergin, 2000). In order to load a program and receive the results of a calculation on the ENIAC, a punch card system designed by International Business Machines (IBM) was utilized that went to a relay system. The operational speed of the ENIAC was 100 kHz and was adjustable due to decreased performance of aging vacuum tubes. Due to the long setup time of the ENIAC which required hours of rewiring, the Electronic Discrete Variable Automatic Computer (EDVAC) was designed based on John von Neuman's architecture. Neuman's architecture was based on the "fetch-decode-execute" repeating cycle where instructions were fetched from memory, then decoded, and executed in a processor (Swedin & Ferro, 2005). The EDVAC had the ability to store programs and data equivalent to 2 KB by using mercury delay lines and had an operational speed of 1 MHz as it paved the way for future computers.

The Universal Automatic Computer (UNIVAC) was developed in 1951 by John Eckert and John Mauchly, who envisioned a machine that could be utilized by engineers, mathematicians, scientists, and businesses alike (Swedin & Ferro, 2005). The UNIVAC was considered faster and more reliable due to its combination of a magnetic tape media storage system and stored program architecture. Additionally, the UNIVAC contained fewer vacuum tubes and utilized solid state transistors which made it faster and more reliable. Due to magnetic tape technology, the amount of data storage capability increased due to the desire to run complex applications.

As specialized computers, the EDVAC, ENIAC, and UNIVAC had specific tasks designed primarily for the military, IBM started to control the commercial market by leasing computers such as the IBM 701 for $8100/month in the 1950s (Swedin & Ferro, 2005). The companies who leased the IBM 701 formed a group called SHARE, which allowed user groups to share programs to operate the IBM 701. SHARE forced IBM to create software products because each machine was the exactly the same and each program did not need to be specifically wrote for that computer (Swedin & Ferro, 2005). The SHARE program paved the way for IBM to receive consumer feedback and helped in development of desired programs for their customers.

As computers became more complex, Jack Kilby developed integrated circuits when employed for Texas Instruments in 1959 (Swedin & Ferro, 2005). The theory behind integrated circuits was to increase the operating speed of the transistors. With the transistors separated by numerous and longer wires, the computer would operate slower. Therefore, by putting all of the electronic components on a single board, it would increase the overall speed of the circuit. This eventually led to the development of microchips which were used by The National Aeronautics and Space Administration (NASA) and the U.S. Air Force for Minuteman intercontinental ballistic missiles (Swedin & Ferro, 2005).

The first microprocessor was developed by Intel Corporation in the early 1970s (Swedin & Ferro, 2005). Robert Noyce and Gordon Moore founded Intel and created 64-bit static random access memory (RAM) to replace magnetic core memory. The disadvantage of the RAM was if power was lost, the memory would be gone unlike the magnetic core memory, which retained its memory regardless of the power state of the machine. However, the static RAM chip was cheaper to manufacture with quicker operating speeds and became the memory of choice by computer manufacture.

The early development of the Internet was developed by Vanevar Bush, who had a vision for universities and the Department of Defense (DoD) exchanging information such as books and films on a system called “memex,” or memory extender. Bush wanted to review his records, books, and communications from a remote desk using nothing more than a keyboard and a set of translucent screens (O’Regan, 2008). In 1965, the

Advanced Research Projects Agency (ARPA) started to develop the Advanced Research Projects Agency Network (ARPANET), and the first wide-area network (WAN) was created in 1965 when a computer in Santa Monica was connected to a computer at the Massachusetts Institute of Technology (MIT) via a dedicated phone line (O'Regan, 2008). The ARPANET was based on packet switching networks when messages were sent on a network; the long messages were split into smaller packets to minimize network congestion. The ARPANET developed what is known as file transfer protocol (FTP) and e-mail today. FTP allowed a remote user to log into another computer and share files between two computers. Ray Tomlinson, who was one of the support engineers for ARPANET, learned that messages could be sent by combining programs named "CPYNET" and "SNDMSG" to users on the same computer and other computers which became the basis for e-mail (Swedin & Ferro, 2005).

2. Client Server Fundamentals

In a client server network, there are two basic parts, a server and a remote client. The server is a powerful computer on a network that consists of more physical processors, memory, and storage space compared to a personal computer, which are designed to facilitate communication and resource sharing between other computers on a network, known as clients. The clients are personal computers that run applications, save data to their local hard disks, and use shared applications on the server as needed. Communications among multiple computers on a network was accomplished through the Open Systems Interconnection (OSI) model. The OSI model was developed in the 1980s 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, 2009).

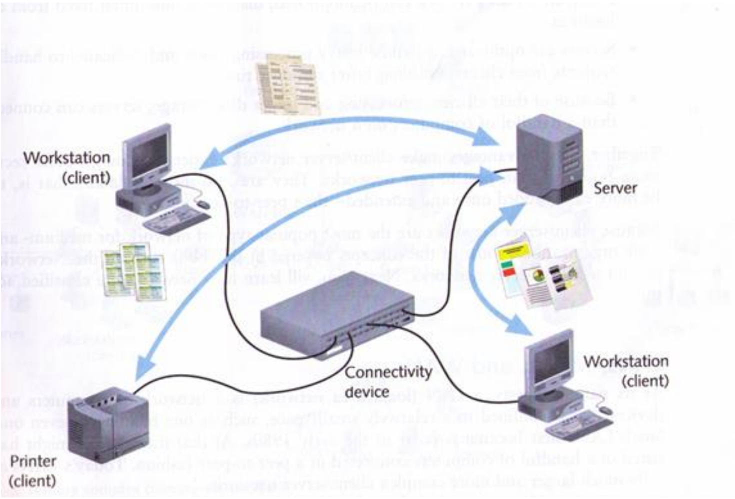

Figure 1. Basic Client Server Network (From Dean, 2009)

The seven layers of the OSI model coordinate how hosts on the same network can communicate with one another in the following manner:

- The Application layer is the top layer in the OSI model and it provides an interface between software application and network services. It does not provide data for such software applications such as Google Chrome or Internet Explorer. Instead it interprets an applications request to and from network so it can be properly formatted and utilized on the network (Dean, 2009).
- The Presentation layer is a translator for the Application layer data from and formats it into a common language so all hosts and applications can interpret it (Dean, 2009).
- The Session layer coordinates and maintains communications between two or more nodes on a network. A session is defined as a connection for ongoing data exchange between two parties (Dean, 2009).
- When data is split into packets for transmission, it needs to be correctly sequenced. Therefore, the Transport layer ensures data is accurately delivered between two hosts and delivered in the correct sequence with no errors (Dean, 2009).
- For data to be exchanged from one host to another, the Network layer translates the network address to the physical address of the host. The Network address is a unique number which is added to the data packets in order to properly route information (Dean, 2009).
- The Data Link layer packages data into frames to transmit on the Physical layer (Dean, 2009).

- The Physical layer manages the physical data transmission on the network through the use of a network interface controller (NIC) (Dean, 2009).

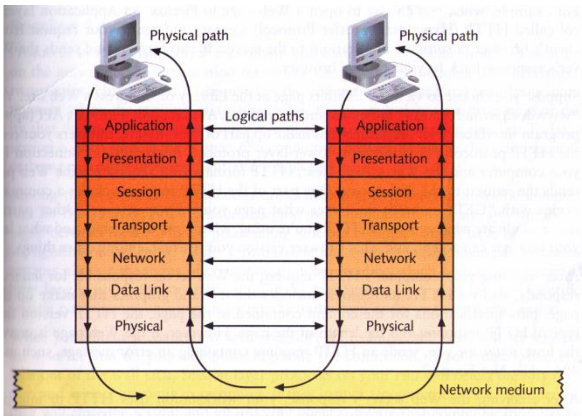

Figure 2. OSI Model Data Flow (From Dean, 2009)

In order for clients to communicate, each client must have a network adaptor or a NIC. The NIC is responsible for transmitting and receiving data on both the Physical and Data link layers of the OSI model because it contains the physical address of the client known as the media access control address (MAC address). The MAC address is a 12-character string which uniquely identifies the client on the address (Dean, 2009). Once data leaves the NIC, it must go through a connectivity device such as a switch or router to reach its destination. A switch divides a network into sub networks that can provide better security and performance since each device connected to the switch has its own dedicated channel. However, a switch has its disadvantages since it holds and buffers data when there are large data bursts on the network. If the switch becomes overwhelmed with data, it could lock up. A router is a multiport device that determines the best path for data travel based on network congestion by either using static or dynamic routing. Static routing is used by the Local Area Network (LAN) Administrator to manually program specific routes for the data. If the routes for the data become unusable, the router will become dysfunctional and data will be lost on the network. To prevent network disruption, dynamic routing is utilized and the router automatically determines the best path to the host based on network congestion (Dean, 2009).

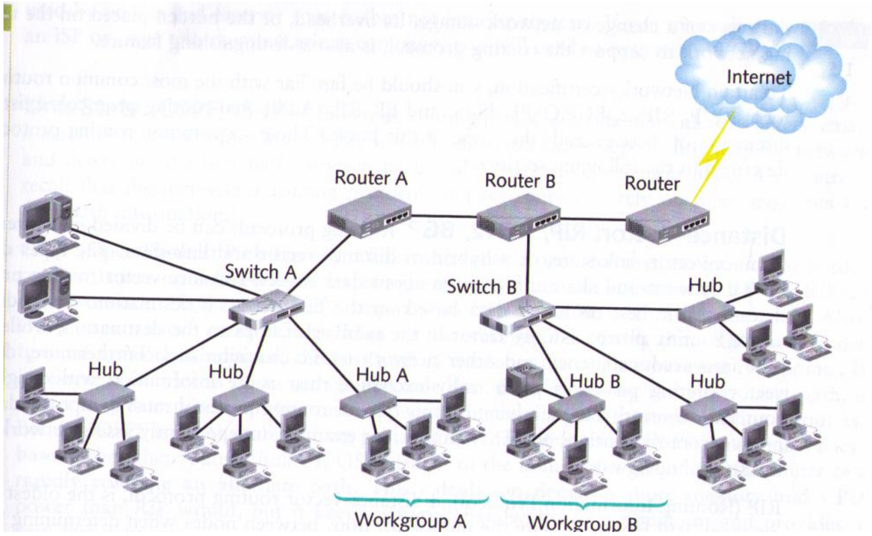

Figure 3. Routers and Switches in a Network (From Dean, 2009)

3. Thick Client vs. Thin Client

A thick client is considered a personal computer that is on a network. All of the processing and storage takes place on the personal computer and the server is used primarily for extra data storage or additional programs that are not on the computer's hard disk. Another advantage of a thick client is the ability to respond quicker while running applications due to its independence of relying on a server to run applications for the client. However, thick clients are considered tougher to deploy and update from a management perspective, since each machine has to be managed individually (John & Han, 2006).

Thin clients are nothing more than a solid state hard drive that has a browser-based application installed that acts as a graphic user interface (GUI) when connected on a network. Thin clients depend on a network connection in order to operate properly because all data processing is done on a centralized server. Due to the centralized server, thin clients are easy to update and deploy, since all operational and application software is on a single server. One major disadvantage of a thin client is that it relies on a constant network connection to operate. If the network is lost, the user may lose any unsaved data. Depending on the network congestion or how many users are utilizing the server simultaneously, thin-client performance will depend on how much server resources are available (John & Han, 2006).

4. Prior Research

Britt (2011) analyzed the advantages and disadvantages of the thin client-server architecture and web-based applications as viable solutions for the Department of the Navy (DON). He discussed how software and hardware developments have transformed the landscape for information technology (IT) so organizations can be competitive and efficient. Britt also identified TCO as an effective method of cost analysis to compare thick and thin client costs. Energy Consumption was identified as an effective means of a direct cost comparison using the average power consumed by a thick client versus thin client network architecture. Additionally, thick client-server architectures were analyzed and compared to the advantages vs. disadvantages of migrating to thin client-server architecture. The TCO of thin clients from the Gartner Research Group case study showed a cost savings of 32% over unmanaged thick clients and 3.3% over thick-managed clients. The Power Consumption of the thick-client with monitor tested was 170 watts versus an average of 70 watts for the thin clients with monitor (Britt, 2011).

Britt's analysis of thin client-server architecture revealed the advantages of migrating to Thin Clients facilitate the Department of the Navy Deputy Chief Information Officer's (DDCIO (N)) data center reduction and server utilization objectives by shifting computer processing to the server. Thin-client devices lack processing capability therefore it increases server utilization significantly, while reducing the overall Total Cost of Ownership (TCO). Britt states software and application licensing are one of the largest costs associated with thick clients. By migrating to thin client architecture, this allows software and applications to be moved to the server and the organization only pays for the server side licensing and the user access it needs. Additionally, the migrating of applications and software to the server also decreases the software maintenance and management burden on IT staff, further reducing costs (Britt, 2011).

B. VIRTUAL TECHNOLOGY AND CLOUD COMPUTING

1. Cloud Computing Background

There is no standard definition for “Cloud Computing.” Liang-Jie and Qun (2009) states that cloud computing has involved from a set of key technologies to address resource sharing based on business requirements. The first computer resource sharing started when computers were “clustered” in order to form a supercomputer. This allowed one computer to configure the clustered computers via protocol and conduct a balanced computational load across the machines (Rittinghouse & Ransome, 2010).

The concept of grid computing was developed in the 1990s by Ian Foster and Carl Kesselman who envisioned selling computer processing similar to an electrical utility company. Grid computing was thought to be a cost-effective solution for companies who would rather lease computing resources vice purchase and manage their own computer resources (Rittinghouse & Ransome, 2010). The disadvantage of grid computing is where the physical processing and data resides. Since some data could be thousands of miles from the computer, it could have data latency and cause the computer to run inefficiently and have significant data delays. Similarly, in 2000, Volunteer Computing was developed to share computer resources over the Internet for science research. However, the owner of the computer was not responsible for the results of their machine since the resources were volunteered (Kondo & Bahman, 2009).

One of the first cloud computing services was Amazon’s S3 (Simple Storage Service), which debuted in 2006. Amazon’s storage service was a web service that users could store and retrieve data from anywhere in the world as long they were connected to the Internet. This service by Amazon has suffered some setbacks such as in February 2008, when its servers became overloaded with service requests and the cloud service temporary failed (Rittinghouse & Ransome, 2010). Even though the idea of trusting another company with one’s data seems feasible, the user is at the mercy of the cloud service provider for securing and accessing stored data.

2. Virtualization Fundamentals

Virtualization is the ability to run multiple operating systems on a single computer simultaneously. In order to run this, there needs to be a control program such as Parallels, which allows a Windows operating system to run on an Apple MAC computer (Rittinghouse & Ransome, 2010). Programs such as Parallels or VMware create what is known as a virtualization layer, which manages and hosts the virtual machines. The virtualization is utilized in either hosted or hypervisor architecture. In a hosted architecture, the virtualization runs on top of the operating system and supports hardware configuration. For a hypervisor architecture, the virtualization layer is installed on a clean computing system where this is no host operating system to compete for computing resources (Li & Jiang, 2010). The hypervisor architecture in essence controls how the hardware and computer resources are used for each guest operating system and imitates the physical ports such as video, network, and printer ports for each virtual machine.

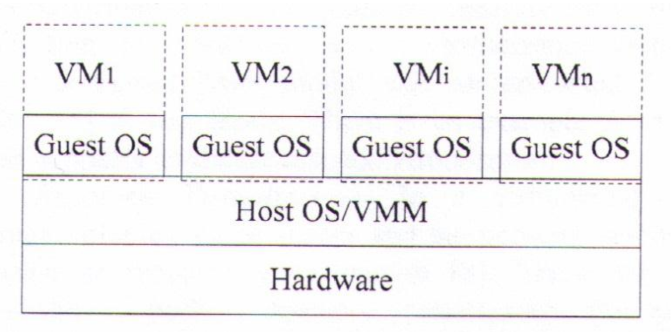

Figure 4. Virtual Machine Architecture (From Li & Jiang, 2010)

3. Cloud Computing Services and Models

There are different models of cloud computing to fit individual needs. Cloud computing can be as simple as providing remote storage backup or provide a complete network that is accessible anywhere over the Internet. The following are different services and models available for cloud computing (Strowd & Lewis, 2010):

- Infrastructure as a Service (IaaS) includes computing power, storage, operating systems, and networks in a virtualized environment. One advantage of this is that everything is centralized and only the main server software is upgraded vice upgrading all of the clients. An example of this is Amazon’s EC2 cloud service. For a fee, a user can choose their operating system, amount of memory, processor speed, and subnet a network if desired. The user saves money by not having to purchase operating disks, licenses, and IT staff to maintain the network (Strowd & Lewis, 2010).

