2001-03

Virtual Desktop Infrastructures (VDIs) Supporting Agile Sea Basing A Study on Improving Embarkable Integration Onboard Amphibious Flagships

Matthias, Robert D.; Becker, John P.

Monterey, California. Naval Postgraduate School

https://hdl.handle.net/10945/6829

NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

VIRTUAL DESKTOP INFRASTRUCTURES (VDIs) SUPPORTING AGILE SEABASING: A STUDY ON IMPROVING EMBARKABLE INTEGRATION ONBOARD AMPHIBIOUS FLAGSHIPS

by

Robert D. Matthias
John P. Becker

March 2012

Thesis Co-Advisors: Shelley Gallup
William J. Welch

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			*Form Approved OMB No. 0704–0188*
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202–4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704–0188) Washington DC 20503.			
1. AGENCY USE ONLY *(Leave blank)*	**2. REPORT DATE** March 2012	**3. REPORT TYPE AND DATES COVERED** Master's Thesis	
4. TITLE AND SUBTITLE Virtual Desktop Infrastructures (VDIs) Supporting Agile Sea Basing: A Study on Improving Embarkable Integration Onboard Amphibious Flagships			**5. FUNDING NUMBERS**
6. AUTHOR(S) Robert D. Matthias, John P. Becker			
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** A

13. ABSTRACT (maximum 200 words).

Given the deliberate planning and effective practice of seabasing in support of the DoD's broadening Range of Military Operations (ROMO), the U.S. Navy (USN) has found itself playing host to a litany of disparate organizations working together in highly dynamic environments. No USN platform is better suited to meet the challenge of organizational integration than our amphibious units. Designed for and well-practiced in the embarkation, deployment, and debarkation of a two-thousand-strong Marine Expeditionary Unit (MEU), amphibious leadership and supporting personnel are still hindered by inefficient IT integration processes, hardware incompatibilities, and resulting security measures. Network integration for embarkable personnel and their deployed equipment has been identified as a priority requiring improvement. Changing institutionalized architectures and their supporting processes cannot deliver sufficient agility. Re-engineering calls for a technology insertion as well. Virtual Desktop Infrastructures (VDIs) may prove a viable option for enhanced interoperability onboard amphibious ships in the near future.

14. SUBJECT TERMS Range of Military Operations (ROMO), Network Integration, Network Integration, Virtualization, Virtual Desktop Infrastructure			**15. NUMBER OF PAGES** 81
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	**18. SECURITY CLASSIFICATION OF THIS PAGE** Unclassified	**19. SECURITY CLASSIFICATION OF ABSTRACT** Unclassified	**20. LIMITATION OF ABSTRACT** UU

NSN 7540–01–280–5500

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

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

VIRTUAL DESKTOP INFRASTRUCTURES (VDIs) SUPPORTING AGILE SEABASING: A STUDY ON IMPROVING EMBARKABLE INTEGRATION ONBOARD AMPHIBIOUS FLAGSHIPS

Robert D. Matthias
Lieutenant Commander, United States Navy
B.S., Naval Academy, 1999

John P. Becker
Lieutenant, United States Navy
B.S., The Citadel, 2005

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN INFORMATION TECHNOLOGY MANAGEMENT

from the

NAVAL POSTGRADUATE SCHOOL
March 2012

Author: Robert D. Matthias
John P. Becker

Approved by: Dr. Shelley Gallup
Thesis Co-Advisor

William J. Welch
Thesis Co-Advisor

Dr. Dan Boger
Chair, Department of Information Sciences

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Given the deliberate planning and effective practice of seabasing in support of the DoD's broadening Range of Military Operations (ROMO), the U.S. Navy (USN) has found itself playing host to a litany of disparate organizations working together in highly dynamic environments. No USN platform is better suited to meet the challenge of organizational integration than our amphibious units. Designed for and well-practiced in the embarkation, deployment, and debarkation of a two-thousand-strong Marine Expeditionary Unit (MEU), amphibious leadership and supporting personnel are still hindered by inefficient IT integration processes, hardware incompatibilities, and resulting security measures. Network integration for embarkable personnel and their deployed equipment has been identified as a priority requiring improvement. Changing institutionalized architectures and their supporting processes cannot deliver sufficient agility. Re-engineering calls for a technology insertion as well. Virtual Desktop Infrastructures (VDIs) may prove a viable option for enhanced interoperability onboard amphibious ships in the near future.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

LIST OF FIGURES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

ADP	Automated Data Processing
ARG	Amphibious Readiness Group
C2	Command and Control
CCE	Common Computing Environment
COTS	Commercial Off-the-Shelf
COMPOSE	Common PC Operating System Environment
CPU	Central Processing Unit
DoD	Department of Defense
DSR	Data Systems Repair
EW	Electronic Warfare
ESG	Expeditionary Strike Group
FFC	Fleet Forces Command
FIRE	FORCENet Innovation and Research Enterprise
GB	Gigabyte
Gb	Gigabit
Gb/s	Gigabit per second
GO	Government Organization
HDD	Hard Disk Drive
IO	Information Operations
I/O	Input/Output
I/Ops	Input/Output per second
ISR	Information, Surveillance, Reconnaissance
IT	Information Technology
JNOC	Joint Networks Operation Center
Kbps	Kilobytes per second
KVA	Knowledge Value Added
LAN	Local Area Network
LCS	Littoral Combat Ship
LFOC	Landing Forces Operations Center
LHA	Amphibious Assault Ship (General Purpose)

LHD	Amphibious Assault (Multi-Purpose)
MET	Mission Essential Task
MEU	Marine Expeditionary Unit
MBps	Megabytes per second
NET	Networks
NGT	Next Generation Technologies
NGO	Non-Government Organization
NPS	Naval Postgraduate School
OS	Operating System
QA	Quality Assurance
ROMO	Range of Military Operations
SPAWAR	Space and Naval Warfare System Command
SA	Situational Awareness
SA	System Administrator
SG	Strike Group
SSD	Solid State Drive
TTP	Tactics, Techniques, and Procedures
TYCOM	Type Commander
TW	Trident Warrior
VDI	Virtual Desktop Infrastructure
VM	Virtual Machine
UNCLAS	Unclassified
USMC	United States Marine Corps
USN	United States Navy

ACKNOWLEDGMENTS

We would like to show our gratitude to all the positive influences we have experienced prior to the finalization of this thesis. A special acknowledgement goes to Vice Admiral Kendall L. Card, who introduced the conceptual application of thin-client technologies to address a litany of DoD IT issues. We are grateful to all the Sailors and Marines who have endured with their existing resources, maintained successful exercises, operations, and deployments, and have taken the time to share their perspectives with us.

We would like to thank Trident Warrior's Dr. Shelley Gallup and NPS Lecturer William Welch. Both underway and ashore, Dr. Gallup provided us with invaluable mentorship, every possible resource available to him, and a constant fueling of our confidence to maintain the ownership of this research. Mr. Welch's positive approach to information systems proved to be a continuous source of inspiration and his guidance was instrumental toward making our thesis clear, compelling, and relevant to the U.S. Navy.

Finally, we are so very thankful for the individual support of our wives and family. Through their love and sacrifices, we have not only been able to achieve our goals, but enjoy this entire experience.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

In the late spring of 2007, a combat weary 15th Marine Expeditionary Unit (MEU) re-embarked on their Expeditionary Strike Group (ESG) ships to begin a steady voyage home to Camp Pendleton, CA. Nearly 2000 Marines loaded themselves and their operational gear onboard the ESG's amphibious flagship, USS BOXER (LHD-4). Finally removed from the insurgency hot-bed within Al Anbar province, the administration of a twice-extended, nine and a half month deployment was just about to begin. Prior to reports, evaluations, awards, and e-mails reconnecting families could even be considered, the MEU's communications/IT personnel required some 300 unclassified workstations to integrate back onto the BOXER's Local Area Network (LAN). Although routinely practiced onboard amphibious ships the large-scale network integration was slow and burdensome.

Surrounded by arguably the latest IT throughout the ship, both Sailors and Marines could not benefit from technology to aid the arduous process of how workstations were integrated. Not only restricted to frequently used Marine equipment, LAN integration is exceedingly difficult with any group the capable amphibious platform could sea-base for missions, anticipated or unexpected. LHDs have the potential to conduct a multitude of military and humanitarian tasking allowing Government or Non-Government Organization (GO/NGO) causes, respectfully, to seek out assistance from the U.S. Navy. To be viable, their technology will have to be able to circumvent traditional USN enterprise views of hardware and software standardization and compliance policy. Current architecture, technology, and business practices have become distracted with such policy where existing technology inhibits rather than enhances overall productively.

Whether planning for a Marine landing onto an opposed beachfront or supporting Doctors Without Borders during Humanitarian Assistance/Disaster Relief operations, the dynamic environment the U.S. Navy operates within today calls for the ability to lead through hosting traditional and unanticipated disparate organizations. When time is a critical resource, Naval afloat units lack the ability to rapidly and efficiently facilitate inorganic equipment that enables those who embark to support U.S. objectives. This chapter provides a background on the observed problem and states the thesis purpose. Additionally, research questions, scope, and the approach to this thesis are presented.

A. BACKGROUND

Today's United States Department of Defense is bracing for a future where institutional beliefs are directly questioned in efforts to better posture its service members for improved readiness. Top government and military leadership have indicated that effective pursuits of national interest cannot be achieved one-dimensionally through a single branch of military service or solely through U.S. joint military involvement. The dynamic environment the present world has to offer requires joint, coalition, interagency, state, and non-state coordination. The latest U.S. National Military Strategy pushes this theme further through the encouragement:

> To facilitate interagency and enable international interoperability before crises occur" (Mullen, 2011). While considering *The Future of Warfare*, Hammes (2005) attempts to pinpoint a fundamental need "to shift our focus from pure technology to technology that supports the human elements required for the long-term, interagency efforts required to win future conflicts. (p. 277)

Arguably, the first aspect to begin moving toward improved interoperability is to focus on how the U.S. military conducts operations jointly, from service to service. Seamless coordination internally through our branches of military can serve as benchmarks to foster interagency and foreign unity of intended efforts. Furthermore, communication between organic and inorganic personnel should be the first priority. U.S. Joint Operations Doctrine (2008) requires that "assigned forces should be capable of complementary mutual support and full communications interoperability" (p. 189).

When asked about technology needs and capabilities wished for (in her future budgeting requirements/wish list), former Expeditionary Strike Group Two (ESG-2) Commander, Rear Admiral Michelle Howard stated a need for "interoperability as we move around theater, not just within Navy forces, but with other services, joint, coalition partners and non-governmental organizations" (CHIPS, 2011). Rear Admiral Howard had the opportunity to experience how redefined interoperability within and outside traditional means was necessary for mission success, yet, in reality, could only address the shortfall within her wish list and not within her provisioned budget. As Commander of ESG-2, she stood witness to the U.S. Navy's emerging requirements to effectively collaborate and coordinate through commonality in technology. Moreover, leaders like Rear Admiral Howard have come to the realization that the Navy must also take a leading role among disparate entities. Often that role is best facilitated through the actual onboard hosting of organizational and operational units.

The U.S. Navy has a unique ability to sea-base, establishing major operational centers upon international and coalition waters. Sea-basing facilitates the structure for operations and grants central platforms for agreed intentions. "Sea-based operations use revolutionary information superiority and dispersed, networked force capabilities to deliver unprecedented offensive power, defensive assurance, and operational independence to Joint Force Commanders" (Clark, 2002). Although naval platforms are tasked to support this invaluable operational agility, formidable challenges of force integration arise when addressing joint or coalition interoperability. Currently, manning levels, evolved processes, and enterprise technology result in unrealistic integration expectations, costly man-hour usage, restricted battlefield preparation, and compromised Situational Awareness (SA).

Could recent advances in information technology and virtualization benefit joint integration with currently deployed elements and their systems? A step in the right direction would be to enhance the surface fleet's ability to economically communicate with any trusted agent able to add value to meet mission objectives. This collaboration and invaluable information exchange may come in very conventional or unconventional forms and over unanticipated platforms. The ability to move with a desired unconventional uniformity will hinge on our military's ability to improve its communications interoperability. Arguably, the most experienced armed services conducting combined operations are the U.S. Navy and Marine Corps (USMC).

Although within the same department, the proud tradition and unique capabilities distinguishing the Navy and Marine Corps team also harbors differences in personnel, procedures, and equipment.

"Joint operational flexibility will be greatly enhanced by employing pre-positioned shipping that does not have to enter port to offload" (Clark, 2002). Few other services embrace the missions and daily practice of joint interoperability better than the U.S. Navy's Amphibious Readiness Groups (ARGs), former Expeditionary Strike Groups. Primarily established for the amphibious landing of Marine Expeditionary Units (MEUs), ARGs share a wealth of other missions absolutely requiring joint, coalition, federal, and non-governmental partnership to include (22d MEU, 2011):

- Peacekeeping/Enforcement
- Humanitarian/Disaster Relief
- Security Operations
- Noncombatant Evacuation Operations
- Reinforcement Operations
- Amphibious Raids/Assaults/Demonstrations
- Tactical Deception Operations
- Airfield/Port Seizures
- Show-of-Force Operations
- Reconnaissance and Surveillance
- Seizure/Recovery of Offshore Energy Facilities

To accomplish these multiple types of mission, a MEU will embark over 2200 Marines composed of (typical example provided by 22d MEU):

1. Command Element (CE)

169 personnel. Serves as the headquarters for the entire unit and allows a single command to exercise control over all ground, aviation, and combat service support forces (22d MEU, 2011).

2. Ground Combat Element (GCE)

1200 personnel. Provides the MEU with its main combat punch. Built around a Marine infantry battalion, the GCE is reinforced with tanks, artillery, amphibious vehicles, engineers, and reconnaissance assets (22d MEU, 2011).

3. Aviation Combat Element (ACE)

417 personnel. The ACE consists of a composite medium helicopter squadron containing transport helicopters of various models and capabilities, attack helicopters and jets, air defense teams, and all necessary ground support assets (22d MEU, 2011).

4. MEU Service Support Group (MSSG):

275 personnel. Providing the MEU with mission-essential support such as medical/dental assistance, motor transport, supply, equipment maintenance, and landing is the mission of the LCE (22d MEU, 2011).

The most routine ARG evolution encompasses the embarkation and debarkation of the Marines in support of “amphibious, security, noncombatant evacuation, humanitarian assistance, and special operations” (15th MEU, 2011). Because MEUs ride their respective amphibious surface ships for months while steaming toward their theater of operations, a successful embark will facilitate a successful debark.

Prior to the transit, embarking Marines are required to integrate with the ship’s infrastructure in order to seamlessly continue their planning for upcoming operations. Differing standards between host ship and MEU computing platforms create a quagmire of identifying compatibility and ensuring security standards are upheld. General attempts at streamlining surface ships network management have fallen short for addressing inorganic shipboard riders.

In a recent 2nd Fleet hosted ESG-2 and 2nd Marine Expeditionary Unit exercise, Bold Alligator 2011, multiple incompatibilities were discovered raising concerns about their “ability to respond to a dynamic threat” (Anderson, 2011). Marines and other military or civilian organizations have typically been required to perform a full scan or, most likely, a hard drive wipe and reimage, if they wish to connect to a host ship’s network. Illustrated in a later chapter, this current practice includes all Navy Marine Corps Intranet (NMCI) workstations for the 2200 Marines within an embarking MEU on their assigned ships for deployment and causes a substantial delay for required C2 planning prior to operations.

From a Navy Staff and Marine perspective, basic network interoperability and incompatibility lies with their legacy NMCI workstations onboard the ships they intend to embark. Since its introduction, NMCI has been met with mixed reviews. Much of the frustration has stemmed from the lack of compatibility between U.S. Navy platforms and the shore based NMCI inorganic workstations. The incompatibility of U.S. Navy afloat platform networks with shore based NMCI and inorganic workstations has been shown to inhibit information workflow, integration of planning, and timeliness of action in support of assigned operational objectives for staff, crew and embarked personnel onboard U.S. Navy platforms.

It is important to first understand what the Navy Marine Corps Intranet is and why it exists. The U.S. Navy NMCI program was an ACAT I program of record, with a contract cost of $9.3 billion extending from 1999 to 2010. A new contract has since been implemented extending remnants of NMCI into the foreseeable future. The Navy and Marine Corps advertised the intranet as the answer to all their computing and communication needs in the online realm. NMCI was initially created to "provide an interoperable command and control network needed for transitioning to a net-centric environment" and consolidated over 6,000 disparate networks into one (Hewlett-Packard, 2011). When NMCI was delivered, it became clear that it was not available to the most important group in the Navy and Marine Corps, the deployed war fighter. According to Christopher Ragano (2006) of ONR, "deployed forces, although they may have NMCI equipment, find that the equipment does not support their mission, and they have to carry out work-arounds using other equipment" (Jordan, 2006, p. 9). As such, U.S. Navy personnel have been unsatisfied with the overall performance of NMCI. The GAO (2006), Government Accountability Office, found that "the end user satisfaction rate was approximately 74 percent, below the target of 85 percent" (Jordan, 2006, p. 10). The GAO also found that operational users (commanders and network operators) had a much lower satisfaction rate (Jordan, 2006, p. 10). These numbers are up from 60 percent, which were reported in 2006 but are still well below the target (Jordan, 2006, p. 10).

Over a decade after its inception, potential technology finally exists that may provide solutions to the incompatibility issues of NMCI, as well as non-NMCI users, while improving information workflow, integration of planning, and timeliness of action in support of assigned operational objectives for staff, crew and embarked personnel onboard U.S. Navy platforms. Review of the current deployment of virtualized computer technology and solid state memory technology can bolster arguments as to the potential solutions to the NMCI and its general incompatibility issue. Virtualization, VMware, thin-client desktops and solid state memory may provide the answers to these issues.

B. PROBEM STATEMENT

The incompatibility of U.S. Navy afloat platform networks with shore-based NMCI and inorganic workstations, along with choke points in the embarkation processes, inhibits information and workflow, integration of planning, and timeliness of action in support of assigned operational objectives for staff, crew, and embarked personnel of U.S. Naval vessels.

C. PURPOSE STATEMENT

The purpose of this study is to assess whether if virtualization technology, paired with solid state drives and process re-engineering onboard U.S. Naval afloat platforms, can improve work flow, integration of planning, and timeliness of action in support of nominal or assigned operational objectives for staff, crew, and embarked personnel of U.S. Naval vessels.

D. RESEARCH QUESTIONS

1. Do hosting Naval platforms have adequate facilities to foster full integration and collaboration of embarking personnel spanning all operational possibilities?

2. How can Virtual Desktop Infrastructure (VDI) and solid state memory equipment be leveraged to achieve higher level of integration amongst disparate organizations?

3. How can the integration of embarking personnel and their necessary equipment be modeled and analyzed a for better understanding of current and improved practices?

4. What are the risks and benefits resulting from a ship-hosted Virtual Desktop Infrastructure (VDI) for embarkables?

5. What research has been done or is there current practice advancing the use of virtualization or VDIs either commercially or militarily?

6. What are the technical capabilities of major VDI commercial organizations?

7. How could a revised model aid in predicting host vessel hardware requirements for a VDI in order to better understand practical costs?

E. SCOPE

This study will focus on the effectiveness and practically of VDI and SSD for embarkables on U.S. Navy ships and how the integration of new technology can improve current processes. It will focus specifically on risks and benefits to information and work flow, integration of planning, and timeliness of action in support embarked personnel on U.S. amphibious platforms.

F. APPROACH AND METHODOLOGY

Literature review, system analysis, application testing, and SME interviews will be used to collect data and information. Equally large amounts of literature review, and data and system analysis will be required to adequately frame the problem statement and fully address our research questions. Application testing and a majority of our interviews have been conducted on and around our time participating with Trident Warrior 2011. With the knowledge we acquire through these methods, we will begin to draw conclusions and detail our understanding on the effectiveness and practicality of VDI onboard Naval afloat platforms.

Action Research, "concerned to create organizational change and simultaneously study the process," is best suited for this method of research (Baskerville & Myers, 2004, pp. 329–330). Literature research, prior experience, and system testing onboard WASP will help identify underlying causes and further develop theoretical assumptions. System analysis of embarkable integration will foster an understanding of current practice and allow for suggestions toward an improved end state. Action taken, or intervention, will encompass the Trident Warrior Next Generation Technology (NGT) Virtual Desktop Infrastructure (VDI) for embarkables efforts.

G. THESIS ORGANIZATION

I. Introduction. Brief overview of the background, problem, and purpose of the research.

II. Virtualization Technology. Provides a general understanding of virtualization, its background, applications and benefits.

III. Solid State Drive Technology. Provides a general understanding of solid state drives, its background, applications and benefits.

IV. Amphibious Integration. Used to develop an understanding of shipboard integration processes, networking practices and technology needs.

V. Assessment of Trident Warrior'11 VDI onboard USS WASP (LHD Analysis of virtualized desktops as tested through the 2011 Trident Warrior initiative.

VI. Re-engineering Amphibious Integration. Present a model to assess the practicality of VDI and SSD onboard USN platforms in the adjustment of shipboard embarkation processes.