- Platform as a Service (PaaS) is an application development that links hardware and software components so users can leverage resources of established organizations to create and host applications of larger scale (Strowd & Lewis 2010). In essence, it allows users to use third party applications over the Internet without having to purchase the hardware and software. Google App Engine provides this type of cloud service where e-mail, chat, and video services are provided over the Internet vice having to purchase software and load it on their computers.
- Software as a Service (SaaS) provides users a specific program to utilize over the Internet without having to purchase the software. The platform is provided and the user can use software remotely on the Internet without installing it on their computer and the user just pays a fee for the service. Examples of SaaS providers are Turbo Tax and Sales Force (Strowd & Lewis, 2010).
- When the organization owns the computing resources, software, and is utilized behind a firewall to prevent unauthorized outside access, this is considered a Private cloud. The Private cloud model does not rely on Internet access and is fully managed by the organization who wants more control of their data and network (Strowd & Lewis, 2010).
- Public cloud computing is when a third party controls the cloud resources so external users can utilize the computing resources for a fee over the Internet.

4. Risks and Benefits of Cloud Computing

For enterprises that consider upgrading from a client-server architecture to a cloud environment, security is a key concern. The cloud network is exposed to both old and new vulnerabilities yet to be seen. But cloud computing does have security advantages. One advantage of cloud computing is the centralization of data for a private cloud network. The network managers can upgrade components, monitor hardware, and encrypt data since it has a smaller foot print and the IT infrastructure is essentially in one space. This also helps with physical security to ensure the equipment and data are secure. The ability to mirror data and have multiple storage locations is another advantage for cloud computing. Google for example stores data in three physical separate locations which allows for greater redundancy and to assist in data recovery.

Another advantage is the rapid ability of security automation for auditing patch management. Patches can be tested easier before being released due to virtual technology. Additionally, the patches can be released to the virtual users just as quick (Antedomenico, 2010). Public cloud computing has potential security savings since its maintained by a third party and saves operating costs, it does have deficiencies such as relying on those same cloud providers for securing data.

Google has 36 data centers across the world (Minqi, Ron, Wei, & Aoying 2010). A copy of one's data can be stored in U.S., Canada, and possibly China. Not knowing the actual physical location of your data nor the physical security in the storage location creates great concern, since the data center may catch on fire, flood, or be compromised. Plus, other countries have different laws which may give them access to a user's data if it is stored in their country.

Currently, in the United States, there are no Federal laws specifically protecting data in the cloud network. The Electronic Privacy Act of 1986 is the closest thing keeping data from being accessed by the government. With the passage of the Patriot act in 2001 and amended in 2005, the FBI can access stored data, but will need a Federal court order. To protect data from this problem is to ensure the entity has a user agreement with the cloud service provider and to encrypt the data before storing it in the cloud. Even though cloud storage service is similar to a bank, the data is not insured in the same manner cash is protected in a bank and the FDIC. Additionally, if a cloud provider goes out of business or is in dispute in another country, their laws may enable them to review stored data if it is stored in their country.

In addition to the physical storage, virtual technology also has deficiencies that can corrupt a cloud network. A virtual computer's operating system is vulnerable to malware, viruses, Trojan horses, and worms like a personal computer. Therefore, virus protection is needed for each virtual client operating in the cloud. This leads to possible cloud cross contamination. For example, if two cloud users were utilizing the same server and one did not protect their software and caught a virus, that virus can negatively affect the cloud's performance because the virtual layer is now infected. The virtual technology software is also vulnerable to worms and viruses to take down a cloud host and lead to destruction of data in a cloud network. There are Virtual Machine Intrusion Detection Systems available to protect the virtual layer from possible virus infection. But it is unknown if the cloud provider is protecting their virtual layer to save on costs and not provide the maximum protection for their users' data (Li & Jiang, 2010).

Cloud computing is economical from a business model. Companies who use cloud services save on capital costs of not investing in an IT network and staff to maintain it. Cloud computing allows smaller businesses to setup a LAN with basic Internet, or use their PDAs to access everything they need in a cloud environment. Since everything in a cloud environment is service for a fee, companies do not need to purchase expensive software, maintain the licenses, and purchase expensive servers costing upward of $1,000 depending on the server capacity and rating. Plus some cloud services also provide software that caters to the smaller business saving anywhere $50 to $1000 per licensed copy. With a virtual network provided by an IaaS cloud, office space is better utilized by not having a dedicated server room. That room utilized for the server can be used for additional inventory or save a company in rental costs because of not having the extra floor space (Colaner, 2010).

The last potential savings of a cloud network is energy costs. When thin clients are used, they consume no more than 78 watts of electricity compared to a desktop computer that utilizes up to 250 watts of electricity when operating at full capacity. Depending on the type of cloud service or model being utilized, these savings can be minimal as discovered in a research test (Baliga, Ayre & Hinton, 2010). The test discovered that a public cloud consumes three to four times more energy than private cloud networks due to consumption of transport of data. But these energy costs are still lower than saving data on a local hard drive. It was determined that if low end laptops were used for routine tasks and cloud services for heavy computation tasks, that there would be approximately 13% in energy savings (Baliga et al., 2010). However, under some circumstances though, the cloud can consume more energy than using a desktop computer as argued by Jennings (2010). Jennings (2010) suggested that cloud computing only transfers energy from the customer to the cloud service provider in order to power the servers and supporting equipment in order to meet user demand. By increasing demand, the energy consumption will increase for the provider (Jennings, 2010). With advances in virtualization technology, energy consumption can be reduced and it depends on how the cloud is setup and utilized to maximize potential (Baliga et al., 2010).

5. Prior Research

Lam (2010) showed, that, by deploying a thin client desktop and virtualization-based server solution, the TCO over the next seven years would be lower than that of the current plan for thick client computers for the U.S Navy's OCONUS network ONE-NET. Since building a software image is identical for both a thin and thick client, there is no change for the operational and security requirements. Lam (2010) stated that security requirements are robust for thin clients since the data installation is done in a central location. In Table 1, Lam (2010) showed the overall TCO breakdown between a thin and thick client for ONE-NET.

Table 1. TCO Comparison Between Thick and Thin Clients (After Lam, 2010)

Alternative Solutions	Labor $M	Hardware $M	Software $M	Transport $M	Power & Cooling $M	Virtualization $M	TCO $M
Thick Client	$545	$97	$21	$125	$15.4	$0	$803
Thin Client	$335	$39	$21	$125	$3.8	$41	$565

Even though there is an additional $41M investment for virtualization, the hardware, labor, and power savings make up the difference and still reduced the overall TCO vice staying with the Thick client network. With these savings, Lam (2010) stated that thin client solutions are feasible with the current infrastructure in place and thin client technology coupled with virtualization can reduce TCO of an enterprises network (Lam, 2010).

Tiglao (2010) researched the application of VM technology to create models and simulations of current IT capabilities used by military operating forces. Tiglao (2010) determined that DoD C4I Support Centers would inherit much of the costs for virtualization and cloud operations, making research and development costs expensive for initial investments of a new system.

C. CONSOLIDATED AFLOAT NETWORKS AND ENTERPRISE SERVICES (CANES)

1. Background

There are 64 legacy systems onboard U.S. Navy vessels under the current IT-21 program. Those legacy systems have 17 variants of hardware, 6 separate operating systems, and 380 applications of software used for various tasks on a daily basis. With various networks on a single ship, none of them communicate with one another and create “stovepipes” of information (Rognlie, 2010). The CANES program is designed to incorporate five legacy networks into one functional network. The stated goals of CANES are as follows:

- Build a secure afloat network required for Naval and Joint operations.
- Consolidate and reduce the number of afloat networks through the use of mature cross-domain technologies and Common Computing Environment infrastructure.
- Reduce the infrastructure footprint and associated costs.
- Provide increased reliability, application hosting, and other capabilities to meet current and project requirements.
- Federate Net-Centric Enterprise Service and Afloat Core Services to the tactical edge to support overall Department of Defense Command, Control, Communications, Computers, Intelligence, Surveillance, and Reconnaissance applications migration to a Service Oriented Architecture environment (Rognlie, 2010).

The legacy networks which CANES will incorporate are the Integrated Shipboard Network System (ISNS), Submarine Local Area Network (SUBLAN), Combined Enterprise Regional Information Exchange (CENTRIXS-M), and the Sensitive Compartmented Information Local Area Network (SCI LAN). The ISNS is the current legacy network in the fleet and provides basic LAN services afloat. ISNS supports both UNCLASSIFIED and TOP SECRET classification levels on a ship by utilizing an air gap network to prevent electronic spillage of classified information on an unclassified network. When data is routed off the ship or needs bandwidth, the ADNS server is used. The ADNS server is a router that provides ship-to-shore and ship-to-ship communication for voice and data information via a satellite connection (Rognlie, 2010). For Special Intelligence data, the SCI LAN is able to handle data, voice, and video information. The SCI LAN is currently a separate network afloat and meets the U.S. Navy's security requirements. SUBLAN is a similar network to the ISNS that is designed for the submarine force and handles the same classification levels (U.S Navy, 2005). Lastly, when operations are conducted with foreign allies, the CENTRIXS-M network is a separate network that is utilized. CENTRIXS-M provides e-mail, web, and chat capabilities with foreign partners without compromising data. This network is also routed through the ADNS terminal (U.S Navy, 2005).

2. Elements of CANES

In order to meet the program goals of CANES and integrate the legacy systems into one functional network, the following three elements are critical for incorporating into CANES: Common Computing Environment (CCE), Cross Domain Solutions (CDS), and Afloat Core Services (ACS).

The CCE consolidates a ship's network hardware into a common networking core. By having a common core, this reduces scattered duplication and creates centralized management of software and program updates. The CCE also hosts legacy applications that are virtualized without having to maintain older equipment in order to run effectively (Rognlie, 2010). One of the advantages of the CCE according to Rothenhaus (2011), by centralizing the computing and storage to more focused nodes reduces network traffic as the number of hops from server to server or off-ship enclave router is reduced. Therefore, the need to update infrastructure such as switches and networking cable to handle larger data is not needed since the data is moved internally at the centralized server.

The next element is the CDS which allows unclassified and classified information to be viewed on a single workstation (Rognlie, 2010). CDS will be running a trusted Solaris session server as the hosting environment to partition the enclaves. According to

Rothenhaus (2011), the user will be able to view classified and unclassified data simultaneously, but users will not be able to cut and paste information between sessions. This is currently being utilized on a U.S. Navy CENTIX system today which fields a Multi-level tactical client that is actually a zero client (Rothenhaus, 2011).

The ACS is the last element and uses a service oriented architecture to decouple hardware from designated software. By decoupling the hardware and software, software developers can now develop software efficiently without having to rewrite and can utilize existing plug-in solutions (Rognlie, 2010).

3. Manpower Reduction with CANES

In 2009, RAND researched and reviewed the current legacy systems and the drive for specialized network administration training to maintain those systems for the Information Systems Technician (IT) and Electronics Technician (ET) ratings. It was determined that one IT Navy Enlisted Classification (NEC) Sailor can be reduced for a DDG and possibly four IT Sailors for a carrier due to the integration of legacy systems in the new CANES program. However, with the influx of new technology, RAND recommended that IT's have longer "A" and "C" schools so technicians would have a Level One IA certification in order to successfully meet the needs of CANES (RAND, 2009). With the increased training time of IT's and new curriculums that need to be developed, this may temporary increase the TCO of CANES to support manpower needs.

4. Estimated Energy Savings

With the consolidation of five legacy networks into one single network, this presents the opportunity to reduce operational costs while in port and at sea. In an early analysis by PEO C4I, the consolidation of the legacy networks could remove eight server racks from an Arleigh Burke Class Destroyer. Each rack is estimated to use 3.3 kW of power. With each rack removed, a ship would save over 3,204 gallons of fuel every year assuming continuous operation (Rognlie, 2010). The fuel savings are based on using an Allison AG9140 Gas Turbine Generator that currently generates electricity for Ticonderoga Class Cruisers and Arleigh Burke Class Destroyers. The fuel usage rate of the AG9140 gas turbine generator is 15,375 BTU/kW-hr and the energy content of DFM is 138,700 BTU/gallon, which equates to saving .365 gallons/hour with each rack removed. In port, these power savings would be approximately $2,890 if in full operation and electricity priced at $.10 kW-hr for each server rack removed. However, these assumptions are considering that the server is under constant load in order to operate at 3.3 kW. There is also additional savings that are not calculated by replacing thick clients with thin clients that consume less energy.

5. Prior Research

Rognlie (2010) calculated the Return on Investment (ROI) for CANES in a thesis. He calculated the ROI for CANES to be at 73% for a service oriented architecture in which CANES is modeled after. For CANES to remain at 73%, manpower reduction is the sensitive variable in order to gain such a high ROI value. However, the ROI is not sensitive to changes in installation costs, phase-out costs of current systems, and costs associated with equipment refresh hardware and software of CANES ships (Rognlie, 2010).

THIS PAGE INTENTIONALLY LEFT BLANK

III. RESEARCH METHOD-SHIPBOARD AND SYSTEMS APPLICATIONS

A. TOTAL COST OF OWNERSHIP AND NAVY VISION

1. Secretary of the Navy's Vision for Energy Conservation

In FY2005, the U.S. Navy reported that it used 41 million barrels of oil for all purposes during that operating period (U.S. Library of Congress, 2006). Of that budget, $2.83 billion is reserved for fuel, and the U.S. Navy utilized 32% of the budget purchasing fuel for its aircraft and ships (U.S. Library of Congress, 2006). Since then, the U.S. Navy has taken measures to reduce fuel consumption across the fleet. One instance of the U.S. Navy reducing fuel consumption is by installing an efficient hybrid propulsion plant on the newly commissioned USS Makin Island (LHD-8) (Mabus, 2011).

On October 13, 2011, Secretary of the Navy Ray Mabus discussed the Navy's vision for energy conservation at the Navy Energy Forum in Washington, D.C. Mabus's (2011) vision for the U.S. Navy by the year 2020 is that half of the command's energy sources, both afloat and ashore, would come from non-fossil-fuel sources. The reason for the decline in fossil-fuel sources is for strategic purposes, since the U.S. Navy purchases fuel from potentially unstable countries. If fuel was not purchased from these unstable countries, the military would not be able to conduct operations and meet the assigned mission (Mabus, 2011). In addition to reducing usage of fossil-fuels, Mabus (2011) is also changing how contracts are awarded and will hold contractors accountable to ensure energy efficiency targets of new equipment are as advertised.

2. Total Cost of Ownership Background

The Gartner Group developed the concept of Total Cost of Ownership (TCO) to determine the direct and indirect costs for the life cycle of a project (West & Daigle, 2004). TCO is a tool to help organizations manage and determine if projects are worth investing and assist in managing overrun costs during the life cycle. In order to develop a successful TCO model, more items that are measured and managed ensure success during the life cycle of a system.

When determining life-cycle costs for a TCO model, the following five phases must be considered: acquisition, implementation, operations, maintenance, and replacement. By predicting costs over these five phases, decision makers can manage future budgets. One flaw is that acquisition costs drive decisions about implementation of new projects and the remaining phases are ignored causing contingency costs. Therefore, it is important to evaluate all phases of the TCO model to ensure proper project and budget management (West & Daigle, 2004).

The identification of direct and indirect costs is important in development of a successful TCO model. Direct costs are budgeted and tangible. Hardware, software, maintenance, labor, recycling, and research are good examples of direct costs. Indirect costs are unbudgeted costs such as downtime costs and affect the end user. Once the direct and indirect costs are identified for the five phases of the TCO life cycle, budget managers can run "what if" simulations and conduct a TCO analysis to determine a possible implementation strategy (West & Daigle, 2004).

B. U.S. NAVY ENCON AND NEURS PROGRAM

1. U.S. Navy's Incentivized Energy Conservation Program

The Incentivized Energy Conservation (i-ENCON) program was established in 1993 for the Pacific and Atlantic Fleet ships and was implemented fleet wide in 1999 in order to reduce energy costs onboard U.S. Navy vessels. The i-ENCON program is a guide for shipboard personnel on how to reduce energy consumption in operational environments. By reducing fuel consumption, ships benefit by conducting underway replenishments less frequently, less maintenance, and can earn cash rewards that go toward their Optimal Scheduling and Operating Target (OPTAR) budget (NAVSEA, 2010).

Each year, the Secretary of the Navy presents energy conservation awards to two categories of ships based on hull size. The Large Hull category is for crews greater than 400 personnel and the ship will receive $30,000 that can be used their budget with no restrictions. For the Small Hull category of crews less than 400 personnel, the same award is $20,000. Even though a ship may not win these awards, a ship that under burns its respective fuel allocation may earn cash from $1,000 to $50,000 that goes into their OPTAR budget and is awarded by the fleet Type Commander (TYCOM). Ships that make outstanding contributions in conserving both energy and water mandated by the Energy Policy Act of 1992 and Executive Order 13123 of June 1999 are eligible for the Department of Energy (DOE) award. The ship will receive an award and plaque from the DOE (NAVSEA, 2010).

2. Navy Energy Usage Reporting System

The Navy Energy Usage Reporting System (NEURS) is governed by OPNAVINST 4100.11C and it details how ships will report their fuel usage (DON, 2007). The NEURS report is a monthly report that must be submitted by all afloat commands to their TYCOM detailing fuel inventory, sale, and usage over a monthly period. The TYCOM is responsible for monitoring the fuel consumption and conservation within their fleets. Besides being tracked by the TYCOM, the NEURS report is utilized by all levels of Navy management to formulate energy policy, measure energy conservation progress, support operational scheduling, generate budget requirements, and assist in ship life-cycle cost estimates and analysis (DON, 2007).

C. GAS TURBINE AND DIESEL GENERATORS

1. Ships Service Gas Turbine Generators

The U.S. Navy currently has 22 Ticonderoga Class Cruisers and 60 Arleigh Burke Class Destroyers that use Allison Ships Service Gas Turbine Generators (SSGTG) to provide electrical power throughout the ship (U.S. Navy, 2012). Each ship is equipped with three SSGTGs and each generator provides continuous 2500 kW, 440 volts, three-phase power (Rolls Royce, 2008). The SSGTG utilize the Brayton Cycle in order to produce power and turn the coupled generator. Air is first compressed in a fourteen-stage axial compressor which feeds six flow-through canular combustors where fuel and compressed air is ignited. After combustion, the exhaust gas drives a two-stage turbine that is coupled to the generator (Lane, 2001). See Figure 5.

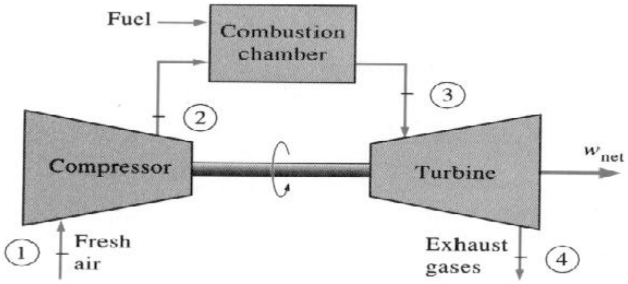

Figure 5. Basic Brayton Cycle (From Lane, 2001)

When the ship is underway or auxiliary steaming in port, how much fuel each generator burns in relation to the ships electrical load is shown in Figure 6. For every 100 kilowatt-hour increase in electrical load, fuel consumption will increase approximately eight gallons per hour. Ships will generally run two out of the three SSGTGs for redundancy purposes in case one generator should fail and prevent the ship from going completely “dark.” Additionally, the curve for “bleed air” is not used during normal underway steaming unless direct by Engineering Operational Sequencing System (EOSS). In the event bleed air is utilized, the compressor portion of the SSGTG will spin at higher operating speeds to compensate for the loss of cooling air for the turbine when build bleed air is utilized, therefore decreasing fuel efficiency of the SSGTG.

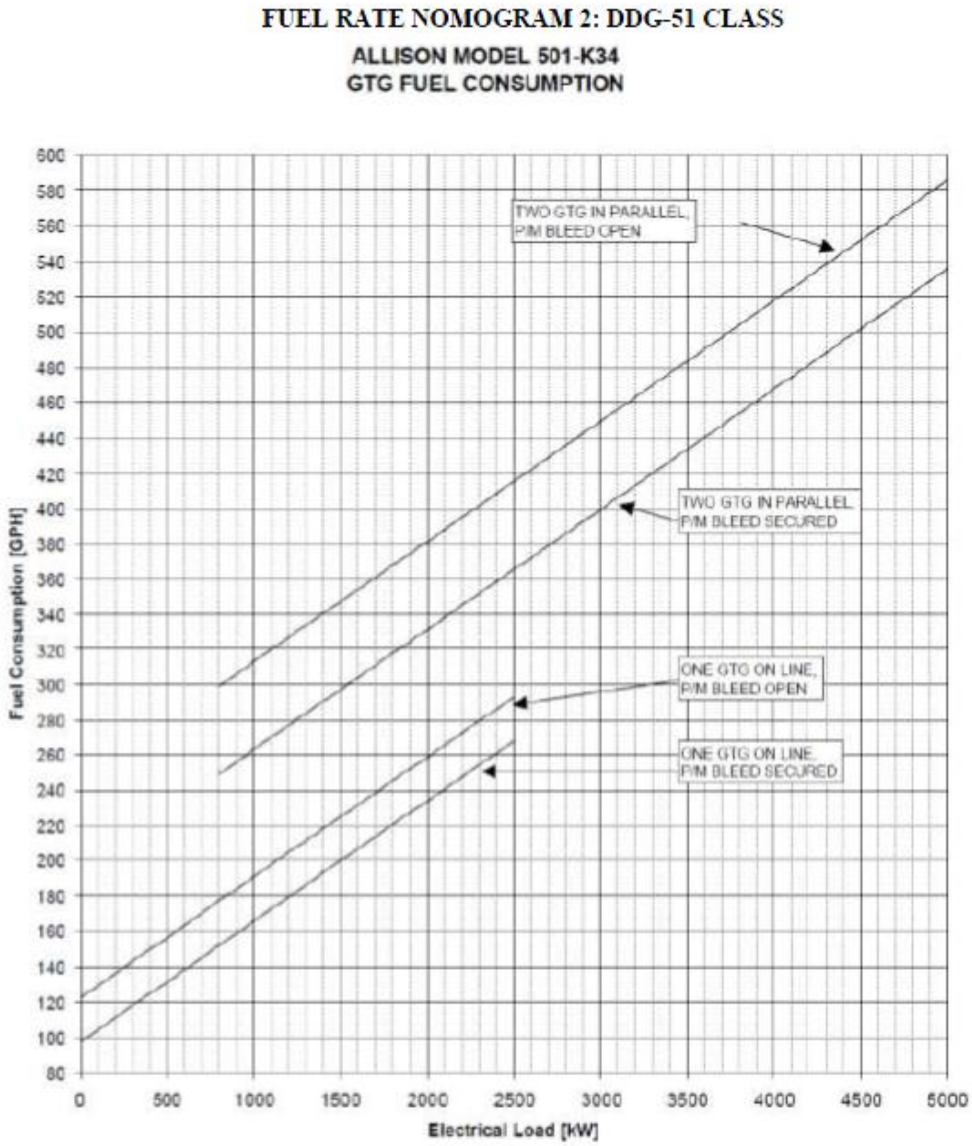

Figure 6. Allison SSGTG Hourly Fuel Burn Rate (From NAVSEA, 2010)

2. Ships Service Diesel Generators

For ships without advanced weapons and radar systems that demand high energy requirements provided by SSGTGs, these ships are equipped with Ships Service Diesel Generators (SSDG). The U.S Navy currently has 47 ships equipped with SSDGs, ranging from frigates to the newest amphibious assault ship USS Makin Island (LHD-8) (U.S. Navy, 2012). SSDGs are considered more fuel efficient compared to SSGTGs, but lack the ability to produce and maintain high electrical loads unlike a SSGTG. Figure 7 shows the hourly fuel usage of a set of SSDGs on a frigate. For every 100 kilowatt-hour increase of electrical power, the SSDG fuel consumption will increase seven gallons per hour.

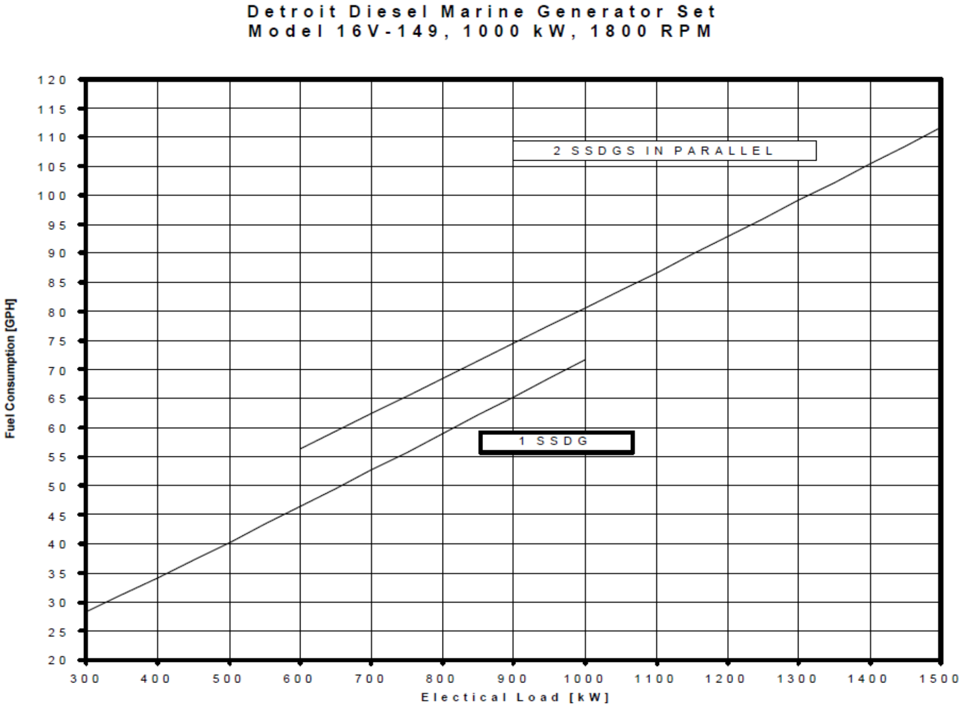

Figure 7. Frigate SSDG Fuel Curve (From NAVSEA, 2010)

D. INTEGRATED SHIPBOARD SYSTEM NETWORK

The Integrated Shipboard System Network (ISNS) is the main network on ship that Sailors use to conduct their daily operations. On a ship with a crew less than 400 people, there are four AN/USQ-153 servers that have unique functions. Each AN/USQ-153 server rack contains a set of blade servers, storage devices, and appropriate power supplies. Two of the AN/USQ-153 servers are used for the UNCLASSIFIED and CLASSIFIED networks that are air-gapped. These server racks are used primary for e-mail and network storage so the crew can store technical manuals and required publications.

The Naval Tactical Command Support System (NCTSS) server hosts a set of applications to assist in logistics information, personnel, maintenance, and finances in order to operate the ship (SPAWAR, 2012). The following is a description of each program:

- Relational Administrate Data Management (R-ADM) is a database used to manage personnel and create the appropriate reports. R-ADM can assist users in managing watch bills, gun qualifications, training records, and recall information for the whole command (SPAWAR, 2012).
- Relational Supply (RSupply) provides the tools so personnel can order and receive parts, maintain financial records, conduct inventory, and track the status of ordered parts (SPAWAR, 2012).
- Organizational Maintenance Management System-Next Generation (OMMS-NG) is used by maintenance personnel to ensure ship readiness. OMMS-NG is used to open jobs, track maintenance hours, order parts, and request outside assistance to fix equipment. OMMS-NG also interfaces with RSupply for tracking parts and checking onboard inventory (SPAWAR, 2012).

The Navy Information Application Product Suite (NIAPS) server is a private web based server that is used for training and management applications. The NIAPS server is useful for conducting Internet based General Military Training when there is limited bandwidth and replicates data to shore based servers when bandwidth is available or in port. Additionally, NIAPS also has a Distance Support website where Sailors can ask for assistance in troubleshooting equipment (SPAWAR, 2012).

E. PRIVATE CLOUD MODEL AND VMWARE BASICS

With new procurements such as the Consolidated Afloat Networks Enterprise System (CANES) under development, it is important to understand how a basic private cloud functions. The hardware and software of a private cloud network at Naval Postgraduate School (NPS) will be reviewed to comprehend how virtual technology functions.

1. Chassis and Power Supply

The basis of the infrastructure for the NPS private cloud is the Dell PowerEdge M1000e modular chassis blade enclosure. The M1000e chassis can hold a maximum of

16 half-height blade server modules, six network and storage input/output (I/O) interconnect modules, and has comprehensive I/O options that support dual links of 20 GB/s connectivity between storage arrays and the network (Loffink, 2008). This allows for scalability and flexibility between server modules and future upgrades.

The Chassis Management Controller (CMC) provides management access to the chassis and installed modules from a local or remote location (Loffink, 2008). With direct management, administrators can manage firmware, firewall traffic, and view real time information such as chassis temperature and power usage. In addition to chassis management, the Integrated Dell Remote Access Controller (iDRAC) function is used to manage each server enclosure individually. The iDRAC incorporates the use of a Baseboard Management Controller (BMC) so a Virtual Keyboard/Video/Mouse (vKVM) and Virtual Media can interface with each module over the network. A vKVM is a remote user’s keyboard, monitor, and mouse that are used to interact with the individual server blade. With the iDRAC assigned to an Internet Protocol (IP) address, remote power management of the server enclosure can be accomplished manually or automatically depending on the software (Loffink, 2008).

Figure 8. Dell PowerEdge M1000e Chassis and CMC (From Loffink, 2008)

The servers interface with one another through the use of a midplane and fabric I/O integration. The midplane is a large circuit board that provides power distribution, system management, and fabric connectivity. The I/O fabric integration includes storage, networking, and interprocess communications (Loffink, 2008). The following are basic terms associated with the different fabrics used in the M1000e chassis:

- A fabric is a method of encoding, transporting, and synchronizing data between devices (Luffink, 2008). The fabrics are used between the server module and I/O modules through the midplane.
- A lane is a single fabric data transport between I/O and end devices (Loffink, 2008).
- A link is a collection of multiple fabric lanes used to form a single communication transport path between I/O end devices (Loffink, 2008).
- A port is a physical I/O end interface of a device to a link with single or multiple lanes of fabric I/O connected to it (Loffink, 2008).

In the M1000e chassis, there are three multi-lane fabrics utilized for communications and make up the high-speed I/O architecture. Fabric A is dedicated to the Gigabit Ethernet. The midplane can support up to four Gb Ethernet links per server module on Fabric A for a potential data bandwidth of 4 Gb/s per server module. Fabrics B and C are identical and fully customizable fabrics. Fabric A and C are routed as two sets of four lanes from the mezzanine cards on the server modules to the I/O modules in the rear of the chassis. The supported bandwidth ranges are from 1 to 10 Gb/s per lane depending on the fabric type used. Figure 9 illustrates the individual paths of the high-speed I/O architecture.

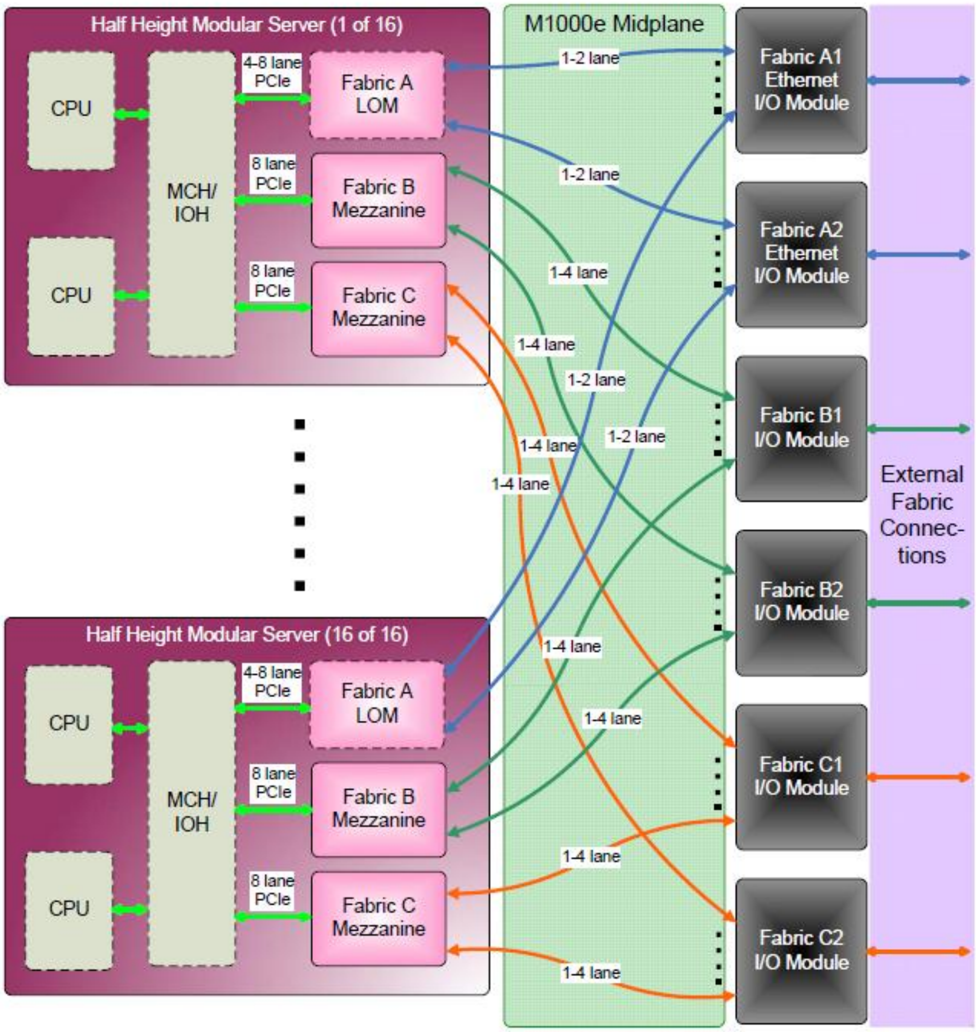

Figure 9. High Speed I/O Architecture (From Loffink, 2008)

Energy conservation of the M1000e chassis is enabled by the Dynamic Power Supply Engagement (DPSE) feature that will toggle power supplies on and off depending upon the actual power draw to maximize energy usage. Prior to any blade being powered on manually or by the Wake on LAN feature (WOL), the CMC will take a power budget inventory from each blade server's iDRAC based on its configuration. Once the information is received, the CMC will coordinate with the power supplies before powering on the blade server. The CMC is important in controlling the power usage of the M1000e chassis and it can also set power priorities for each blade enclosure as needed. In addition to managing power, the CMC can perform "power capping" and throttle individual servers and reduce the performance of the server if excessive power is drawn (Dell, 2010).