VII. Conclusions. Discusses key areas where virtualization technology promotes agile enterprise architecture and improvements to observed business processes supporting integration.

THIS PAGE INTENTIONALLY LEFT BLANK

II. VIRTUALIZATION AND THIN CLIENT TECHNOLOGY

A. GENERAL DISCUSSION

In the last decade, virtualization has become a popular subject that is touted as possibly revolutionizing networking and computing. In its most basic form, virtualization allows a PC or server to run multiple operating systems, a guest and a host, on one machine and is illustrated in Figure 1, which compares a non-virtualized system against a virtual system (Vaughn, 2006). Through the use of a hypervisor, a software technology that "emulates a hardware device for each virtual operating system and handles each operating system's communications with the CPU, the storage medium, and the network," a server or PC can run many operating systems simultaneously (Vaughan, 2006, p. 12). This allows for the complete optimization of a PC or server and can provide multiple virtual machines, or a virtual desktop infrastructure (VDI), to the end users. The virtual machines, in turn, can provide all the benefits of a standalone PC without the cost of acquisition or physical set-up and maintenance. VDIs can be created ad hoc as computing needs change.

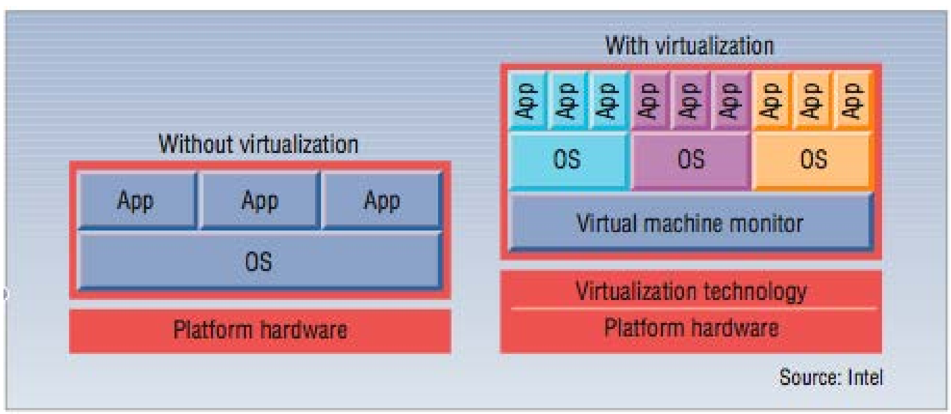

Figure 1. Non-Virtualized Platform vs. Virtualized Platform (From Vaughn, 2006)

Most modern computers run dual and quad core processors. The Department of Defense and the general public are buying computer systems believing that they are obtaining superior performing machines but this may not be the case. The truth is, operating systems routinely do not take full advantage of the processing power of today's dual and quad core servers and PCs. According to the IDC, the International Data Corporation, servers are only “utilizing 10–15% of their total capacity” (VMware, 2011). This indicates companies, the general public, and our armed forces are paying for capabilities that they possibly will never utilize. Virtualization can remedy this underutilization and save money wasted on hardware. Organizations have reported “60–80% utilization rates” through the use of virtualization (VMware, 2011). This may become extremely critical as the United States government, commercial organizations, and citizens in general are looking for ways to accomplish more while spending less.

In addition to the underutilization of processing, an underutilization of memory exists as well. A study conducted by Intel confirmed this apparent underutilization in a study that showed that only “50 percent of approximately 3,000 servers tested used no more than 1 GB of memory” (Intel, 2010, p. 2) and by “excluding the top quartile, analysis showed that for 75 percent of servers, maximum memory consumption averaged about 1 GB,” (Intel, 2010, p. 2) and is depicted in Figures 2 and 3 (Intel, 2010, p. 2).

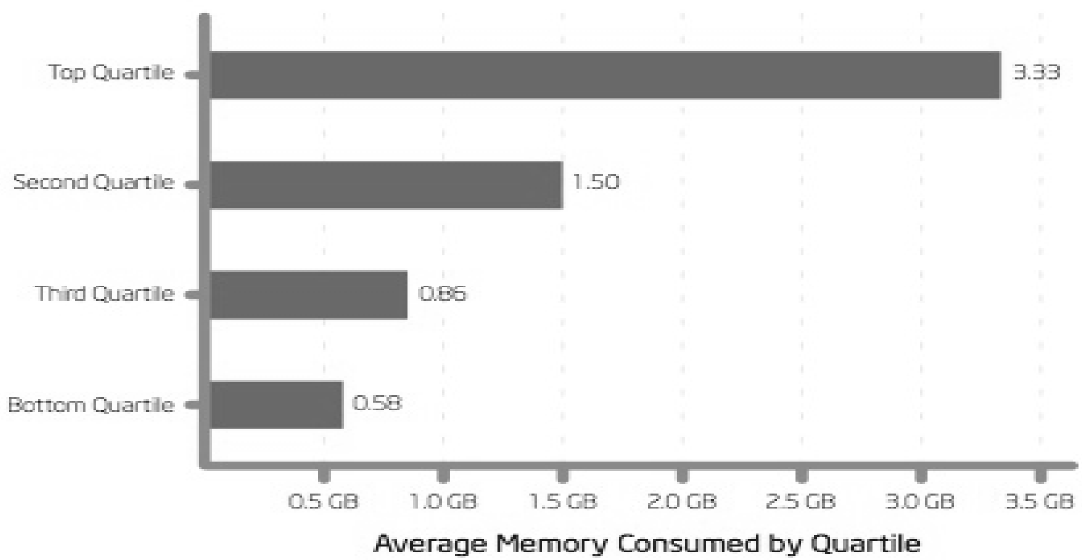

Figure 2. Average Memory Consumed by Quartile (From Intel, 2010, p.2)

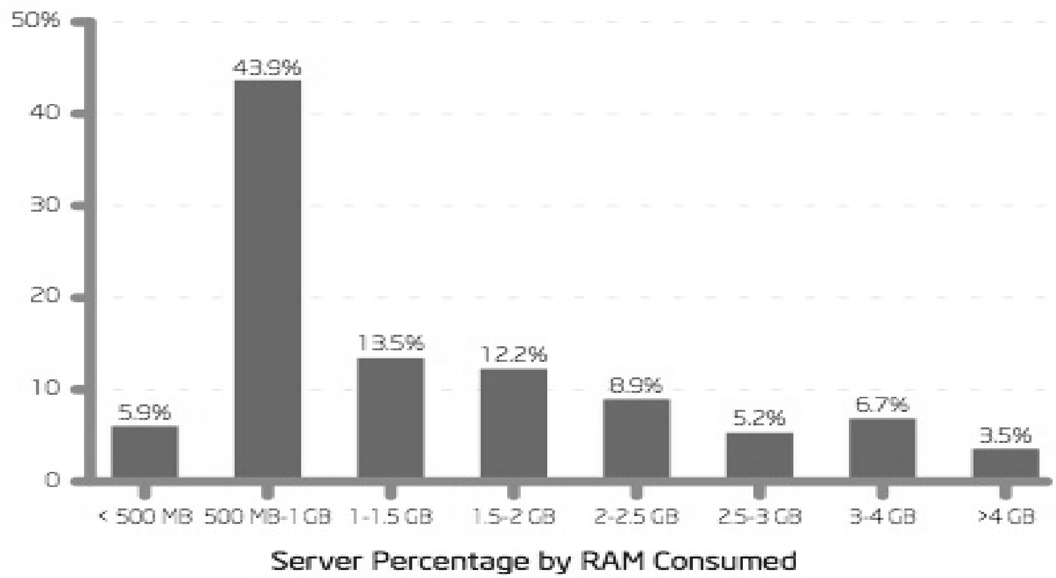

Figure 3. Server Percentage by Ram (From Intel, 2010, p. 2)

In general, computers today are far more powerful, have far more memory, and, as a whole, are much more capable than their counterparts from just a couple of years ago. Today's machines are "often under-used, incur significant space and management overhead, and the increased functionality that had made operating systems more capable has also made them fragile and vulnerable" (Rosenblum & Garfinkel, 2005, p. 39). Because of this, virtualization "will be less a vehicle for multitasking, as it was originally, and more a solution for security, reliability and integration." (Rosenblum & Garfinkel, 2005, p. 39). As the Navy looks to expand its mission capability by conducting more and more other than war operations, it will need to look at the best and most cost effective means of integrating outside organizations while protecting itself from intrusion into its networks.

B. HISTORY

Virtualization was originally used in the mid-1960s. During this decade, personal computers were not available, meaning computing only "belonged to large, expensive mainframe hardware" systems (Rosenblum & Garfinkel, 2005, p. 39). IBM was an early dominant force in computing and used virtualization heavily within mainframes since early systems did not have near the processing power of today's servers or even today's

PCs. Virtualization provided a "way to logically partition mainframe computers into separate virtual machines" (VMware, 2011). This partitioning allowed mainframes to "run multiple applications and processes at the same time" (VMware, 2011). With little or no processing power, miniscule amounts of memory compared to today's standards and overall crushing costs, getting the most out of the mainframes then was a must.

The reduction of costs for hardware and the rise of modern multi-tasking operating systems lead to the demise of virtualization. By the 1980s, the world moved away from virtualization all together. The "broad adoption of Windows and the emergence of Linux as server operating systems in the 1990s established x86 servers as the industry standard" and lead to the retreat from virtualization technology and methods (VMware, 2011). With the adoption of x86 servers new challenges arose. These challenges included:

1. Low Infrastructure Utilization

Typical x86 server deployments achieve an average utilization of only 10% to 15% of total capacity, according to International Data Corporation (IDC), a market research firm. Organizations typically run one application per server to avoid the risk of vulnerabilities in one application affecting the availability of another application on the same server (VMware, 2011).

2. Increasing Physical Infrastructure Costs

The operational costs to support growing physical infrastructure have steadily increased. Most computing infrastructure must remain operational at all times, resulting in power consumption, cooling and facilities costs that do not vary with utilization levels (VMware, 2011).

3. Increasing IT Management Costs

As computing environments become more complex, the level of specialized education and experience required for infrastructure management personnel and the associated costs of such personnel have increased. Organizations spend disproportionate time and resources on manual tasks associated with server maintenance, and thus require more personnel to complete these tasks. According to Gartner Research (2004), for every dollar an organization spends on hardware, it spends an additional $3.50 to support it over its lifetime (VMware, 2011).

4. Insufficient Failover and Disaster Protection

Organizations are increasingly affected by the downtime of critical server applications and inaccessibility of critical end user desktops. The threat of security attacks, natural disasters, health pandemics and terrorism has elevated the importance of business continuity planning for both desktops and servers (VMware, 2011).