To fully power the M1000e chassis, each power supply receives power from the Power Distribution Unit (PDU). The PDU receives input power from either single-phase or 3-phase power configurations that feed three 2360 Watt +12 VDC power supplies (Luffink, 2008). Increased redundancy is available in the event of a power outage or power supply failure by utilizing six 2360 Watt power supplies as shown in Figure 10.

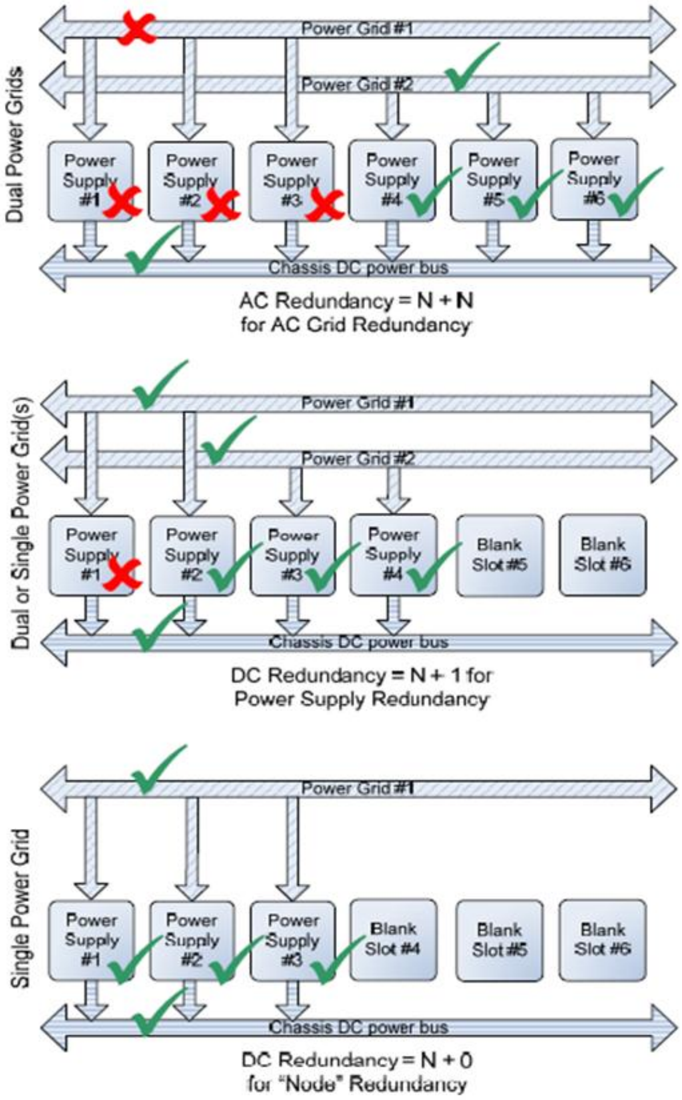

Figure 10. Power Redundancy Configuration (From Loffink, 2008)

The following are the possible configurations of the 2360 Watt power supplies:

- The N+N configuration provides maximum system protection against input power loss. Input power is provided from two different grid sources and if a grid source is lost, the other power supplies would be able to maintain operation (Luffink, 2008).
- The N+1 configuration provides protection only against power supply failures. If input power is lost, the chassis will lose power (Luffink, 2008).
- The N+0 configuration provides no input power or backup power protection (Luffink, 2008).

2. Blade Server Module Specifications

There are nine Dell PowerEdge M610 power blade servers in the M1000e chassis. Each blade server has two 2.5″ Solid State hard drives, 24 gigabytes (GB) of system memory, and two Intel Xeon quad-core processors (E5540 @ 2.53 GHZ). To boost virtual technology capabilities, the M610 can handle a total of 192 GB of system memory (Dell, 2011).

The Intel Xeon E5540 processor has a base clock speed of 2.53 GHZ and can be increased to 2.8 GHZ if the Intel Turbo Boost option is enabled in the server's Basic Input/Output System (BIOS). By increasing the clock speed, it increases the processor's operating capacity from 40.5 Giga Floating Operating Points Per Second (GFLOPS) to 45 GFLOPS (Intel, 2011e). A FLOP is used to measure the performance of a computers processor (Tech Terms, 2012). Processors are normally characterized by their clock speed, but a FLOP measures how many calculations the processor can accomplish within a second. With the E5540 processor having four cores, or four individual processors in one chip, each core can process 10 GFLOPS making it ideal for virtual technology since each core can be assigned to an individual virtual machine and efficiently handle any assigned task.

Virtual technology on the Intel Xeon processor has the ability to maximize system resources by abstracting the software from the hardware. This is enabled by processor extensions that optimize virtualization software efficiency by improving interrupt handling of the processor (Righini, 2010). With improved interrupt handling, processor resources can either be shared among tasks or allocated for specific tasks. This will ensure that proper processing power is always available.

With the current configuration, the private cloud at Naval Postgraduate School has 72 processing cores and 216 GB of system memory by utilizing only nine M610 server blades. With proper configuration, this private cloud can host 72 or more virtual clients and properly manage them based on user demand. The additional specifications of the M610 server blade are displayed in Figure 11.

3. Storage Array Network

To store virtual machine profiles and user data storage, a Dell Internet Small Computer System Interface (iSCSI) Storage Array Network (SAN) is utilized. The iSCSI is attached to a set of stacked switches in slots B1 and B2. This will coincide with Fabric B, which is attached to each server enclosure.

4. Thin Clients

A Samsung Syncmaster NC240 and Wyse P20 will be used as thin clients for the private cloud to compare the advantages and disadvantages of each. The Samsung Syncmaster uses Personal Computer over Internet Protocol (PCoIP) technology that sends encrypted pixel data over the network that is then decoded and displayed by the monitor (Samsung, 2009). Additionally, a keyboard and mouse can plug directly in the Syncmaster without needing another peripheral device to connect. Due to its simplicity of utilizing a solid state hard drive to hold display program information, the Samsung Syncmaster is completely independent of any application, host, or client. This makes it ideal for increasing the life cycle of a network because the firmware only needs updated. By only upgrading the firmware, the monitor can be used for any future system upgrades.

The Wyse P20 also utilizes PCoIP technology and has a 128 MB solid state hard drive that holds program information (Wyse, 2011). Unlike the Syncmaster, the P20 needs an external monitor to display received data. However, the P20 can save on replacement costs of legacy systems by utilizing the current keyboard, mouse, and monitor.

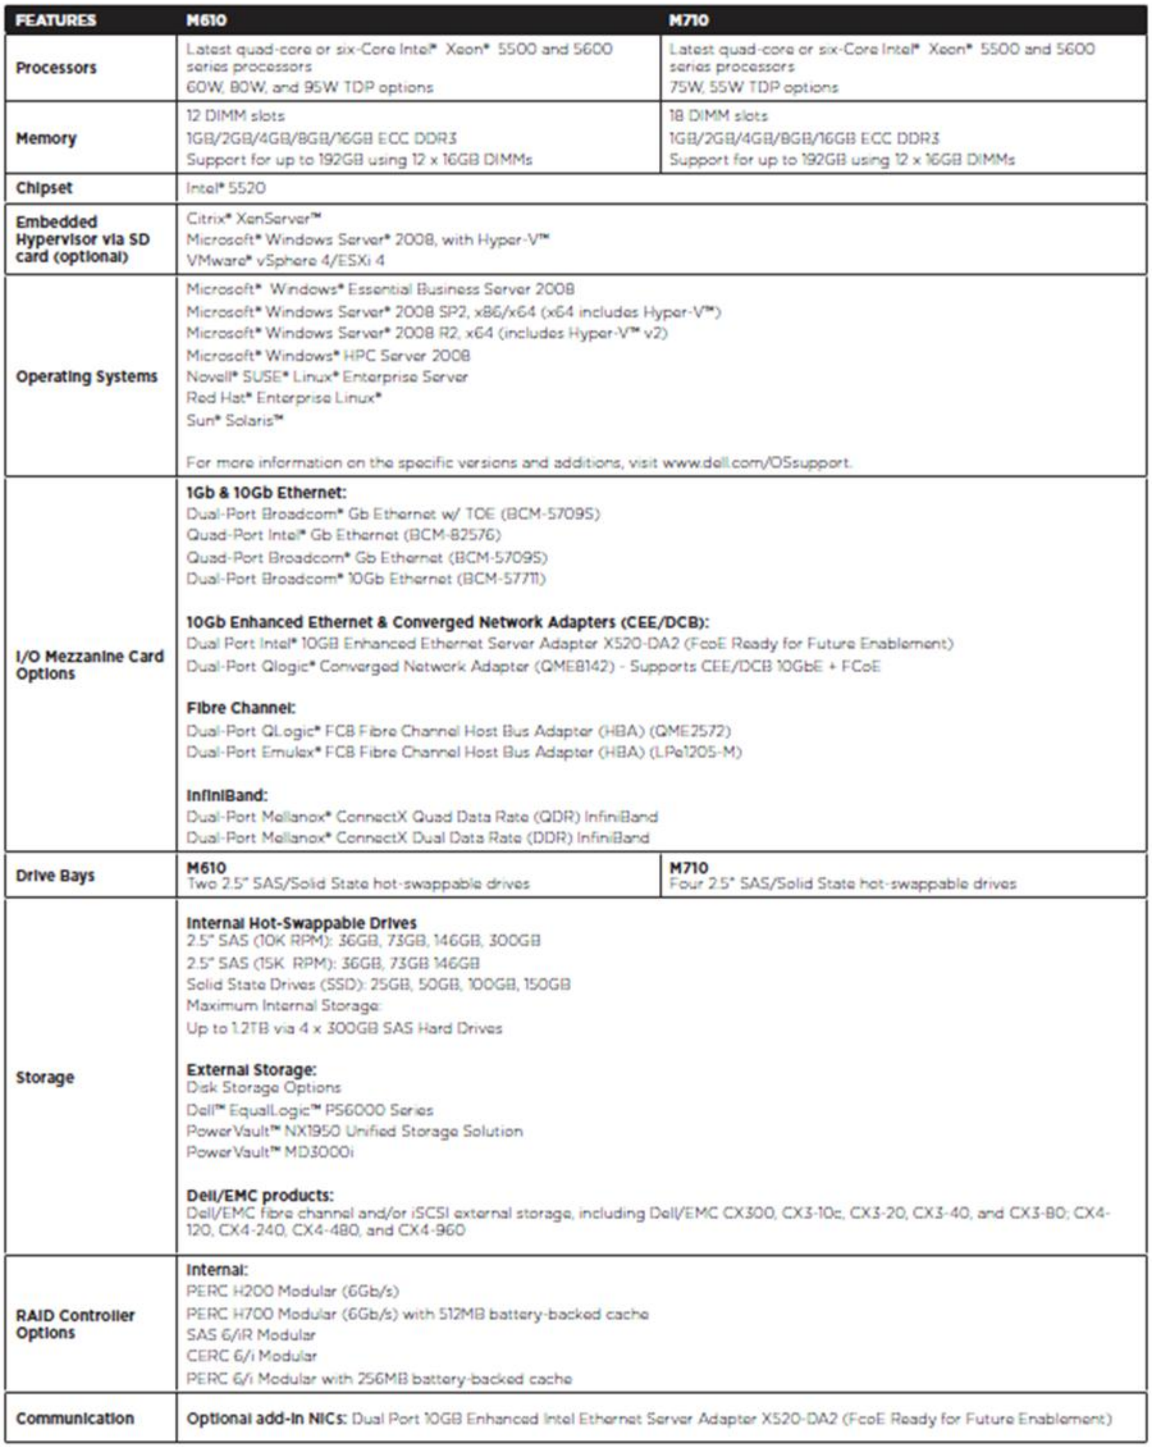

FEATURES	M610	M710
Processors	Latest quad-core or six-Core Intel® Xeon® 5500 and 5600 series processors 60W, 80W, and 95W TDP options	Latest quad-core or six-Core Intel® Xeon® 5500 and 5600 series processors 75W, 55W TDP options
Memory	12 DIMM slots 1GB/2GB/4GB/8GB/16GB ECC DDR3 Support for up to 192GB using 12 x 16GB DIMMs	18 DIMM slots 1GB/2GB/4GB/8GB/16GB ECC DDR3 Support for up to 192GB using 12 x 16GB DIMMs
Chipset	Intel® 5520	
Embedded Hypervisor via SD card (optional)	Citrix® XenServer™ Microsoft® Windows Server® 2008, with Hyper-V™ VMware® vSphere 4/ESXi 4	
Operating Systems	Microsoft® Windows® Essential Business Server 2008 Microsoft® Windows Server® 2008 SP2, x86/x64 (x64 includes Hyper-V™) Microsoft® Windows Server® 2008 R2, x64 (includes Hyper-V™ v2) Microsoft® Windows® HPC Server 2008 Novell® SUSE® Linux® Enterprise Server Red Hat® Enterprise Linux® Sun® Solaris™ For more information on the specific versions and additions, visit www.dell.com/OSsupport.	
I/O Mezzanine Card Options	**1Gb & 10Gb Ethernet:** Dual-Port Broadcom® Gb Ethernet w/ TOE (BCM-57095) Quad-Port Intel® Gb Ethernet (BCM-82576) Quad-Port Broadcom® Gb Ethernet (BCM-57095) Dual-Port Broadcom® 10Gb Ethernet (BCM-57711) **10Gb Enhanced Ethernet & Converged Network Adapters (CEE/DCB):** Dual Port Intel® 10GB Enhanced Ethernet Server Adapter X520-DA2 (FcoE Ready for Future Enablement) Dual-Port Qlogic® Converged Network Adapter (QME8142) - Supports CEE/DCB 10GbE + FCoE **Fibre Channel:** Dual-Port QLogic® FC8 Fibre Channel Host Bus Adapter (HBA) (QME2572) Dual-Port Emulex® FC8 Fibre Channel Host Bus Adapter (HBA) (LPe1205-M) **InfiniBand:** Dual-Port Mellanox® ConnectX Quad Data Rate (QDR) InfiniBand Dual-Port Mellanox® ConnectX Dual Data Rate (DDR) InfiniBand	
Drive Bays	**M610** Two 2.5" SAS/Solid State hot-swappable drives	**M710** Four 2.5" SAS/Solid State hot-swappable drives
Storage	**Internal Hot-Swappable Drives** 2.5" SAS (10K RPM): 36GB, 73GB, 146GB, 300GB 2.5" SAS (15K RPM): 36GB, 73GB 146GB Solid State Drives (SSD): 25GB, 50GB, 100GB, 150GB Maximum Internal Storage: Up to 1.2TB via 4 x 300GB SAS Hard Drives **External Storage:** Disk Storage Options Dell™ EqualLogic™ PS6000 Series PowerVault™ NX1950 Unified Storage Solution PowerVault™ MD3000i **Dell/EMC products:** Dell/EMC fibre channel and/or iSCSI external storage, including Dell/EMC CX300, CX3-10c, CX3-20, CX3-40, and CX3-80; CX4-120, CX4-240, CX4-480, and CX4-960	
RAID Controller Options	**Internal:** PERC H200 Modular (6Gb/s) PERC H700 Modular (6Gb/s) with 512MB battery-backed cache SAS 6/iR Modular CERC 6/i Modular PERC 6/i Modular with 256MB battery-backed cache	
Communication	**Optional add-in NICs:** Dual Port 10GB Enhanced Intel Ethernet Server Adapter X520-DA2 (FcoE Ready for Future Enablement)	

Figure 11. Dell PowerEdge M610 Specs (From Dell, 2011)

Figure 12. Samsung Syncmaster NC240 and Wyse P20

5. VMware ESX Software

The virtualization foundation for the private cloud will be VMware's ESX software. It installs directly on the server without a host operating system such as Windows Server 2008. By installing directly on the server without a host operating system, the use of the server's resources are maximized and strictly used for virtualization (Marshall, Beaver, & McCarty, 2009). Server virtualization is useful to perform functions such as server consolidation and legacy application support. By consolidating servers into one unit, it allows for efficient use of the servers resources, unlike having separate servers with specific functions which only use 12% of the servers resources (Marshall et al., 2009).

With legacy applications that run on older systems without hardware and software support, server virtualization is a viable option to upgrade the server's hardware and still maintain the legacy operating system and application. This will help save enterprises time and money without having to relearn a new application, lose data from information migration, and create hidden costs due to possible system downtime (Marshall et al., 2009).

There are limitations to VMware's ESX software such as graphic intensive applications. When these conditions exist, server virtualization is not recommended. Graphic intensive applications require large performance requirements that will utilize a majority of a server's resources for that one application. For high graphic intensive applications, it is recommended that a stand-alone computer be utilized (Marshall et al., 2009).

The use of specialized peripheral cards or creating a virtual machine with an AMD processor on an Intel-based platform system is not currently possible. The emulation of specialized devices in a virtual environment is not available unless it is physically there. Additionally, the creation of a virtual server inside of another virtual environment is also not recommended. By creating virtual environments on top of each other, double slicing of a servers resources and the virtual environment will become unusable (Marshall et al., 2009).

To develop a virtual environment, the Virtual Center is the command console used in configuring all aspects of virtualization. With Virtual Center, an administrator can manage up to 200 hosts and 2000 virtual machines (Marshall et al., 2009). The following five components make up a Virtual Center:

- Virtual Sphere Client (vSphere Client) is the center control point for all configuration and management of the VMware environment (Marshall et al., 2009).
- Virtual Center Database is used for all storage regarding the physical server, resource pools, and managed virtual machines (Marshall et al., 2009).
- Virtual Infrastructure Client is the administrative client that connects Virtual Center Management Server directly (Marshall et al., 2009).
- Virtual Center Agent is the ESX server agent that connects to Virtual Center on VMware ESX server (Marshall et al., 2009).
- Virtual Infrastructure Web Access allows virtual machine management of consoles without the use of a client (Marshall et al., 2009).

The Virtual Center has a set of core feature that enable the administrator to rapidly develop and deploy a virtual infrastructure. With VMware, the cloning of virtual machines and creating templates is possible. This allows the administrator to save virtual machines in storage and configure them as needed for updating and patching purposes. Once a template is developed, it can be cloned and deployed to other servers. This can save hours of configuration if done on a normal client (Marshall et al., 2009).

Another feature of VMware is the Live Migration (vMotion) feature. This has the ability to migrate a running virtual machine from one physical host to another with no impact by the end user (Marshall et al., 2009). This feature enabled with the Distributed Resource Scheduler (DRS) and Distributed Power Management (DPM) features ensures that physical hosts are being utilized efficiently. DRS will manage the physical host and ensure that the load utilization is even across all servers. If the load is light, the Live Migration feature will shift virtual machines from one server to the other and the DPM will power down the under-utilized servers. Additionally, if the virtual machine load increases, the DPM powers on additional host servers to ensure sufficient resources are available (Marshall et al., 2009).

The last core feature of VMware is the High Availability feature that protects virtual machines against hardware or software failures. In the event of a failure, the DRS will determine which host to start. It will then restart the virtual machine automatically on the new host. Therefore, if a host server goes down in the middle of the night, the High Availability feature automatically migrates the remaining virtual machines to another host and allows the administrator to correct the hardware issue at their convenience (Marhsall et al., 2009).

F. BUILDING A PRIVATE CLOUD

Before building a private cloud, the firmware and BIOS for each server was updated to the latest version to ensure hardware and software compatibility. Once BIOS upgrades were completed, a host name and static IP address were assigned to each individual server in accordance with Table 2. The AEGIS host name has no affiliation with the AEGIS weapon system and is simply a server naming convention.

Table 2. Private Cloud IP Address Table

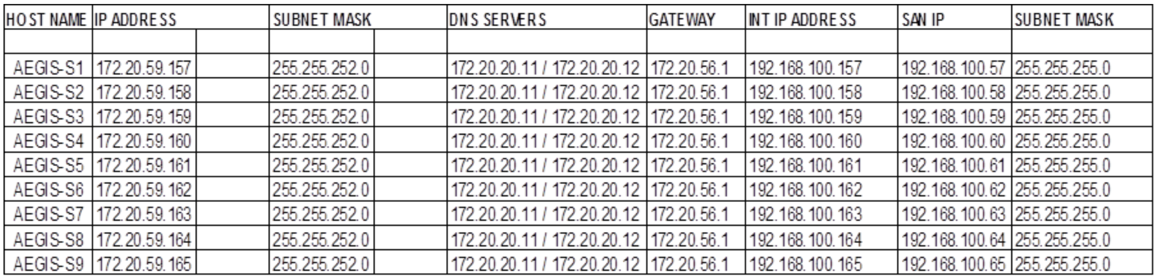

HOST NAME	IP ADDRESS		SUBNET MASK		DNS SERVERS	GATEWAY	INT IP ADDRESS	SAN IP	SUBNET MASK
AEGIS-S1	172.20.59.157		255.255.252.0		172.20.20.11 / 172.20.20.12	172.20.56.1	192.168.100.157	192.168.100.57	255.255.255.0
AEGIS-S2	172.20.59.158		255.255.252.0		172.20.20.11 / 172.20.20.12	172.20.56.1	192.168.100.158	192.168.100.58	255.255.255.0
AEGIS-S3	172.20.59.159		255.255.252.0		172.20.20.11 / 172.20.20.12	172.20.56.1	192.168.100.159	192.168.100.59	255.255.255.0
AEGIS-S4	172.20.59.160		255.255.252.0		172.20.20.11 / 172.20.20.12	172.20.56.1	192.168.100.160	192.168.100.60	255.255.255.0
AEGIS-S5	172.20.59.161		255.255.252.0		172.20.20.11 / 172.20.20.12	172.20.56.1	192.168.100.161	192.168.100.61	255.255.255.0
AEGIS-S6	172.20.59.162		255.255.252.0		172.20.20.11 / 172.20.20.12	172.20.56.1	192.168.100.162	192.168.100.62	255.255.255.0
AEGIS-S7	172.20.59.163		255.255.252.0		172.20.20.11 / 172.20.20.12	172.20.56.1	192.168.100.163	192.168.100.63	255.255.255.0
AEGIS-S8	172.20.59.164		255.255.252.0		172.20.20.11 / 172.20.20.12	172.20.56.1	192.168.100.164	192.168.100.64	255.255.255.0
AEGIS-S9	172.20.59.165		255.255.252.0		172.20.20.11 / 172.20.20.12	172.20.56.1	192.168.100.165	192.168.100.65	255.255.255.0

1. Installation of Software

The installation of Windows Server 2008 was accomplished on the server labeled AEGIS S1. Windows Server was installed on this server so the use of Dell Server Manager software can be utilized to monitor real time energy consumption of each individual server enclosure and chassis. By using the Dell software, this is a simple and effective way to monitor power usage and monitor how the VMware software manages the server chassis.

The installation of VMware ESXi V4.1 was installed directly on servers AEGIS S2 through AEGIS S5 without a host operating system. These servers will be the hosts for which the virtual machines will operate. Lastly, VMware vSphere Client was installed on AEGIS S1 to build and manage the virtual environment.

2. Creating a Virtual Network

In order to utilize vMotion capabilities, a VMkernel was created and configured for AEGIS S2 in the vSphere Client interface. A VMkernel controls and manages the physical resources on the server including the processor, memory, and networking interfaces (Marshall et al., 2009). Once the VMkernel was established, a virtual network distributed switch called “dvswitch0” was created and will be used for the management network. Lastly, a virtual Network Interface Controller (vNIC) was created, assigned an IP address, and virtually connected the dvswitch0. All server management functions such as DPM and DRS will be coordinated on this network.

To test the DPM feature and the network, AEGIS S2 was added to a “cluster” and was manually powered off and on utilizing the vSphere Client interface to ensure that the management network and DPM feature is functioning correctly. The previous steps were repeated for AEGIS S3 through AEGIS S5.

To prevent network bottlenecking and maximize the networking capabilities, a separate virtual distributed switch was created to attach to the SAN, handle vMotion functions, and connect to the Internet on the ERN domain. Figure 13 shows how AEGIS S2 through S5 is connected utilizing vNICs and virtual distributed switches.

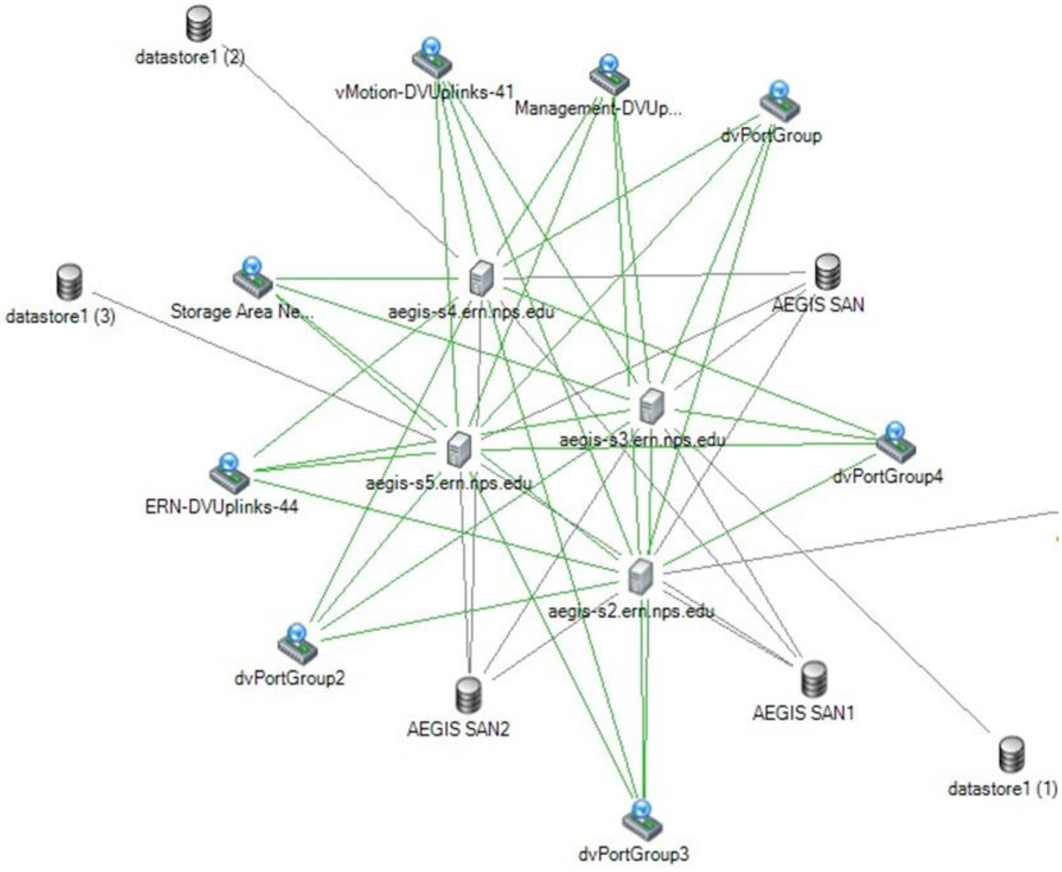

Figure 13. Virtual Network Configuration

3. Adding Servers and Virtual Desktops to the Virtual Network

In order for the virtual network to operate, a virtual server was created having two processors, 6 gigabytes of memory, two 60 GB hard disk drives, and Windows Server 2008 operating system installed. Once all software patches, virus protection, and user configurations set, the server was set as a template and cloned to create three new virtual servers.

The first server named AEGIS-Virtual Desktop Infrastructure (VDI), will be the VMware View Manager server. This server will allow administrators to configure the View Connection Server, deploy and manage desktops, and control which users have access to the virtual desktops. The second server named AEGIS-SQL, will have Microsoft SQL installed and will create databases and manage the event logs for the View Manager server and View Composer Server (VMware, 2010). The last server, AEGIS-SQL2 will be another database server with Microsoft SQL installed for future use.

Before creating virtual desktops, VMware View Composer was installed on the AEGIS-S1 server. View Composer can create a pool of linked clones from a specific template. By having each cloned desktop share a base image, this reduces overall data storage and allows easy deployment of updates and patches without affecting users since only the base image is getting the upgrade (VMware, 2010).

With all the necessary software configured; a virtual desktop was created with Windows 7 Professional operating system. Once the required software and patches were installed on the virtual desktop, the desktop was set as a template to be cloned. Due to limits on licenses, View Composer was utilized and two pools of 25 virtual machines were created for a total of 50 virtual desktops.

IV. ANALYSIS

A. INTEL XEON HISTORICAL DATA ANALYSIS

Over the past 17 years, virtual technology is possible because of server processors such as the Intel Xeon series. The Xeon processor evolved from being a single core processor to having 10 cores in a processor. The key factor in doing virtualization is processing power and being able to handle multiple tasks simultaneously.

In doing the analysis, the processor clock speed was compared against the processors rated Floating Operating Points Per Second (FLOPS) for Xeon processor family. The Intel Xeon Pro 200 was the first server processor released in 1995 followed by the Pentium II Xeon 450 in 1998 (Intel, 2011a). Due to the number or Xeon processors, random processors were selected starting with Xeon 3000 series, which was released in 2006 to the latest Xeon E7–8800 series released in 2011 (Intel, 2011b, d, f, h). Figure 14 shows the relationship between processor speed and FLOPS over a 17-year period. In the first three Xeon series, processor speed increased from 200 MHz to 800 MHz over a six-year period increasing by a magnitude of four. For the processors FLOPS, they increased from .2 GFLOPS to 1.6 GFLOPS for a magnitude of eight. The increase of processor speed has a direct correlation to the increase of GFLOPS.

Randomly selected Xeon processors from 2006 to present have a peak processor clock speed of 3.2 GHz and a mean of 2.4 GHz (Intel, 2011b, d, f, h). However, their GFLOPS has increased from 14.9 GFLOPS to a peak of 96 GFLOPS. Due to the increase of GFLOPS, there is no significant correlation between GFLOPS and processor clock speed since the processor clock speed is relatively steady.

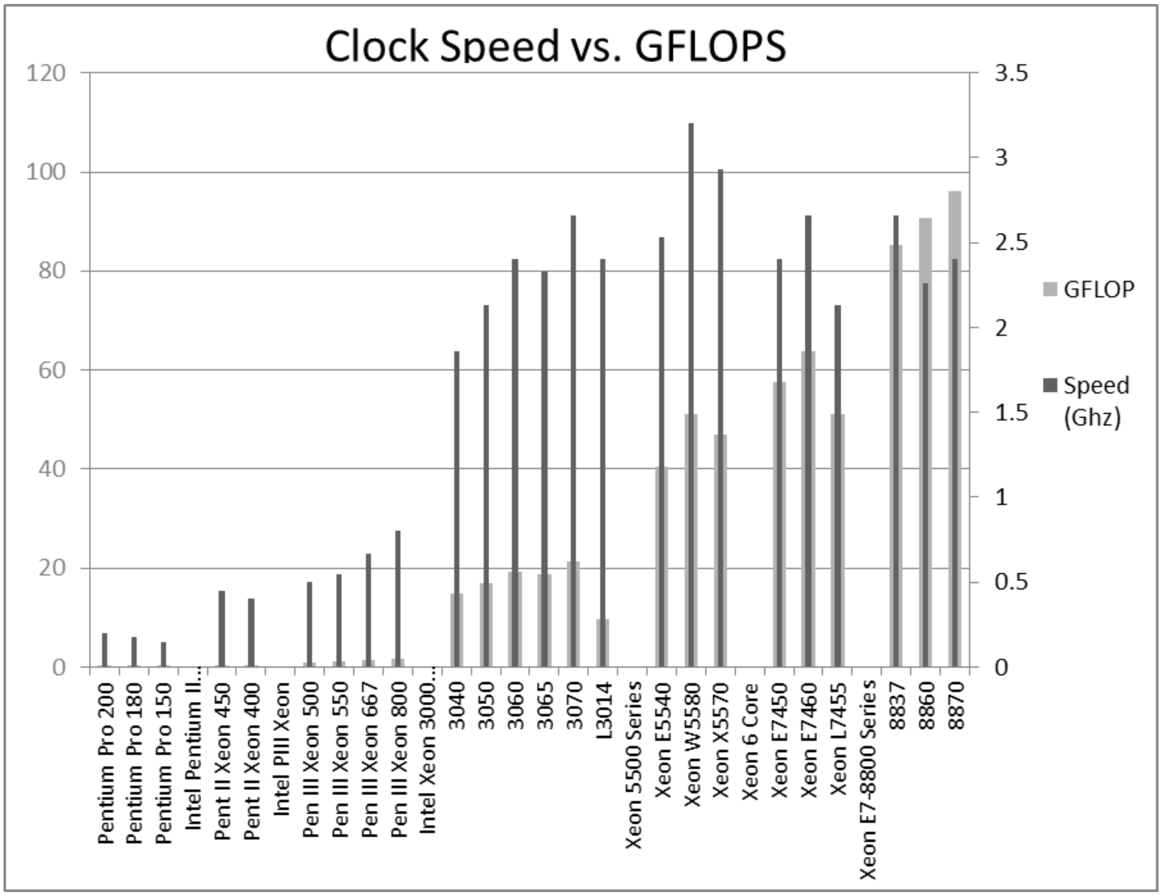

Figure 14. Intel Xeon Processor Clock Speed vs. GFLOPS (After Intel, 2011a, c, e, g, i)

To account for the increased GFLOPS with a steady processor clock speed, the GFLOPS was compared against the number of processor cores in Figure 15. The Xeon 3000 series are the first multicore processor in this sample with a median of 8.9 GFLOPS. With the multicore processors, increasing clock speed has significant variance in increasing their GFLOPS performance. As processors advanced up to 10 cores per processor, the median increased to 9.6 GFLOPS per processor core and vary according to the processor clock speed for the 2006 to present processors. With the combination of additional cores and clock speed, a processor's performance will increase and handle multiple tasks in a virtual environment (Intel, 2011c, e, g, i).