5. High Maintenance End-User Desktops

Managing and securing enterprise desktops present numerous challenges. Controlling a distributed desktop environment and enforcing management, access and security policies without impairing users' ability to work effectively is complex and expensive. Numerous patches and upgrades must be continually applied to desktop environments to eliminate security vulnerabilities (VMware, 2011).

To combat these new challenges caused by the adoption of the x86 servers as the standard, alternative methods have to be adopted. Not a new method by any means, though one that has been used in the past with success, virtualization is again proving to be the answer to flaws and challenges presented with current computing needs.

C. STATE OF VIRTUALIZATION

Some giants of computing, Microsoft, Sun Microsystems, Intel, and IBM are all jumping into the business of virtualization and expect to earn billions of dollars in the sales. The sales of virtualization systems in 2008 totaled $6.7 billion and were expected to grow to $11.7 million by 2011 (Vaughan, 2008, p. 13). In the coming five years "North America and European virtualization aggregate markets will top $218 Billion" (Market Intel Group, 2011). It is estimated that "almost 75 percent of enterprises in the U.S. have already deployed virtualization in one form or another, and the virtualization market is growing by approximately 26 percent [annually] on average" (Vaughan, 2006,

p. 12). It is estimated that the “number of servers worldwide using virtualization has grown from just under 175,000 in 2004 to just over 1 million in 2009 (Vaughan, 2006, p. 12).

These trends in virtualization system sales indicate that virtualization adoption is not a small niche movement. The commitment of key corporations to develop, produce and sell systems and the expected growth of these systems is a clear indicator of the ever-growing demand for virtualization. For this kind of growth and commitment to take place, virtualization must present a compelling case for use and has done so to this point in time.

D. ADVANTAGES OF VIRTUALIZATION

As discussed previously, virtualization allows for multiple operating systems to be run on single servers or PCs but it also provides many other advantages as well. According to VMware there are five main advantages to the use of virtualization:

Get more capability, performance, and value out of existing resources: Pool common infrastructure resources and break the legacy “one application to one server” model with server consolidation. (VMware, 2011)

Reduce datacenter costs by reducing physical infrastructure and improving server to admin ratio: Fewer servers and related IT hardware means reduced real estate and reduced power and cooling requirements. Better management tools let you improve your server to admin ratio so personnel requirements are reduced as well (VMware, 2011).

Increase availability of hardware and applications for improved business continuity: Securely backup and migrate entire virtual environments with no interruption in service. Eliminate planned downtime and recover immediately from unplanned issues (VMware, 2011).

Gain operational flexibility: Respond to market changes with dynamic resource management, faster server provisioning and improved desktop and application deployment (VMware, 2011).

Improve desktop manageability and security: Deploy, manage and monitor secure desktop environments that users can access locally or remotely, with or without a network connection, on almost any standard desktop, laptop or tablet PC (VMware, 2011).

E. VIRTUALIZATION AND THE USN/USMC

As illustrated in the introduction, there exists a significant demand for efficient interoperability onboard ARGs. This plea extends to both professional and personal needs. Given the traditionally limited amount of time available to integrate workstations compared to the sheer number of personnel that embark, a clear call for virtualization presents itself. The embarkation of MEUs allows for the greatest potential of process and physical improvement through the use of virtualization technology.

The Marines embark ARGs more frequently than any other group but they are not the only type of unit which embarks. Navy Command Staffs, government organizations and non-governmental organizations may all embark on an ARG separately or together. At the mercy of integration processes illustrated later, each of these organizations requires joint operability and connectivity to efficiently plan for, and execute their assigned, often extremely short-notice, tasking. Through the use of virtualization, all embarking groups could have the ship's approved operating system and updated security pushed directly their carry-on systems. This would greatly reduce the time required from the ship's company for integration and allow ship riders to get to work almost immediately. During crisis response situations, the ability to work as soon as boots hit the deckplates is critical to overall mission success.

Virtualization allows groups outside the normal realm of the Department of Defense to embark without concern for encroachment of the ship's actual network. By providing virtual systems to embarked groups, a layer between the main network and the virtual system exists through the use of a hypervisor. According to VMware (2011):

> While virtual machines can share the physical resources of a single computer, they remain completely isolated from each other as if they were separate physical machines. Isolation is an important reason why the

> availability and security of applications running in a virtual environment is far superior to applications running in a traditional, non-virtualized system.

Those with virtual machines will not have a direct path to sensitive information that may exist within the network or a path to bring down the network. Also, each virtual machine's permissions can be quickly and precisely controlled, quelling any concerns over security. The Navy and Marine Corps could be able to work with an unexpected mission contributor with comparatively less concern for compatibility or security.

Further, virtualization has the potential to reduce the amount of hardware and software required to meet mission tasking. Virtualization optimizes servers, which reduces the need for an overabundance of servers to do the same amount of processing as a virtualized system. Less overall demand for servers processing leads to less hardware usage. This in turns leads to a reduction in the need for manpower and the ability to assign manpower to where it is better served. Along with the reduction in manpower, fewer servers require less power to run as well. Less power onboard Navy vessels equates to less fuel being burned. Onboard Navy vessels, any way to save fuel is looked upon favorably. The Honorable Ray Mabus, Secretary of the Navy, stated:

> We are also doing this (reducing fuel usage) to be better war fighters. A navy ship is at its most vulnerable when refueling. The USS Cole was refueling in the port of Aden in Yemen when it was attacked in 2000. Each of these reductions is a reduction in cost incurred by the Navy. (McKenna, 2011, p. 29)

Through virtualization, servers and computers are optimized. Optimization is not just a physical measure, but also a monetary measure. The Department of Defense is continually looking at ways to save money without losing operational capabilities. Because the Navy and Marine Corps must do more with less, virtualization is a way to achieve this goal. Through virtualization the USN and USMC can expand joint mission capability by opening the availability of hardware and applications. Typically, the DoD will buy a specialized or unique type of hardware or application that works well for the group or unit but it will lack capabilities to jointly work with other groups or bring them on board. Through the improved security and confidence provided by the virtual environment, naval vessels can utilize new or different hardware or applications on the fly.

These new capabilities may allow them to more effectively meet mission requirements or increase joint interoperability. In terms of software, license management is improved significantly. DoD units would only have to acquire minimal copies of required software to disseminate to all virtual machines, vice having to buy multiple or unnecessary enterprise copies and load them on all physical stations. By expanding joint mission capability, reducing the amount of physical hardware and manpower required, and improving security, virtualization may be worth the initial cost in order to save for the future. This will allow for the USN and USMC to rapidly adapt and adjust to ever changing computing and mission requirements. By improving both the range and time agility of the organizations, they become more capable to adapt and overcome at a more affordable and justifiable cost.

As mentioned Navy and Marine differences both distinguish and detract. The use of VDI, and its ability to virtualize numerous desktops, would take the USN and USMC from a department that reflects short-sited diversification to one that reflects modern coordination. Virtualization would allow for organizations to integrate customer services while still allowing the different groups to maintain their unique operational requirements. This ability opens up further possibilities to work with groups outside the DoD and expands the mission capabilities of the USN and USMC to more than just war fighting.

THIS PAGE INTENTIONALLY LEFT BLANK

III. SOLID STATE MEMORY

A. GENERAL DISCUSSION

As discussed in the previous chapter, virtualization and the establishment of sea-based VDI can provide many benefits to any organization afloat. However, this technology is only one piece of the puzzle in implementing a successful solution. The utilization of VDI systems with modifying or replacing typical IT-21 server farms presents many inherent challenges. Two of these challenges stand out above others. They are:

1. Performance. The I/O constraints of disk-based systems result in poor performance. If a virtual desktop fails to perform as well or better than a physical desktop, end users complain and productivity suffers. Unfortunately, delivering the performance most enterprises need requires racks of hardware that quickly make VDI impractical (Fusion io, 2011, p. 1).

2. High and Unpredictable Costs. The initial outlay to purchase NAS or SAN systems is enormous. On top of this, system administration, power, cooling, and colocation fees become an ever-increasing operating expense.

To fully realize the potential benefits provided by a virtual desktop infrastructure, storage considerations along with cost must be weighed, determined and implemented effectively. According to Siebert (2011), "implementing a virtual desktop infrastructure (VDI) involves many critical considerations, but storage may be the most vital" (Siebert, 2011).

The most difficult and troublesome challenge faced when setting up a virtual desktop infrastructure environment "is accommodating the periods of peak usage when storage I/O, input/output, is at its highest" (Siebert, 2011). Input/output refers to "any program, operation or device that transfers data to or from a computer and to or from a peripheral device. Every transfer is an output from one device and an input into another" (Siebert, 2011). The largest periods of storage I/O issues when running a VDI exist during "boot storms," or when large numbers of users boot up applications or operating systems at the same time. “Initial startup of a desktop is very resource-intensive activity with the operating system and application doing a lot of reading from disk” (Siebert, 2011). Once a storm has been started it may continue for minutes or last up to hours before it stops. During the storm, the VDI cannot and will not accept any inputs and cannot produce any outputs until the storm has cleared. This leaves the end users without a fully operational computer system for possibly hours.

The same issue can occur during peak usage times and during shut-down storms. “Events like patching desktops, antivirus updates/scans” and heavy applications workload can all lead to storage I/O issues (Siebert, 2011). To deal with these issues, it is critical to determine the data storage infrastructure required to handle peak period challenges.

When setting up a VDI the “the key measurement for storage is IOPS,” or the amount of input and output operations per second (Siebert, 2011). With the use of virtualization comes the creation of “a highly random I/O pattern from each host. Adding additional VMs to a host results in an exponential increase in random I/O operations per second. The result is that every server in the environment is a potential IOPS consumer meaning that storage system IOPS is more important than air” (Crump, 2010, p. 1). With the creation of a handful of VMs there may only be a moderate amount of resources consumed. As stated before, when many VMs are created and are in use, the amount of I/Os in use grows exponentially. Using the VMware rule of thumb that a VM uses around fifty IOPS per VM, we can determine that a host providing fifteen VMs will need around 750 IOPS. Based upon Trident Warrior experiments and observation, the number of IOPS required per VM is closer to 150 IOPS. This means that the same fifteen VM infrastructures would produce closer to 2,250 IOPS. If an organization were to run 25 VMs, that would bring the amount of IOPS to 3,750. A medium-sized virtual infrastructure of 50 VMs would reach 187,000 IOPS for the virtual environment (Crump, 2010, p. 2).

The amount of IOPS for a VDI should be estimated, as it may be impossible to actually determine, prior to its implementation. Since IOPS are “*sustained* levels of I/O demand and I/O capacity, the average usage of a system needs to be considered.” The average I/O loads, however, should not be the benchmark in which a VDI’s storage needs are based upon. Instead the peak I/O loads based upon the amount of expected usage and the amount of VMs required must be considered and factored into the development of a VDI's storage needs in order to provide the best user experience.