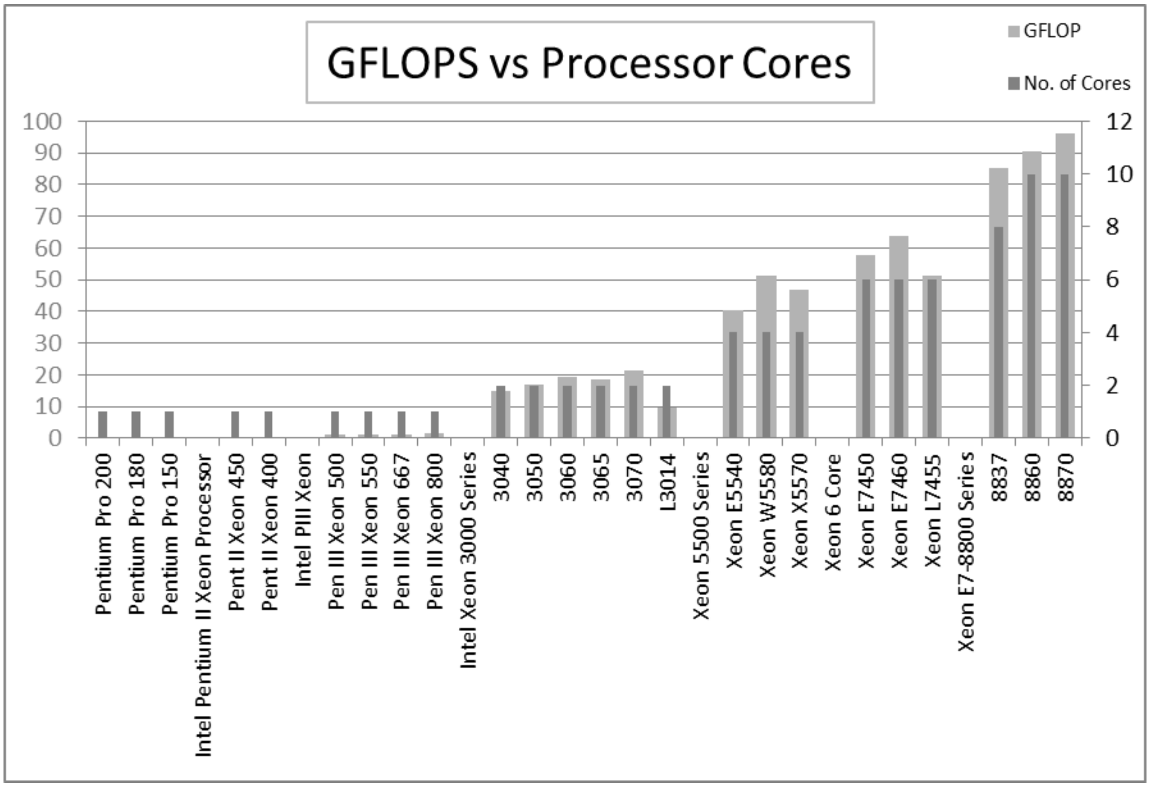

Figure 15. Intel Xeon GFLOPS vs. Processor Cores (After Intel, 2011a, b, d, f, h)

Increasing processor capabilities creates increasing electrical power consumption that adds to operating costs for regular operation and cooling. Instead of measuring actual power consumption of a processor, the heat dissipated in watts will be compared. The Thermal Design Power (TDP) is the maximum amount of heat that a thermal cooling solution must be able to dissipate from the processor so it will operate under normal operating conditions (Intel, 2011j). As a processor operates under a continuous load, it will generate heat that must be dissipated. If the heat is not dissipated, the processor will overheat and cause failure. In Figure 16, the TDP ratio to GFLOPS and processor clock speed was compared. The early Xeon processors generated 38 Watts of heat and had a processing capability of .2 GFLOPS; therefore, the graph is skewed right. For the newer processors, they dissipate between 65 and 130 Watts of heat depending on the number of cores and processing power. However, their Watt/GFLOP ratio varies from 1.5 to 4. Therefore, the 3000 series and newer processors are energy efficient by emitting less heat per GFLOP increase (Intel, 2011a, b, c, d, e, f, g, h, i).

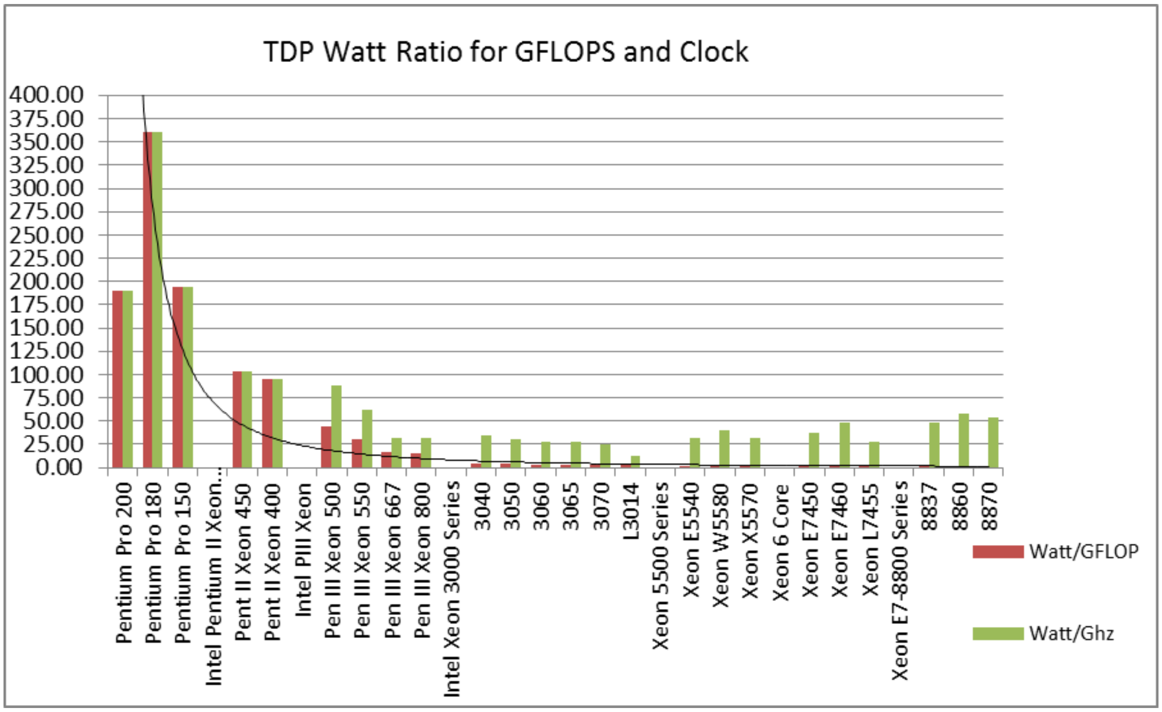

Figure 16. Intel Xeon TDP Ratio per GFLOP and Ghz (After Intel, 2011a, c, e, g, i)

B. TOTAL COST OF OWNERSHIP ANALYSIS

1. Assumptions

With the U.S. Navy developing a private cloud for naval vessels in the Consolidated Afloat Network and Enterprise Services (CANES) program, a Total Cost of Ownership analysis will be done for an Arleigh Burke Class Destroyer that uses Ships Service Gas Turbine Generators (SSGTG), and an Amphibious Dock Landing ship that uses Ships Service Diesel Generators (SSDG) for electrical power distribution. The analysis will determine how much costs could change by transitioning from the current client server architecture to a private cloud network once the acquisition phase is complete. Due to the unknown costs of the acquisition phase of the CANES program, this cycle will be disregarded. The following is the assumptions for the current network infrastructure on ships that utilize SSGTGs and SSDGs:

- There are 175 computers on the UNCLASSIFIED network and 50 computers on the CLASSIFIED network
- Each computer on both the UNCLASSIFIED AND CLASSIFIED networks is a Dell Optiplex GX620 with an Intel Pentium 4 processor with a rated cock speed of 3.2 Ghz, and has 2 gigabytes of random access memory.
- There are four servers on a ship that use an average of 2K watts each.

For the private cloud model at Naval Postgraduate School (NPS), the model has limitations due to licensing constraints of the VMware software and will only be able to license 50 virtual machine workstations. To replicate a proposed model in a shipboard environment, there will be 200 virtual machine workstations due to the Cross Domain Solutions (CDS) of CANES that combines both UNCLASSIFIED and CLASSIFIED networks into one session. Therefore, 25 workstations will be eliminated from the ship. The last assumption will be that the four physical shipboard servers will be combined into the virtual environment and will be simulated with four virtual servers.

According to the U.S. Energy Information Administration (2011), the average price per kilowatt-hour in Virginia and Florida is $.10. However, for ships in California, the average commercial price per kilowatt-hour is $.14 (U.S. EIA, 2011). Therefore, an average of $.12 will be used for calculating shore power costs.

For every change of 100 kilowatt-hours in electrical load, the fuel consumption for both SSDGs and SSGTGs will change eight gallons per hour based on the fuel curves in Figures 6 and 7 from Chapter III. To determine a baseline for how much fuel a ship burns in port and underway, the ships class quarterly average from the Navy Energy Usage Reporting System (NEURS) will be used. Table 3 shows the quarterly average burn rate for a ship in each class during 2011 in barrels. A barrel of fuel equals 42 gallons (NIST, 2012).

Table 3. 2011 NEURS Quarterly Fuel Usage (After, i-ENCON, 2012)

	UNDERWAY			NOT UNDERWAY		
	AVG HRs/QTR	AVG BBLs/QTR	BBLs/HR	AVG HRs/QTR	AVG BBLs/QTR	BBLs/HR
CG 47 Class Average	734	21,214	28.92	241	1,739	7.2
DDG 51 Class Average	779	18,673	23.96	239	1,525	6.38
FFG 7 Class Average	737	8,095	10.98	216	467	2.16
LSD 41 Class Average	618	7,981	12.92	243	933	3.84
LSD 49 Class Average	947	11,790	12.45	231	1,114	4.81

2. Thin Client vs. Thick Client Analysis

To measure the implementation phase of the private cloud which includes hardware, software, and installation costs, the following items will be compared: Samsung Syncmaster NC420, Wyse P20, and a Dell T1600 Workstation to replace the current thick clients in the fleet. The Dell T1600 workstation is a basic workstation that comes with Windows 7 Professional already installed. Table 4 compares the Manufactured Suggested Retail Price (MSRP) of each item and the overall cost for a ship.

Table 4. Hardware Implementation Cycle Costs

Item	MSRP	No. of Machines	Total Ship	Notes
Dell T-1600 (w/o monitor)	$760	225	$171,000	To upgrade current network
Samsung Syncmaster NC240	$550	200	$110,000	
Wyse P20 (w/o monitor)	$450	200	$90,000	

The Samsung Syncmaster and Wyse P20 is $61K and $81K dollars cheaper respectively compared to the workstation upgrade. To save on replacement costs for the legacy system, the current monitors will be retained if they are flat screen liquid crystal display (LCD) monitors. If the Syncmaster NC240 was used, existing peripherals such as the keyboard and mouse can be retained, but electronic waste recycling fees could increase the replacement phase costs for legacy systems, since states such as California charge $8 per monitor for recycling fees (CA Board of Equalization, 2012).

For the maintenance phase of the thin clients, there is no required maintenance or software patches required unlike thick clients. Thin clients have no operating systems that require updating since all the computing is done on the cloud server. If upgrades are needed for the operating system on the virtual machines, the updates can be done to the template, tested, and then cloned to the virtual machine pool without affecting any users. Unlike thin clients, thick clients need to be individually upgraded with patches that can create indirect costs due to users down time.

To calculate the difference in operating costs for the operating phase of the TCO analysis while on shore power, the input power of each device was measured in watts for the following power states in Table 5: off, standby, idle, and normal operation.

Table 5. Energy Cost and Usage Comparison of Thin Clients vs. Thick Clients on Shore Power

	Power State	No. of Machines	Load per Kilowatt-hr	Cost per Kilowatt-hr @ $.12
	Off			
Syncmaster NC240	1.2	200	0.240	$0.03
P20 w/monitor	1.1	200	0.220	$0.03
Dell Optiplex GX620 w/monitor	3.1	225	0.698	$0.08
	Standby			
Syncmaster NC240	20.5	200	4.100	$0.49
P20 w/monitor	15.3	200	3.060	$0.37
Dell Optiplex GX620 w/monitor	19	225	4.275	$0.51
	Idle			
Syncmaster NC240	61.7	200	12.340	$1.48
P20 w/monitor	46.5	200	9.300	$1.12
Dell Optiplex GX620 w/monitor	119.5	225	26.888	$3.23
	Normal Operation			
Syncmaster NC240	63	200	12.600	$1.51
P20 w/monitor	46.5	200	9.300	$1.12
Dell Optiplex GX620 w/monitor	180.5	225	40.613	$4.87

By reducing the number of thick-client workstations and shifting away from the current client server architecture, thin clients on a private cloud reduce energy costs. The hourly energy usage costs can be reduced by $3.76 during peak working hours in port if using the Wyse P20 thin client. If thick clients are not shut down and stay idled for the evening due to user neglect, a $2.11 hourly savings still can be achieved since thin clients use less energy when at idle.

To calculate the difference when the ship is underway on ship's power, only the idle and normal operation power state will be compared in Table 6. For every change of 12.5 kilowatts, one gallon of fuel will be saved using either a SSDG or SSGTG.

Table 6. Thick Client vs. Thin Client on Ship's Power

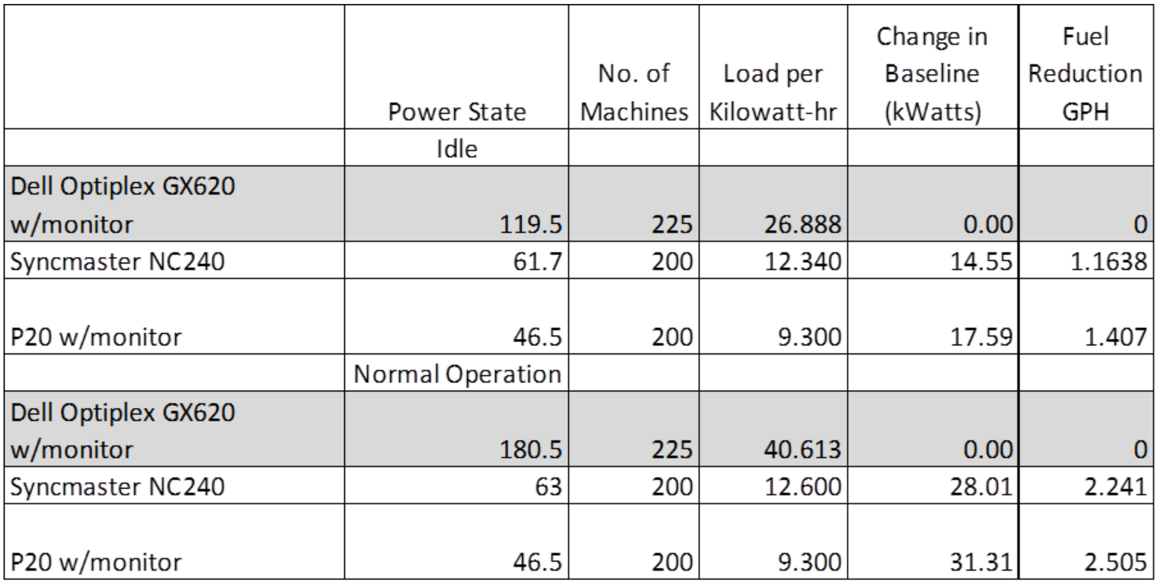

	Power State	No. of Machines	Load per Kilowatt-hr	Change in Baseline (kWatts)	Fuel Reduction GPH
	Idle				
Dell Optiplex GX620 w/monitor	119.5	225	26.888	0.00	0
Syncmaster NC240	61.7	200	12.340	14.55	1.1638
P20 w/monitor	46.5	200	9.300	17.59	1.407
	Normal Operation				
Dell Optiplex GX620 w/monitor	180.5	225	40.613	0.00	0
Syncmaster NC240	63	200	12.600	28.01	2.241
P20 w/monitor	46.5	200	9.300	31.31	2.505

During an operating quarter while underway, if the workstation was used 12 hours a day and idle for the other 12 hours, an estimated 1,989 gallons of fuel will be saved each quarter for an average destroyer underway 779 hours. When the same ship is in port and not on shore power, an estimated 466 gallons can be saved for a total quarterly savings of 2,455 gallons of fuel.

3. Private Cloud Elasticity and Energy Analysis

To test Jennings's (2010) theory that energy is transferred from the customer to the cloud or if cloud computing is efficient and reduces energy consumption, the private cloud at Naval Postgraduate School (NPS) will be tested. To measure the energy consumption of the private cloud, the real time energy consumption of the chassis will be measured by using the Dell M1000e chassis monitoring software. The software can also measure the real time and peak energy consumption for each physical blade server over a specified period of time.

Before starting the test and enabling the Data Power Management (DPM) function of VMware, a baseline energy reading was taken. In Table 7, the component AEGIS-S1 is only used for managing the virtual machines and has no virtualization capability. AEGIS-S2 through S5 has the VMware ESXi software that allows virtual hosts to run on the physical servers. The remaining 233 watts is used by the chassis for cooling and running peripheral devices.

Once the DPM software was enabled with the three virtual management servers and 12 virtual desktops at idle, another power reading was taken. With the DPM enabled, AEGIS-S2 and AEGIS-S4 were automatically placed into Standby by the VMware software and reduced energy consumption by 230 watts. Even though there were 12 virtual desktops online and in Standby, the energy consumption did not increase for AEGIS-S3 and AEGIS-S5. Since the virtual machines are online and not using any of the hosts' resources, there is no change in power consumption.

Table 7. AEGIS VM Server's Baseline Energy Usage

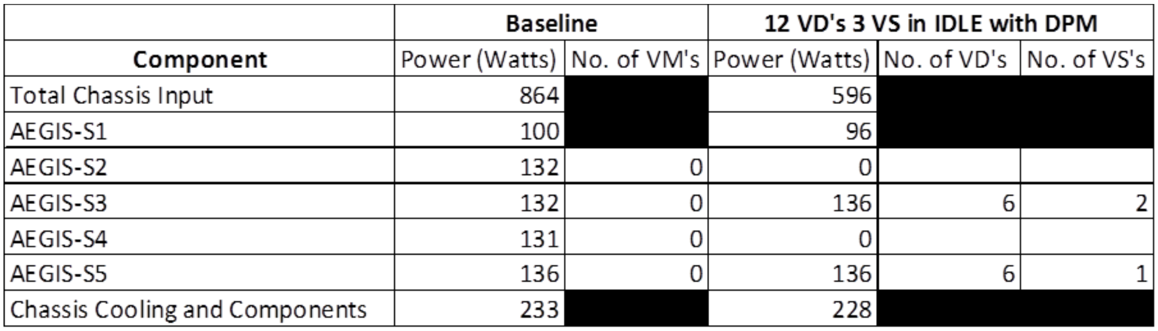

	Baseline		**12 VD's 3 VS in IDLE with DPM**		
Component	Power (Watts)	No. of VM's	Power (Watts)	No. of VD's	No. of VS's
Total Chassis Input	864		596		
AEGIS-S1	100		96		
AEGIS-S2	132	0	0		
AEGIS-S3	132	0	136	6	2
AEGIS-S4	131	0	0		
AEGIS-S5	136	0	136	6	1
Chassis Cooling and Components	233		228		

To test how well the Data Resource Scheduler (DRS) and DPM work together, virtual desktops were gradually added to the cluster and placed in normal operation. Due to each physical host only having 24 gigabytes of memory, when 12 virtual desktops are running and operating, warning indications of low memory resources of a physical host is displayed. The DRS will automatically notify the DPM to start another physical host server, and it will be online to start accepting virtual desktops within 10 minutes. The long startup time for a new host server is because the server is starting from a cold boot because there is no standby mode of operation for the Dell M610 server blade. Additionally, it takes five minutes for the physical host server to be configured for High Availability, which protects the virtual machines from hardware and software failures (Marshall et al., 2009). Once the new physical host server is online, the vMotion feature will transfer virtual desktops to the new host in order to balance the memory resources and energy. Each physical host can manage approximately 12 virtual desktops before warning indications of high memory usage is displayed to the administrator.

While adding virtual desktops to the host of clusters, the power consumption was measured throughout the process. When each virtual desktop was online and working, power consumption for each host server increased approximately 5.5 watts from the baseline as displayed in Figure 17. Once desktops 40 through 48 were added, the curve started to flatten since all four physical server hosts were online. The positive slope from 20 to 40 is attributed to the baseline of 135 watts when each physical host was started and placed online.

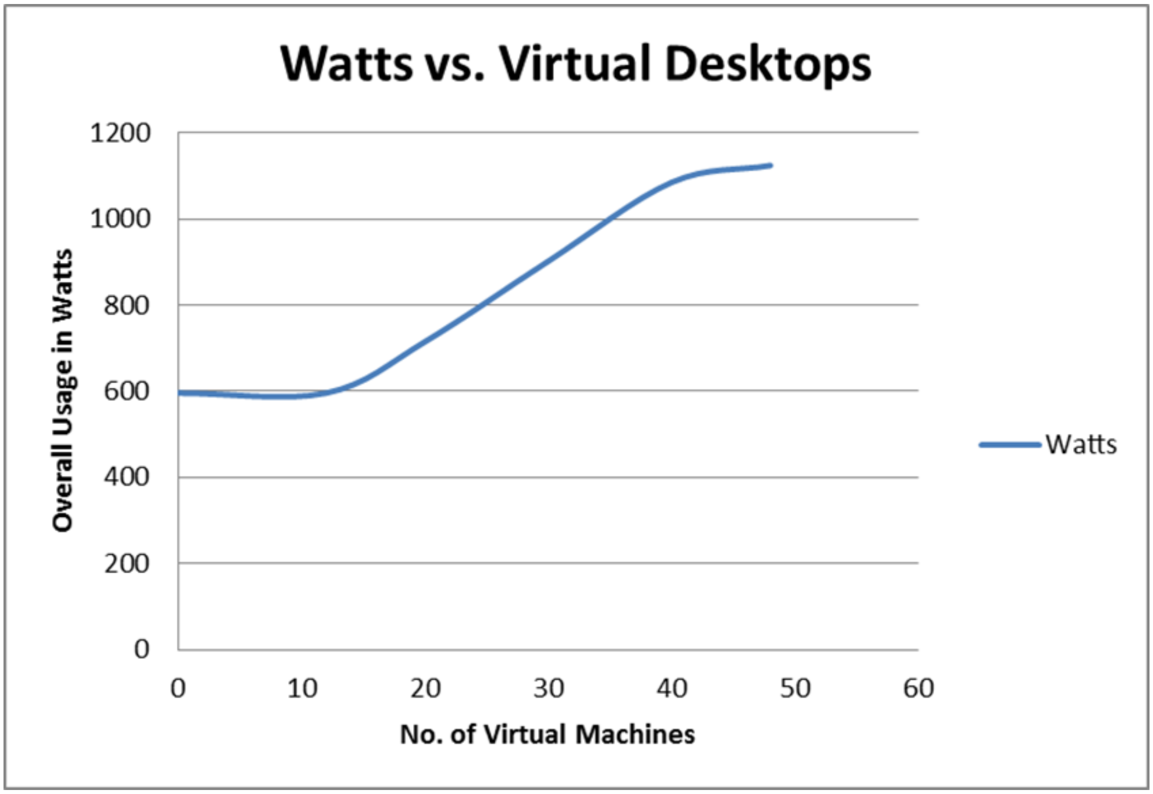

Figure 17. Overall AEGIS VM Server Energy Consumption

With the current configuration of each Dell M610 blade server having only 24 gigabytes of memory, 200 virtual desktops would consume 4.5 kilowatts of electricity. To reduce electrical consumption even more, additional memory can be added to each physical host server to increase the number of virtual desktops in can host. By consolidating servers and workstations into a virtual computing environment, overall server energy consumption was reduced by 3.5 kilowatts compared to the baseline of 8 kilowatts of a ship's current configuration.

4. Overall Total Cost of Ownership Savings for a Five-Year Period

To estimate the total savings in energy consumption over a five-year period for an Arleigh Burke destroyer and an Amphibious Dock Landing ship, the results from thin versus thick and private cloud analysis will be used. Due to the unknown prices of the Dell M1000e server chassis, only the price of 16 Dell Poweredge M610 server blades with an MSRP of $4,100 each will be accounted for in the analysis. The price of

$3.81 per gallon for F-76 diesel fuel will be the standard price according to Defense Energy Support Center (2012). Table 8 displays the potential savings for a destroyer if upgraded with the Wyse P20 thin client and private cloud. The total operational savings for an upgraded ship would be approximately $44,000 a year with no discount rate. The new Wyse P20 thin clients and Dell M610 server blades would have a payback period of 42 months ($155,600/$44,000) because of the reduced energy usage.

Table 8. Quarterly and Annual Savings for a Destroyer

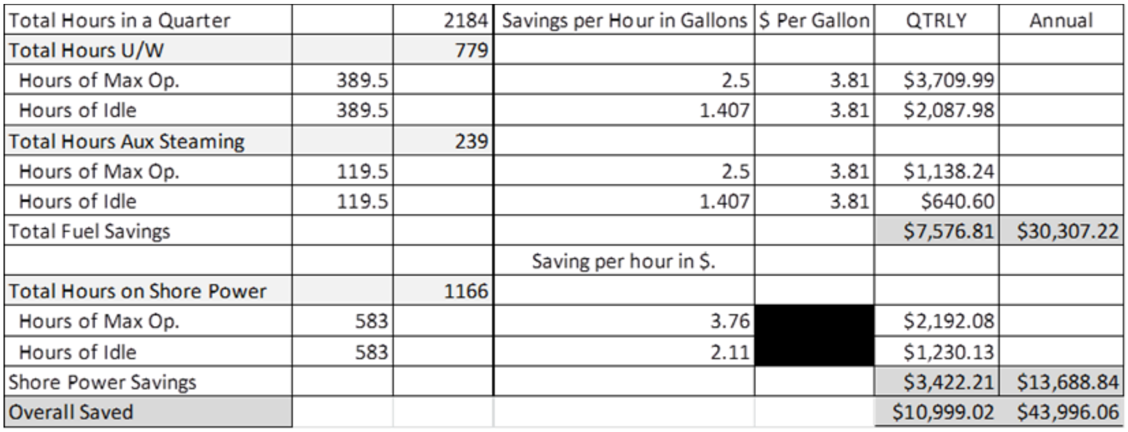

Total Hours in a Quarter		2184	Savings per Hour in Gallons	$ Per Gallon	QTRLY	Annual
Total Hours U/W		779				
Hours of Max Op.	389.5		2.5	3.81	$3,709.99	
Hours of Idle	389.5		1.407	3.81	$2,087.98	
Total Hours Aux Steaming		239				
Hours of Max Op.	119.5		2.5	3.81	$1,138.24	
Hours of Idle	119.5		1.407	3.81	$640.60	
Total Fuel Savings					$7,576.81	$30,307.22
			Saving per hour in $.			
Total Hours on Shore Power		1166				
Hours of Max Op.	583		3.76		$2,192.08	
Hours of Idle	583		2.11		$1,230.13	
Shore Power Savings					$3,422.21	$13,688.84
Overall Saved					$10,999.02	$43,996.06

The overall Net Present Value (NPV) of migrating to a private cloud over a five-year period is displayed in Table 9. The savings in energy costs offset the new hardware thin clients and blade server costs with a NPV of $64,380 with no discount rate. With a 5% discount rate used, the NPV drops to $34,880. Due to the unknown costs of the Dell® M1000e™ chassis server and associated software, the hardware costs of a new server chassis should be recovered within a five-year period. The largest driver for the NPV is price of fuel. If the price falls below $2.80 a gallon, the NPV is less than zero with a 5% discount rate for an Arleigh Burke Destroyer and Amphibious Dock Landing Ship. Due to fewer operating hours, the price of fuel must stay above $2.86 for a Ticonderoga Class Cruiser to keep a positive NPV. As energy prices stay high, a private cloud can recover its hardware costs and save $4M in operating costs over a five-year period for the three classes of ships listed in Table 9.

Since the existing infrastructure of routers, switches, and cabling is already in place; there is no need to upgrade these components since all the data transfer is internal to the server chassis and storage area network. By reducing the amount of network traffic since only pixel data is transferred over the network; current components can be reused and reduce the implementation phase costs.

Table 9. Net Present Value and Internal Rate of Return Over a Five-Year Period

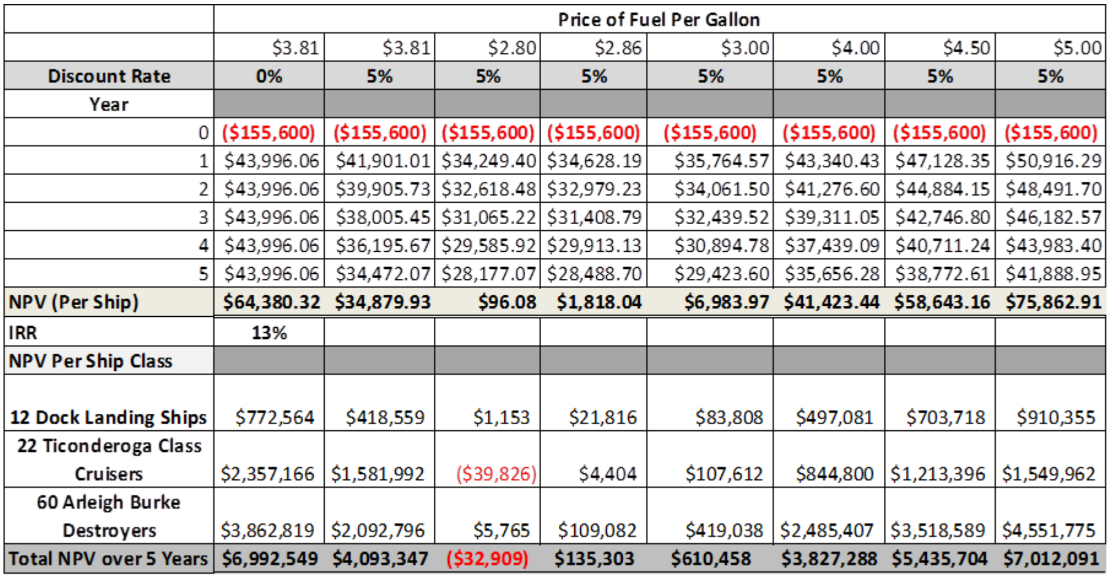

	Price of Fuel Per Gallon							
	$3.81	$3.81	$2.80	$2.86	$3.00	$4.00	$4.50	$5.00
Discount Rate	**0%**	**5%**	**5%**	**5%**	**5%**	**5%**	**5%**	**5%**
Year								
0	**($155,600)**	**($155,600)**	**($155,600)**	**($155,600)**	**($155,600)**	**($155,600)**	**($155,600)**	**($155,600)**
1	$43,996.06	$41,901.01	$34,249.40	$34,628.19	$35,764.57	$43,340.43	$47,128.35	$50,916.29
2	$43,996.06	$39,905.73	$32,618.48	$32,979.23	$34,061.50	$41,276.60	$44,884.15	$48,491.70
3	$43,996.06	$38,005.45	$31,065.22	$31,408.79	$32,439.52	$39,311.05	$42,746.80	$46,182.57
4	$43,996.06	$36,195.67	$29,585.92	$29,913.13	$30,894.78	$37,439.09	$40,711.24	$43,983.40
5	$43,996.06	$34,472.07	$28,177.07	$28,488.70	$29,423.60	$35,656.28	$38,772.61	$41,888.95
NPV (Per Ship)	**$64,380.32**	**$34,879.93**	**$96.08**	**$1,818.04**	**$6,983.97**	**$41,423.44**	**$58,643.16**	**$75,862.91**
IRR	**13%**							
NPV Per Ship Class								
12 Dock Landing Ships	$772,564	$418,559	$1,153	$21,816	$83,808	$497,081	$703,718	$910,355
22 Ticonderoga Class Cruisers	$2,357,166	$1,581,992	($39,826)	$4,404	$107,612	$844,800	$1,213,396	$1,549,962
60 Arleigh Burke Destroyers	$3,862,819	$2,092,796	$5,765	$109,082	$419,038	$2,485,407	$3,518,589	$4,551,775
Total NPV over 5 Years	**$6,992,549**	**$4,093,347**	**($32,909)**	**$135,303**	**$610,458**	**$3,827,288**	**$5,435,704**	**$7,012,091**

THIS PAGE INTENTIONALLY LEFT BLANK

V. CONCLUSION

A. SUMMARY

This thesis researched how virtual technology can reduce Total Cost of Ownership on a naval vessel. The origin and evolution of computing and virtual technology was presented to give the reader a thorough understanding of how information technology has evolved. With advances in both software and hardware, the consolidation of desktops and servers can now be accomplished in a virtual environment. The advancements of computer processors were researched to understand how efficient and powerful they have become to manage a virtual environment. By transforming a multi-processor blade server into a virtual environment, computing processing can be accomplished efficiently vice independently at a thick client workstation. TCO was identified as an effective method of cost analysis to compare current shipboard network infrastructure and a proposed private cloud utilizing thin clients on a naval vessel in the Consolidated Afloat Network Enterprise Services (CANES) program. Fuel consumption and shore power costs were identified as a direct cost comparison of thin and thick clients. Lastly, the total energy savings were projected for a five-year period and analyzed with hardware implementation costs for a naval vessel.