There are other considerations for the selection of the type of storage that is required when setting up a VDI. Performance parameters such as read and write capabilities, and spin rates should be considered. Cost is another factor that must be considered as well and includes factors such as acquisition, maintenance, and overhead costs to operate.

This leads to the type of storage possibilities. The two main types of storage drives, spinning disk drives, typically made up of SAS and SATA, and solid state drives. There are pros and cons associated with each of these types of mediums and a direct comparison is provided in Table 1 (Wikipedia, 2012).

Attribute or characteristic	**Solid-state drive**	**Spinning disk drive**
Spin-up time	Almost instantaneous; nothing mechanical to "spin up." May need a few milliseconds to come out of an automatic power-saving mode.	May take several seconds. With a large number of drives, spin-up may need to be staggered to limit total power drawn.
Data transfer rate	SSD technology can deliver rather consistent read/write speed, typically ranging from about 100MB/s to 500MB/s, depending on model. When accessing individual smaller blocks, the performance will usually degrade from that. In general, the speeds are continuously improving.	When reading or writing a continuous track, a HDD can access data roughly at speed of 100MB/s. However, due to need of seeking, the actual transfer rates will almost always be much lower.

Attribute or characteristic	Solid-state drive	Spinning disk drive
Random access time[57]	About 0.1 ms - many times faster than HDDs because data is accessed directly from the flash memory	Ranges from 5–10 ms due to the need to move the heads and wait for the data to rotate under the read/write head.
Read latency time[59]	Generally low because the data can be read directly from any location; In applications where hard disk seeks are the limiting factor, this results in faster boot and application launch times (see Amdahl's law).	Generally high since the mechanical components require additional time to get aligned
Consistent read performance[61]	Read performance does not change based on where data is stored on an SSD	If data is written in a fragmented way, reading back the data will have varying response times
Fragmentation	There is usually very little benefit to reading data sequentially (beyond typical FS block sizes), making fragmentation a void issue for SSDs. Defragmentation process also makes additional writes on the NAND flash cells that already have a limited cycle life. It is also uncertain whether defragmentation would arrange the data in a truly sequential order, as the drive itself can again remap it to various positions.	File systems on HDDs may fragment after continued operations of erasing and writing data, especially involving large files. Therefore, periodical defragmentation is required to maintain ultimate performance.
Acoustic levels	SSDs have no moving parts and make no sound	HDDs have moving parts (heads, spindle motor) and have varying levels of sound depending upon model

Attribute or characteristic	Solid-state drive	Spinning disk drive
Mechanical reliability	A lack of moving parts virtually eliminates mechanical breakdowns	HDDs have many moving parts that are all subject to failure over time
Maintenance of temperature	SSDs do not usually require any cooling maintenance and they can tolerate higher temperatures than HDDs. High-end enterprise models delivered as add-on cards may include heat sinks to dissipate heat generated by its chips.	Air-forced ventilation is recommended for desktop hard drives to avoid build-up of heat.[66] Otherwise, bad sectors on its media can appear later and/or its lifespan will diminish over time. HDDs designed for laptops do not need as much cooling, but heat issues are a matter of concern with them too.
Susceptibility to environmental factors	No flying heads or rotating platters to fail as a result of shock, altitude, or vibration	The flying heads and rotating platters are generally susceptible to shock, altitude, and vibration
Installation and mounting	As for SSD, as long as it's mounted securely to its place, the position and installation mechanism do not have much of impact to its normal use. Most ordinary SSDs have all components (except for power and data connectors) encased inside.	While installing a hard disk drive, one must take care of sufficient cooling and sturdy mounting. Additionally accessories to dampen vibration, noise and mechanical shocks, can be installed. The printed circuit board underneath of a HDD is usually exposed and any conductive material can not be let to short-circuit the components or electronic contact points.
Magnetic susceptibility	No impact on flash memory	Magnets or magnetic surges can alter data on the media
Weight and size	The weight of flash memory and the circuit board material are very light compared to HDDs	Higher performing HDDs require heavier components than laptop HDDs (which are light, but not as light as SSDs)

Attribute or characteristic	Solid-state drive	Spinning disk drive
Parallel operation	Some flash controllers can have multiple flash chips reading and writing different data simultaneously	HDDs have multiple heads (one per platter) but they are connected, and share one positioning motor.
Write longevity	Flash-based SSDs have a limited number of writes (1–5 million or more) over the life of the drive. Software controllers manage this limitation in such a way that drives can last for many decades before failure. SSDs based on DRAM do not have a limited number of writes.	Magnetic media do not have a similar limited number of writes but are susceptible to eventual mechanical failure.
Secure writing limitations	NAND flash memory cannot be overwritten, but has to be rewritten to previously erased blocks. If a software encryption program encrypts data already on the SSD, the overwritten data is still unsecured, unencrypted, and accessible (drive-based hardware encryption does not have this problem). Also data cannot be securely erased by overwriting the original file without special “Secure Erase” procedures built into the drive.	HDDs can overwrite data directly on the drive in any particular sector.
Cost per capacity	As of February 2011, NAND flash SSDs cost about (U.S.)$.90–2.00 per GB	As of February 2011, HDDs cost about (U.S.)$0.05/GB for 3.5 in and $0.10/GB for 2.5 in drives

Attribute or characteristic	**Solid-state drive**	**Spinning disk drive**
Storage capacity	As of December 2011, SSDs come in different sizes up to 2TB but are typically not larger than 64–256GB, due to their high cost per capacity.	As of December 2011, HDDs are typically up to 1TB in capacity but drives up to 4TB are available.
Read/write performance symmetry	Less expensive SSDs typically have write speeds significantly lower than their read speeds. Higher performing SSDs have a balanced read and write speed.	HDDs generally have slightly lower write speeds than their read speeds.
Free block availability and TRIM	SSD write performance is significantly impacted by the availability of free, programmable blocks. Previously written data blocks that are no longer in use can be reclaimed by TRIM; however, even with TRIM, fewer free, programmable blocks translates into reduced performance.	HDDs are not affected by free blocks or the operation (or lack) of the TRIM command
Power consumption	High performance flash-based SSDs generally require 1/2 to 1/3 the power of HDDs; High performance DRAM SSDs generally require as much power as HDDs and consume power when the rest of the system is shut down.	High performance HDDs generally require between 12–18 watts; drives designed for notebook computers are typically 2 watts.

Table 1. Solid State Memory Drives vs. Spinning Disk Memory Drives (From Wikipedia, 2012)

As shown in Table 1, spinning disk drives are much more affordable than their solid-state counterparts, in terms of initial acquisition costs. Solid state drives, however, provide far more performance and far less cost once acquired due to the reduction in the amount of physical equipment required. This is highlighted in Figure 4, the Fusion-io comparison of a SAN based VDI versus a Fusion-io SSD system (Fusion-io, 2011):

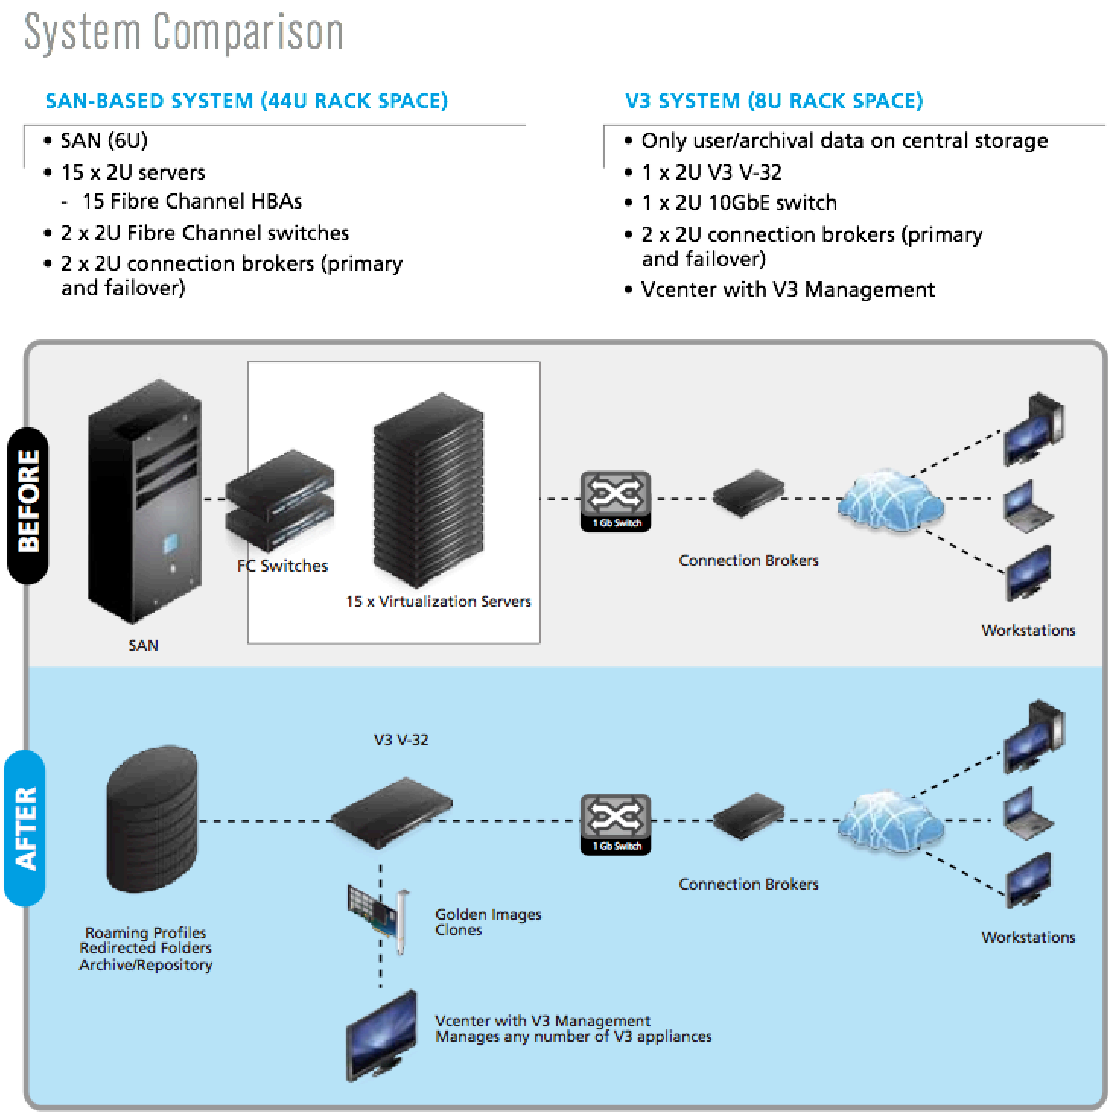

Figure 4. SAN Based System and SSD Based System Comparison (From Fusion-io, 2011)

Additionally, the performance of a spinning disk drive does not scale linearly like that of a SSD. To increase read and write latency, a SSD has to spin faster. The head actuator that moves across the disk does not move faster though. So by spinning a drive 50% faster, the net performance increase is only 30% (Siebert, 2011). This does result in higher IOPS, but at a cost. Having to spin the disk at such a high rate for modest returns will lead to higher than normal wear on the disk, which will require an earlier replacement. SSDs are linear in their performance scalability, which makes them the better candidate for an organization seeking to know exactly the performance and capability they will receive for the price.

Organizations, like the U.S. Navy, are looking to get the most bang for their buck while reducing overhead. In the ever more stringent fiscal environment, the U.S. Navy and others need to be able to track exactly where their expenditures are going. SSDs may provide organizations the performance they want from their VDIs. At the same time all will need to lower overall costs through a reduction in physical infrastructure. SSDs make financial sense when compared to their spinning-disk counterpart.

B. HISTORY

SSD drives first made an appearance in the 1950s, but were short lived as the introduction of the cheaper drum storage units replaced them. They were once again used in the 1970s and 1980s in conjunction with IBM, Amdahl and Cray super computers. Their extremely high costs meant that they were rarely used and were once again shelved in favor of cheaper, more readily available technologies. SSDs have sparsely appeared throughout the years since. All have been unable to take hold until recently, where more and more SSD developers have finally found a foothold in the marketplace (Wikipedia, 2012).

C. THE STATE OF SOLID STATE MEMORY

Again, the concept of SSDs and its benefits are not new. Just now, though, the technology is reemerging on the marketplace in quantities and prices that make it a viable option for organizations. The performance of a SSD versus that of a spinning disk is not exactly one to one which made it hard for organizations to cross examine the two technologies. Since they are not one to one and the cost of spinning disks tends to be cheaper, organizations have had an adverse outlook on SSDs.

When looking at the market now, SSDs still make up only a small percentage of storage units sold. That is not to say that market is not growing. More and more developers have appeared in recent years than there has ever been in the SSD market. The increased numbers of developers are a signal that there is a market for the product as the numerous amounts of vendors must be selling their products to someone. Corporations such as Fusion-io, the largest SSD developer, are trading publicly on the stock market with more SSD companies to follow.

D. ADVANTAGES OF SOLID STATE MEMORY

The benefits of SSDs for VDIs are significant according to Fusion-io, and include the following advantages:

- The ability to support a greater number of virtual desktops per virtual server than spinning media. (Fusion-io, 2011)
- Can guarantee performance even during peak loads due to the ability to sustain high IOPS rates. (Fusion-io, 2011)
- Provides low latency for faster end user desktop experiences. (Fusion-io, 2011)
- Will make VDI cost-effective enough to localize servers and eliminate network latency. (Fusion-io, 2011)
- Reduce reliance on expensive, complex external storage. (Fusion-io, 2011)
- Will lower the overall cost per desktop while providing higher performance. (Fusion-io, 2011)

E. SOLID STATE MEMORY AND THE USN/USMC

As the Navy and Marine Corps looks to expand joint proficiencies through improved computing capabilities, these capabilities must prove to perform better than what is currently in use, must drastically expand mission capabilities, and do so for a cost that is reasonable or justifiable. The Department of Defense is loath to throw money at technology that only provides singular mission capabilities or limited improvements. For new technology to be considered for acquisition by the DoD, it must offer multiple mission capability improvements. VDI in conjunction with solid state drives can provide Naval and Marine units improved computing capabilities that open up a world of possibilities in terms of joint mission tasking. Alone, solid state drives provide many performance improvements in terms of computing but also provide many fringe benefits that are of value to Naval and Marine units that makes them a compelling new technology that is worth investing in.

In terms of performance improvements, solid state drives provide many compared to spinning disks. SSDs provide faster spin up times, faster data transfer rates, quicker random access times, and improved read latency times. When utilized within a VDI, these improvements quickly add up to real world capabilities. Virtual desktops using products like Fusion-io drives can provide the end user with a desktop that rivals high end desktops. Figure 5 compares the times for a virtual desktop using Fusion-io's solid state drives to complete common computing tasks compared to a local desktop (Fusion-io, 2011).

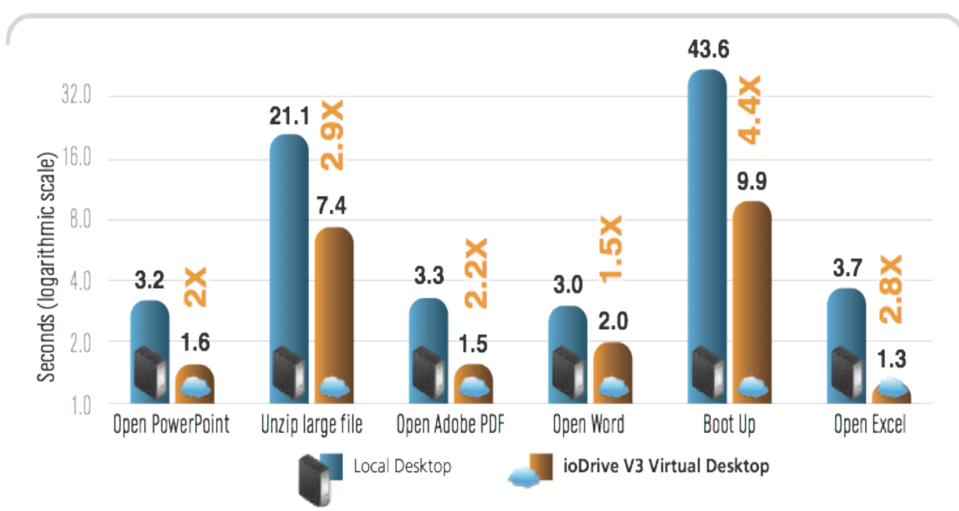

Figure 5. Virtual Desktop Utilizing SSD vs. Local Desktop (Fusion-io, 2011)

The performance improvements are clear when put side by side. These desktops can be instantly created or removed as needs arise providing amphibious vessels additional flexibility in joint missions. This allows for disparate organizations and personnel to embark Naval Amphibious vessels and be given high quality desktop without the expense or overhead of local desktops. Also, considering a VDI can support tens to hundreds of virtual desktops, the cost savings are clear.

Solid state drives not only improve the capabilities of an amphibious vessel's VDI but also provide many other improvements in terms of cost and performance. Solid state drives require far less power than spinning disk drives. "High performance flash-based SSDs generally require 1/2 to 1/3 the power of spinning disk drives" (Wikipedia, 2012). Solid state drives also require far less hardware which means that far less power is used to run the system and to run support equipment such as cooling systems. Any reduction in power consumption onboard naval vessels is of a high value as power at sea equals fuel usage.

Solid state drives, as their name implies, do not have any moving parts. Without moving parts to break, mechanical reliability increases. This leads to costs savings in both re-acquisition and as well as in man-hours required to fix or replace broken drives. This allows manpower to be reduced or reassigned to where it is needed which improves overall mission capability.

With no moving parts, solid state drives have far lower acoustic levels. Although the amount of sound emitted by spinning drives may negligible in real world applications, any reduction in sound levels of equipment aboard ships is applicable. Sound emitted from equipment on board naval vessels is cumulative and all adds to the overall sound signature. Any reduction in sound levels of equipment aboard a Naval vessel will lead to the overall reduction in the sound signature of the vessel.

Another concern in the naval environment is shock and vibration. Naval vessels are susceptible to shock and vibration from equipment onboard, sea states, and in some cases battle damage. Due to the lack of moving parts in solid state drives, they are far less, if at all, susceptible to vibration or shock damage caused by detonations and explosions. This is another factor that improves the reliability of an amphibious vessel's VDI which in turn reduces cost and man-power requirements.

IV. AMPHIBIOUS NETWORK INTEGRATION ANALYSIS

A. OVERVIEW

Initial observations of this LAN integration process began over four years ago with the mentioned 15th MEU re-embarkation, where inorganic workstation integration was witnessed first-hand. Those experienced practices were bolstered through additional observations made from later interactions with multiple Expeditionary Strike Groups (ESGs) and Amphibious Readiness Group (ARGs). Most recently, a draft of the AS-IS architecture was refined and confirmed to be accurate onboard USS WASP (LHD-1) during Trident Warrior 11.

The purpose of this chapter is to apply observations toward a notional analysis of the LHD workstation integration process in effort to conceptually identify areas of improvement and form feasible technology and architecture recommendations. The approach is guided from NPS IS4220 and IS4250 course discussions based upon an Enterprise Architecture Integration Case Study, in-class exercises, and individual product examples provided from the instructor. The following structure is leveraged for the model:

1. Define current view of LHD LAN Integration
2. Define core integration processes and organizing logic
3. Current Processes
4. Generalizations for Modeling

1. Current View Defined

The AS-IS Enterprise Architecture used onboard amphibious flagships has grown overwhelmingly with noted compliance policies. Functional personnel supporting network integration are separately tasked and functional-to-functional interaction has evolved into bottlenecking reciprocal relationships. Individual workstations must separately travel to functional areas, each with tedious and time-consuming compliancy requirements. For example, a computer must be delivered to one workspace for hardware compliance validation, another for operating system verification, and yet another for security update confirmations. In practically every instance, each workstation will not meet any Navy afloat standards, requiring a complete hard drive slick, O/S re-imaging, and comprehensive security patch update conducted. Additional areas for improvement within the current architecture are:

- Strict reliance on limited IT21 afloat approved hardware
- Inflexibility on specific version of O/S and ship network version image
- Requisite lengthy security scanning and updating durations
- Separation of functional area assets and capabilities
- Physical change of custody of workstation to ship
- Waiting time of functional areas in reciprocal workstation processing

Overall, amphibious ships' IT crew continuously integrate MEU, squadron, staff, and individual workstations requesting network services for underway and destination use. The functional personnel and processes dogma is so ingrained that customers/embarkees become lulled into an acceptance of high military standard compliance, even in the context of unclassified or publically accessible work. In situations where rapid action, crisis response, or disaster relief call for immediate integration and critical planning, the present architecture becomes unacceptable and core processes are called into high visibility with direct questioning. Occasionally, what was once mandatory becomes compromised and temporarily modified to fit tighter operational timelines. Through identifying and defining what exactly is core to workstation integration from a ship-hosting enterprise network perspective, an understanding of process integration and standardization contributing to overall improvement becomes visible.