The historical data analysis of the Intel Xeon server processor was conducted to determine how virtual technology is possible because the Xeon series has developed significantly over the past 17 years. The Xeon processor has developed from being a single core processor to having 10 cores in a single processor. The key factor in doing virtualization is processing power and being able to handle multiple tasks simultaneously that is measured in Floating Operating Points Per Second (FLOPS). With increasing processor capabilities creates increasing electrical power consumption that adds to operating costs for normal operation and cooling. Instead of measuring actual power consumption of a processor, the heat dissipated in watts known as Thermal Design Power, was compared against various processor models. The early Xeon processors generated 38 Watts of heat and had a processing capability of .2 GFLOPS and newer processors dissipate between 65 and 130 Watts of heat with 10 cores and 96 GFLOPS of processing power. However, the earlier Xeon series processors have a Watt/GFLOP ratio of 15 to 190, and the newer 3000 series processors vary from 1.5 to 4. Therefore, newer multicore processors enabled for virtualization are energy efficient by emitting less heat per GFLOP increase.

To determine the difference in direct costs, average power consumption was measured for a new thin client and private cloud server against the current legacy thick client server architecture in various power states. By consolidating and reducing the number of workstations on a ship from 225 to 200 workstations due to features in the CANES program that combines UNCLASSIFIED and CLASSIFIED workstations into one, a reduction of 27 kilowatt hours can be achieved. With the reduced electrical load, shore power costs can be reduced by $3.75 an hour during peak usage and $2.11 an hour when machines are not in use. When a ship is underway, 2.5 gallons of fuel can be saved each hour by a ship running either a Ships Service Gas Turbine Generator (SSGTG) or a Ships Service Diesel Generator (SSDG). The reduced direct costs are contributed to the common computing environment where desktops and servers are virtualized instead of running independently and consuming unnecessary electricity.

Over a year, a typical ship with 200 workstations that utilizes an SSDG or SSGTG to produce electrical energy can save $44,000 in operating costs. Within 42 months, the energy savings alone could pay for the new hardware that was installed for a shipboard private cloud. When the savings are calculated over a five year period and given a 5% discount rate and the initial investment was hardware only, the Net Present Value (NPV) of the private cloud model is $34,880. The sensitive variable for having a high NPV is the price of fuel. If fuel price of fuel should drop below $2.80 per gallon, the NPV would go negative with a 5% discount rate.

In addition to being energy efficient, there are intangible benefits of the private cloud. One benefit is the ability to conduct patches and upgrades without affecting any users by using VMware virtualization software. By upgrading and performing tests on a template that is used to clone a virtual machine, this allows the network administrator to ensure the template is performing properly to minimize network downtime before deploying the upgrades. Another software benefit is VMware’s View Composer, which creates a pool of linked clones from a specific template and shares a base image to reduce overall data storage. By reducing data space, this will decrease storage requirements and costs for a private cloud.

Migrating information technology to a private cloud network is a viable solution for reducing operating costs onboard naval vessels. By consolidating legacy networks into one functional network, legacy systems can be extended and run on newer platforms that are no longer supported. Not only will operating costs decrease, but a private cloud is also agile and relevant in meeting future computing needs for warfighters at sea.

B. RECOMMENDATIONS AND FUTURE WORK

This thesis looked at a hypothetical model of a private cloud with key features from the CANES program. To further evaluate how virtualization is evolving, comparing the differences of VMware's ESXi V4.1 to ESXi V5.0 would be beneficial to evaluate how ESXi V5.0 handles high definition graphics where that function is not available on ESXi V4.1. In addition to evaluating different versions of VMware clouds, the new Microsoft Private Cloud using Windows Server, Hyper-V, and System Center software should also be evaluated and compared against VMware.

With more Power Over Ethernet (PoE) peripherals developing where a voltage is sent over a network to power switches, routers, and possibly thin clients in distant places where power receptacles are not present. The study of how PoE can affect TCO models could be beneficial for shipboard and remote environments where electrical power is not readily available.

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

Antedomenico, N. (2010). *Optimizing security of cloud computing within the DoD* (Master's thesis). Naval Postgraduate School.

Baliga, J., Ayre, R. W. A., Hinton, K., & Tucker, R. S. (2011). Green cloud computing: Balancing energy in processing, storage, and transport. *Proceedings of the IEEE, 99*(1), 149–167.

Bergin, T. J. (2000). *50 years of Army computing*. Aberdeen: U.S. Army Research Laboratory.

Britt, J. (2010). *Web application and thin clients in the Navy* (Master's thesis). Naval Postgraduate School.

California Board of Equalization. (2012). Electronic waste recycling fee. Retrieved on January 21, 2012, from http://www.boe.ca.gov/sptaxprog/ewfaqsgen.htm

Colaner, S. (2010, July). Saving dollars by making sense. *PC Today*, *8*(7), pp. 34–35.

CNET. (2012). Samsung Syncmaster NC240 MSRP. Retrieved February 12, 2012, from http://reviews.cnet.com/lcd-monitors/samsung-syncmaster-nc240-lcd/4505–3174_7–34534861.html

Dean, T. (2009). *Network + guide to networks* (5th ed.). Boston: Course Technology.

Defense Energy Support Center. (2012). Standard Fuel Prices in Dollars, FY2012 Rates. Retrieved February 14, 2012, from http://www.desc.dla.mil/dcm/files/Jan01FY12_Standard_Prices_Short-LongList.pdf

Dell. (2011). Manufactured Suggested Retail Price of M610 Server Blade. Retrieved February 12, 2012, from http://configure.us.dell.com/dellstore/config.aspx?oc=becwfk1&c=us&l=en&s=bsd&cs=04&model_id=poweredge-m610

Dell. (2012). Manufactured Suggested Retail Price of Dell Precision T1600 Workstation. Retrieved February 12, 2012, from http://premierconfigure.us.dell.com/dellstore/config.aspx?oc=g2734&c=us&l=en&s=fed&cs=RC1009777&model_id=precision-t1600&~tgt=cfg&customer_id=RC1009777

Department of Navy. (2007, December 17). *Navy Energy Usage Reporting System* (OPNAV Instruction 4100.11C). Washington, DC: Office of the Chief of Naval Operations.

Jennings, C. (2010, February 23–March 1). Is cloud computing really greener? *Computer Weekly*, 18–20.

John, S., & Han, M. W. (2006). A proposed framework for an effective integration of supporting environments for smart client application development. *Computing & Informatics, 2006*. ICOCI ‘06. International Conference on, 1–6.

Intel Corporation. (2011a). Intel Xeon Discontinued Processor Specification Sheet. Retrieved November 2, 2011, from http://www.intel.com/support/processors/sb/CS-020868.htm?wapkw=FLOPS

Intel Corporation. (2011b). Intel Xeon Processor 3000 Series Comparison Sheet. Retrieved November 2, 2011, from http://ark.intel.com/search/advanced/?s=t&FamilyText=Intel® Xeon® Processor 3000 Sequence

Intel Corporation. (2011c). Intel Xeon Processor 3000 Series Specification Sheet. Retrieved November 2, 2011, from http://download.intel.com/support/processors/xeon/sb/Xeon3000.pdf

Intel Corporation. (2011d). Intel Xeon Processor 5000 Series Comparison Sheet. Retrieved November 2, 2011, from http://ark.intel.com/search/advanced/?s=t&FamilyText=Intel® Xeon® Processor 5000 Sequence

Intel Corporation. (2011e). Intel Xeon Processor 5500 Series Specification Sheet. Retrieved November 2, 2011 from, http://download.intel.com/support/processors/xeon/sb/xeon_5500.pdf

Intel Corporation. (2011f). Intel Xeon Processor 7000 Series Comparison Sheet. Retrieved November 2, 2011, from http://ark.intel.com/search/advanced/?s=t&FamilyText=Intel® Xeon® Processor 7000 Sequence

Intel Corporation. (2011g). Intel Xeon Processor 7000 Series Specification Sheet. Retrieved November 2, 2011, from http://download.intel.com/support/processors/xeon/sb/xeon_7400.pdf

Intel Corporation. (2011h). Intel Xeon Processor E7 Series Comparison Sheet. Retrieved November 2, 2011, from http://ark.intel.com/search/advanced/?s=t&FamilyText=Intel® Xeon® Processor E7 Family

Intel Corporation. (2011i). Intel Xeon Processor E7 Series Specifications Sheet. Retrieved November 2, 2011, from http://download.intel.com/support/processors/xeon/sb/xeon_E7–8800.pdf

Intel Corporation. (2011j). Thermal Design Power. Retrieved November 2, 2011, from http://www.intel.com/support/motherboards/server/sb/CS-030213.htm?wapkw=Xeon%20W5580

Kondo, D., Javadi, B., Malecot, P., Cappello, F., & Anderson, D. P. (2009). Cost-benefit analysis of cloud computing versus desktop grids. *Parallel & Distributed Processing, 2009*. IPDPS 2009. IEEE International Symposium on, 1–12.

Lam, M. (2010). *Lower total cost of ownership of ONE-NET by using thin client desktop deployment and virtualization-based server technology* (Master's thesis). Naval Postgraduate School, Monterey, CA.

Lane, D. (2001). Brayton Cycle: The Ideal Cycle for Gas-Turbine Engines In Relation to Power Plants. Retrieved January 28, 2012, from http://web.me.unr.edu/me372/Spring2001/Brayton%20Cycle.pdf

Liang-Jie Zhang, & Qun Zhou. (2009). CCOA: Cloud computing open architecture. *Web Services, 2009*. ICWS 2009. IEEE International Conference on, 607–616

Loffink, J. (2008). *Dell PowerEdge M1000e modular enclosure architecture*. Dell Enterprise White Paper. Retrieved September 21, 2011, from http://www.dell.com/downloads/global/products/pedge/en/pedge_m1000e_white_paper.pdf

Lowe, S., McCarty, J., Johnson, M., (2010). *VMware vSphere™ 4 Administration instant reference*. Indianapolis: Wiley Publishing Inc.

Mabus, R. (2011, October 13). *Remarks by the Honorable Ray Mabus, Secretary of the Navy, Navy Energy Forum.* Washington, DC. Retrieved from http://www.navy.mil/navydata/people/secnav/Mabus/Speech/NavyEnergyForum13Oct2011.pdf

Marshall, D., Beaver, S., & McCarty, J., (2009). *VMware ESX Essentials in the virtual data center*. New York: CRC Press.

Minqi Zhou, Rong Zhang, Wei Xie, Weining Qian, & Aoying Zhou. (2010). Security and privacy in cloud computing: A survey. Semantics Knowledge and Grid (SKG), *2010 Sixth International Conference* on, 105–112.

National Institute of Standards and Technology. (2012). General Tables of Units of Measure. Retrieved February 14, 2012, from http://www.nist.gov/pml/wmd/pubs/upload/AppC-12-hb44-final.pdf

NAVSEA (2010). Shipboard Energy Conservation Guide. Retrieved November 15, 2011, from http://www.i-encon.com/ENCON%20Guide%202010.pdf

O'Regan, G. (2008). *A brief history of computing*. London: Springer-Verlag London Limited.

RAND National Defense Research Institute. (2009). Consolidated Afloat Networks and Enterprise Services (CANES) Manpower, Personnel, and Training Implications. Santa Monica: RAND Corporation.

Righini, M. (2010). *Enabling Intel® virtualization technology features and benefits.* Intel White Paper. Retrieved January 15, 2012 from, http://www.intel.com/content/dam/www/public/us/en/documents/white-papers/virtualization-enabling-intel-virtualization-technology-features-and-benefits-paper.pdf

Rittinghouse, J. W., & Ransome, J. F. (2010). *Cloud computing implementation, management, and security*. Boca Raton: CRC Press.

Rognlie, A. B. (2010). *An analysis of ROI of the Consolidated Afloat Networks and Enterprise Services (CANES) Program* (Master's thesis). Naval Postgraduate School.

Rolls-Royce.(n.d.) Allison AG9140 and AG9140RF Ship Service Generators Fact Sheet. Retried March 5, 2011, from http://www.rollsroyce.com/Images/MMS%20FS%2053%2008%201%20Allison%20AG9140%20and%20AG9140RF%20_tcm92–9324.pdf

Samsung (2009). *Samsung SyncMaster NC240 User Manual*. Retrieved January 27, 2012, from http://downloadcenter.samsung.com/content/UM/201112/20111213132430406/BN59–00954A-11Eng.pdf

Strowd, H. D., & Lewis, G. A. (2010). *T-Check in system-of-systems technologies: Cloud computing*. Hanscom AFB: Carnegie Mellon University.

SPAWAR. (2012). Products and Services. Retrieved January 18, 2012, from http://www.public.navy.mil/spawar/Atlantic/ProductsServices/Pages/NTCSS.aspx

Swedin, E. G., & Ferro, D. L. (2005). *Computers: The life story of a technology*. Westport, CT: Greenwood Publishing Group Inc.

Tech Terms. (2012). Floating Operating Points per Second. Retrieved on January 15, 2012 from, http://www.techterms.com/definition/flops

Tiglao, L. (2010). *Virtual Machine Modules for use by DoD C4I Support Centers* (Master's thesis). Naval Postgraduate School.

U.S. Energy Information Administration. (2011). *Average Retail Price of Electricity to Ultimate Customers by End-Use Sector, by State, December 2011 and 2010*. Retrieved February 18, 2012, from http://www.eia.gov/electricity/monthly/excel/epmxlfile5_6_a.xls

U.S. Library of Congress, Congressional Research Service. (2006, December 11). *Navy Ship Propulsion Technologies: Options for Reducing Oil Use* by Ronald O'Rourke. (CRS Report No. RL33360). Washington, DC: Office of Congressional Information and Publishing.

U.S. Navy. (2005). *Vision Presence Power 2005: A Program Guide to the U.S. Navy*. Retrieved January 14, 2011, from http://www.navy.mil/navydata/policy/vision/vis05/vpp05-ch3_forcenet.pdf

U.S. Navy. (2012). *U.S. Navy Fact File for Cruisers and Destroyers*. Retrieved January 20, 2012, from http://www.navy.mil/navydata/fact_display.asp?cid=4200&tid=800&ct=4

West, R., & Daigle, S. (2004). *Total Cost of Ownership: A Strategic Tool for ERP Planning and Implementation*. Retrieved December 22, 2011, from http://net.educause.edu/ir/library/pdf/ERB0401.pdf

Wyse Technology. (2011). *P20 Specification Sheet*. Retrieved January 25, 2012, from http://www.wyse.com/fulfillment/downloads/Wyse-P20.pdf

Wyse Technology. (2012). Wyse P20 MSRP. Retrieved February 12, 2012, from http://www.wyse.com.au/products/howtobuy/hotproducts.asp

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
   Ft. Belvoir, Virginia

2. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California

3. Prof. Dan Boger
   Chair, Department of Information Sciences
   Naval Postgraduate School
   Monterey, California

4. Prof. Glenn R. Cook
   Department of Information Sciences, Cod IS/Pr
   Naval Postgraduate School
   Monterey, California

5. Professor Albert Barreto III
   Department of Information Sciences, Cod IS/Pr
   Naval Postgraduate School
   Monterey, California

6. LT BJ S. Boothe
   Joliet, Montana