2. Core Processes and Organizing Logic Defined

Observed network integration practices were captured through Business Process Modeling techniques. After an initial introduction to functional personnel, the workstation integration process is described.

Stated observation and fleet verification identified five groups of core personnel supporting network integration: Joint Network Operations Center (JNOC), Data System

Repair (DSR), Automated Data Processing (ADP), Information Assurance (IA), and USMC Information Technology (IT) Quality Assurance (QA). Onboard amphibious command ships, JNOC members serve roles as customer support, receiving requests and managing their resolution, and System Administrators (SAs) to the network servers. DSR personnel are responsible for the operation of the network infrastructure, consisting of the physical connections of servers through backbone and edge switches and finally drops within the space. Also, DSR provides some limited capability for trouble-shooting and repairing IT hardware. Next, the IA team ensures all onboard IT hardware and software configuration and operation meet DoD, Navy, and ship security policy. Finally, USMC QA represents the embarking organization's IT personnel holding comparable training and knowledge to their previously mentioned Navy counterpart groups; however, embarkees typically focus on the end result, an operable workstation on the provided network.

3. Current Process

As experienced in the Fleet, our AS-IS process model can be best described as:

a) JNOC receives a request to integrate a quantified number of workstations onto the ship's network. Servers are checked to verify if the capacity exists to host the requested user profiles and manage the connected workstations. Overall count is determined and DSR validates if the network infrastructure can support the workstations at the requested/assigned ship space. When the server and infrastructure capacities are negotiated, JNOC is able to take inventory, often transferring possession to themselves, of the workstations. This process is illustrated in Figure 6.

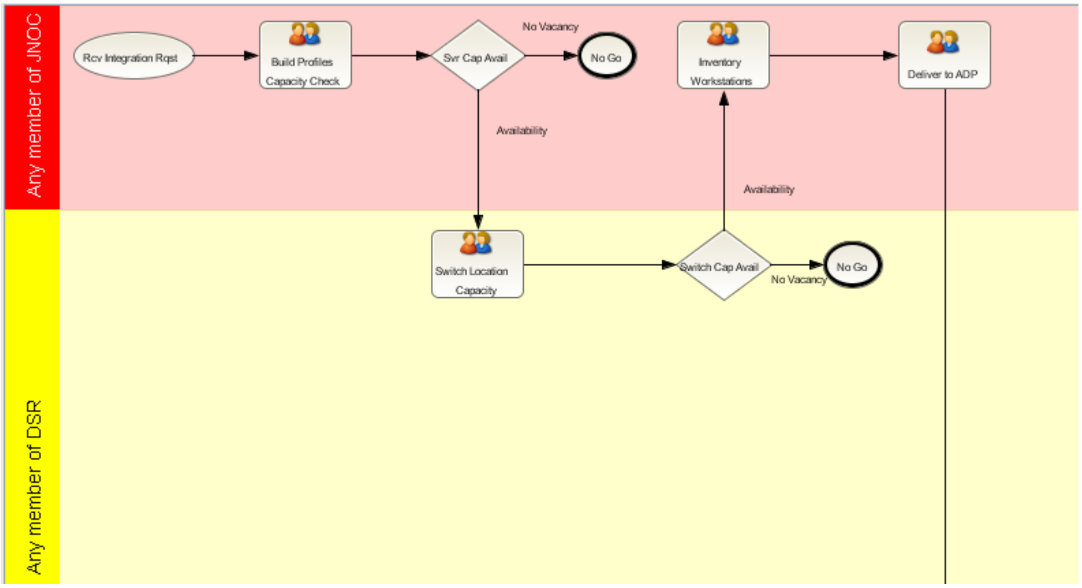

Figure 6. Amphibious LAN Integration Step a)

b) ADP initially confirms that the workstation is functional. If not, either the embarking IT personnel or DSR have an opportunity to trouble-shoot and fix the computer. If operational, the image is checked to ensure it complies with the exact shipboard network image. In nearly all cases policy requires the workstation to be slicked, re-imaged, and loaded with the most recent version of applicable Operating System (O/S). The IA team then loads each workstation with the appropriate software and patches to ensure overall security compliance. This process is illustrated in Figure 7.

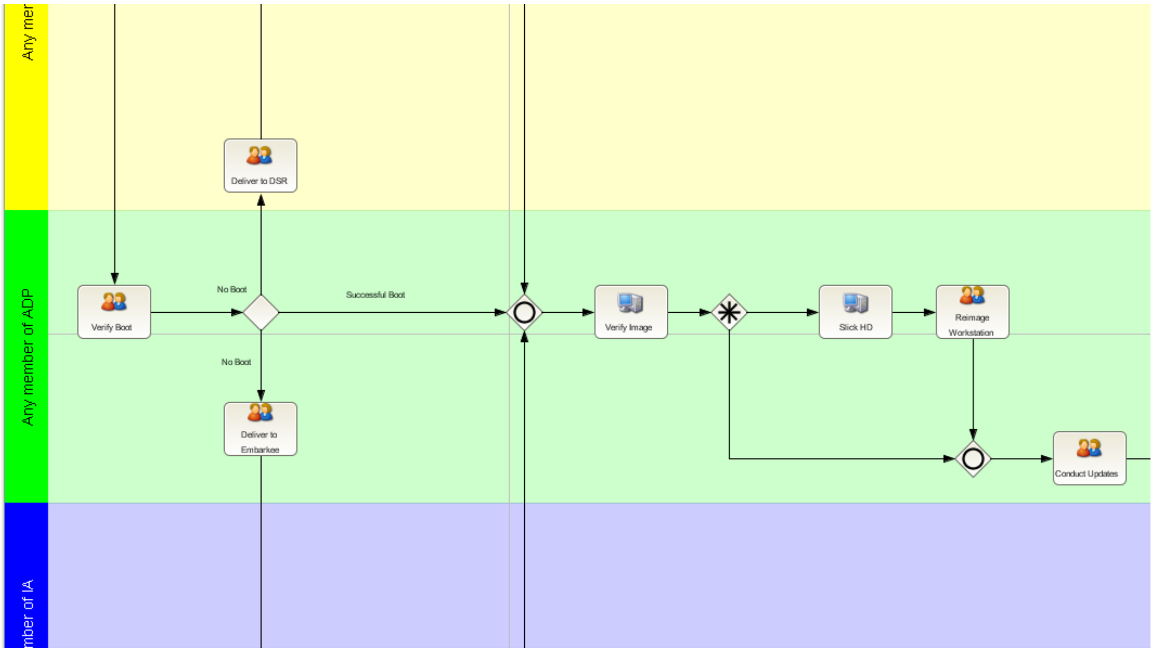

Figure 7. AS-IS Amphibious LAN Integration Step b)

c) JNOC delivers the workstation to its assigned space and DSR enables the specific port on the switch that is wired to the drop. Occasionally, the port and drop link is not initially operational and requires some trouble-shooting by DSR to achieve network connectivity. This process is illustrated in Figure 8.

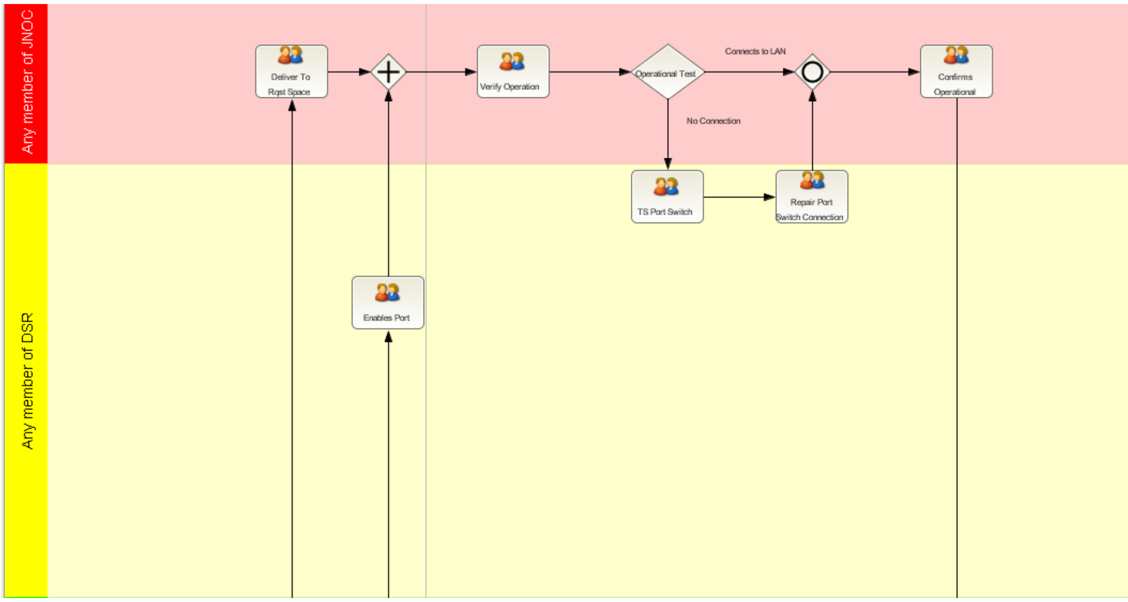

Figure 8. AS-IS Amphibious LAN Integration Step c)

d) Finally, USMC QA performs their overall functionality check on each workstation with an appropriate user where integration is ultimately confirmed.

4. Generalizations for Modeling

To design and model the analyzed integration process, the following generalizations were required:

a. Sub-Process Duration

Each modeled activity consisted of estimated times to complete. Again, previous tours experience managing the actual sub-processes allowed averages to be applied. Activity durations were further reviewed and agreed upon by current WASP IT leadership, as well.

b. Decision-Point Probability

At distinct points throughout the process, probabilities had to be incorporated to represent various realities to the modeled environment. Thresholds to servers and switch capacity often do not reach full visibility until each embark is requested. Additionally, equipment onboard and brought from various locations often is transported and operates in conditions exceeding manufacturers' recommendations. As a result, workstations/IT equipment and their components fail before their anticipated end of life cycle.

c. Volume of Integration

Embarkations may vary from individuals, to staffs, and Marine Expeditionary Units (MEUs). Designed to embark MEUs, amphibious ship networks and their supporting processes are only flexed when a full complement of Marines are onboard. Our current model is best served representing an integration of 300 workstations typically requested by the MEU for the ship's UNCLAS network.

d. Associated Costs

The current MEU UNCLAS workstation integration process is only one of a multitude while beginning an underway period, transit, of a deployment. The litany of preparations to get underway has both Sailors and Marines pulled in several directions. As observed, two members from each group are able to support this process during a workday. Each is required to have the capabilities expected of an E-5 in their IT specialty. Our cost is measured based upon this analysis. The AS-IS amphibious LAN integration metrics, including cost and number of people, developed by the authors are depicted in Table 2.

Simulation Results					
Duration	121:00:15 Time				

Process Time And Cost

Process	Scenario	Instances	Total Cost	Waiting Time (Time)	Total Time (Time)
Matthias_Becker_GruberV1_V1_0	(default)	300	13504.6	22548:51:30	15428:04:45

Resource	Unit	Cost/Unit	# People	Utilization
Any member of ADP	Hour	16.66	2.00	264.60%
Any member of DSR	Hour	16.66	2.00	430.60%
Any member of IA	Hour	16.66	2.00	229.96%
Any member of JNOC	Hour	16.66	2.00	210.92%
Any member of USMC QA	Hour	16.66	2.00	79.69%

Table 2. AS-IS Amphibious LAN Integration Metrics

Running several simulations, the results given in Table 2 were found to be typical. The results shown in Table 2 realistically portrayed outcomes in alignment with our experiences and several Navy and Marine IT support staff input. Total duration of 300 MEU NIPR workstations integration lasted over 15 days, accurately reflecting mentioned embarkations. Also, acceptable utilization of personnel should be less than 80% of their workday, calculating in other tasking and breaks. The business processes used onboard LHD's causes IT staff, characteristic to this platform, to be over-worked during MEU embarks upwards of 5 times what is considered reasonable utilization.

The cursory knowledge of the core personnel and current business processes of network integration allows for a foundation of organizing logic. Process integration and standardization are then prioritized to provide requirements for a TO-BE EA. The AS-IS model clearly enforces high standardization and low business process integration. Both hardware and software must be manipulated to ensure the standard infrastructure integrity is upheld. High replication practices left no room for coordination efforts of nearly every other military and government entity besides, specifically, a Marine Expeditionary-budgeted force with equipment expressed designed for the platform they are embarking.

Given today's broad Range of Military Operations (ROMO), the current organizing logic becomes unsuitable. Rather, a balanced combination of both Coordination and Replication are necessary. Leveraging the two, outside disparate organizations need to become unified beginning with the shipboard LAN/enterprise, into a single business, to support the combined, collaborative operational and decision-making processes demanded from today's challenges. Although technology alone should never be an absolute solution, it can be successfully partnered with some process re-engineering. The targeted architecture involves an upward movement toward higher process integration to obtain unification.

THIS PAGE INTENTIONALLY LEFT BLANK

V. TRIDENT WARRIOR 2011 EXPERIMENTATION

A. TRIDENT WARRIOR DESCRIPTION

Gaps in operational needs and the realities of enterprise planning are not held within dialogues between afloat commanders and their onboard captive audiences. Feedback from deployed Strike Group and ARG lessons learned are routinely captured, reviewed, and forwarded through their respective chains of command to a higher, ultimately, Component Command. Component Commands, in turn, are tasked with anticipating and appropriately manning, training, and equipping their operational units to effectively meet desired Mission Essential Tasks (METs) across multiple warfare domains. Obtaining operational responsiveness and relevance requires venues supporting near-term and developmental inquiries into improving each aspect of manning levels, personnel, team, and command training, and technology available. One means for addressing voiced challenges or closing gaps is through experimentation onboard within an operationally simulated environment.

Recently, U. S. Fleet Forces Command (FFC), a Component Commander for the Navy's Atlantic theater, directed TW11 to be conducted primarily within 2nd Fleet while coordinating efforts and sponsorship with 5th Fleet. "Trident Warrior (TW) is an annual fleet experiment designed to improve war fighting policies and capabilities by providing answers to detailed analytical questions about more than 50+ critical maritime initiatives included in the experiment's execution" (FFC TW, 2011). TW team members are highly skilled in structuring, executing, and analyzing research initiatives requested from the deckplates, addressing urgent needs, to enterprise inquiries on future capabilities. Notably, TW is able to create initiatives for ship's company personnel to be integral research participants within interactive experimentation and data collection. TW final reports warrant the highest visibility and effectively grant a peek at the possible for decisions made impacting the next Strike Group's deployment or shaping strategies.

From August until September 2011, the leadership and bulk of the TW11 team embarked aboard U.S.S. WASP (LHD-1) to conduct experiments spanning the full range of warfare areas. The whole of their assigned initiatives incorporated emerging technologies to enhance the processes and operations that maritime forces currently practice. Additionally, the Trident Warrior process has the ability to introduce capabilities that could develop new Tactics, Techniques, and Procedures (TTPs) for the U.S. Navy. Rallying on FORCENet Innovation and Research Enterprise's (FIRE) portal and database resources, TW11 organized their experimental objectives, resulting planning, assessments, respective analysis, and results. Each critical initiative was grouped into seven focus areas: Command and Control (C2), Electronic Warfare (EW)/Fires, Information Operations (IO), Information Technology (IT), Information Surveillance and Reconnaissance (ISR), Missile Defense, and Networks (NET).

B TRIDENT WARRIOR THREAD NET 11.01

One Networks focus area thread, identified as NET 11.01, directly aligned with the purpose of this thesis. Titled as Next Generation Technologies (NGT) Virtual Infrastructure for embarkables, NET 11.01's single objective question asked, "Can [a] NGT virtual desktop infrastructure (VDI) client serve as an intermediary to interface NMCI workstation client to [an] IT-21 network to provide IT-21 applications and services (e.g., IT-21 patches) for an embarked NMCI machine (capable/usable)?" (FFC TW, 2011). In short, the experimentation thread inquired if VDI could provide a bridge from a hosting ship's network to embarking workstations. The ship's network and operating system, IT-21 and COMPOSE (Common PC Operating System Environment) respectively, and the simulated embarking NMCI workstations all represent the most stringent requirements for network integration. As an objective for experimentation, TW11 sought to "demonstrate NGT's ability to provide a variety of improved network services" (FFC TW, 2011).

C. NET 11.01 TESTING

Of note, the Network's thread specifically keyed more into VDI technology functional capabilities than the processes supporting integration. The preceding and following chapters provide in depth observations and recommendations regarding workstation integration processes. Onboard WASP, NET 11.01's intent was to test the underway performance of NGT with a nominal IT infrastructure, or Common Computing Environment (CCE), while conducting desktop virtualization. Simulating the nominal and scalable infrastructure, TW technical representatives mounted an IBM HS-22 blade server within WASP's JNOC opposed to the standard HS-21 blade server, typical to the location of IT-21 servers on LHDs. For reference, the specifications and structure of HS-22 servers are illustrated in Figures 9 and 10 (IBM, 2011, p. 8 and 13). The specifications and structure of HS-21servers are illustrated in Figures 11 and 12 (IBM, 2007, p. 6 and 10).

Microprocessor: Supports up to two multi-core Intel Xeon microprocessors.
Note: Use the Setup utility to determine the type and speed of the microprocessors in the blade server.

Memory:
- 12 dual inline memory module (DIMM) connectors
- Type: Very Low Profile (VLP) double-data rate (DDR3) DRAM. Supports 1 GB, 2 GB, 4 GB, 8 GB, and 16 GB DIMMs with up to 192 GB of total memory on the system board

Integrated functions:
- Horizontal-compact-form-factor (CFFh) expansion card interface
- Vertical-combination-I/O (CIOv) expansion card interface
- Local service processor: Integrated Management Module (IMM) with Intelligent Platform Management Interface (IPMI) firmware
- Integrated Matrox G200eV video controller
- LSI 1064E SAS controller
- Broadcom BCM5709S dual-port Gigabit Ethernet controller
- Integrated keyboard/video/mouse (cKVM) controller through IMM
- Light path diagnostics
- RS-485 interface for communication with the management module
- Automatic server restart (ASR)
- USB 2.0 for communication with cKVM and removable media drives (an external USB port is not supported)
- Serial over LAN (SOL)
- Redundant buses for communication with keyboard, mouse, and removable media drives

Predictive Failure Analysis (PFA) alerts:
- Microprocessors
- Memory
- Storage drives

Electrical input: 12 V dc

Environment:
- Air temperature:
  - Blade server on: 10°C to 35°C (50°F to 95°F). Altitude: 0 m to 914.4 m (0 ft to 3000 ft)
  - Blade server on: 10°C to 32°C (50°F to 89.6°F). Altitude: 914.4 m to 2133.6 m (3000 ft to 7000 ft)
  - Blade server off: 10°C to 43°C (50°F to 109.4°F). Altitude: 914.4 m to 2133.6 m (3000 ft to 7000 ft)
  - Blade server shipping: -40°C to 60°C (-40°F to 140°F)
- Humidity:
  - Blade server on: 8% to 80%
  - Blade server off: 8% to 80%
  - Blade server storage: 5% to 80%
  - Blade server shipment: 5% to 100%
- Particulate contamination:

  Attention: Airborne particulates and reactive gases acting alone or in combination with other environmental factors such as humidity or temperature might pose a risk to the server. For information about the limits for particulates and gases, see "Particulate contamination" on page 85.

Drives: Supports up to two hot-swap, small form factor (SFF) Serial Attached SCSI (SAS) or Serial ATA (SATA) storage drives

Size (Type 7870 and Type 1936):
- Height: 24.5 cm (9.7 inches) (6U)
- Depth: 44.6 cm (17.6 inches)
- Width: 2.9 cm (1.14 inches)
- Maximum weight: 4.8 kg (10 lb)

Size (Type 1911):
- Height: 24.5 cm (9.7 in)
- Depth: 44.6 cm (17.6 in)
- Width: 14.5 cm (5.71 in)
- Maximum weight: 8.15 kg (40.02 lb)

NEBS Environment
- Air temperature:
  - Blade server on: 5°C to 40°C (41°F to 104°F). Altitude: -60 m to 1800 m (-197 ft to 6000 ft)
  - Blade server on: 5°C to 30°C (41°F to 86°F). Altitude: 1800 m to 4000 m (6000 ft to 13000 ft)
  - Blade server off: -5°C to 55°C (23°F to 131°F). Altitude: -60 m to 1800 m (-197 ft to 6000 ft)
  - Blade server off: -5°C to 45°C (23°F to 113°F). Altitude: 1800 m to 4000 m (6000 ft to 13000 ft)
  - Blade server storage: -40°C to 60°C (-40°F to 140°F)
- Humidity: 8% to 85%
- Particulate contamination:

  Attention: Airborne particulates and reactive gases acting alone or in combination with other environmental factors such as humidity or temperature might pose a risk to the server. For information about the limits for particulates and gases, see "Particulate contamination" on page 85.

Figure 9. HS-22 Blade Server Specifications (From IBM, 2011, p.

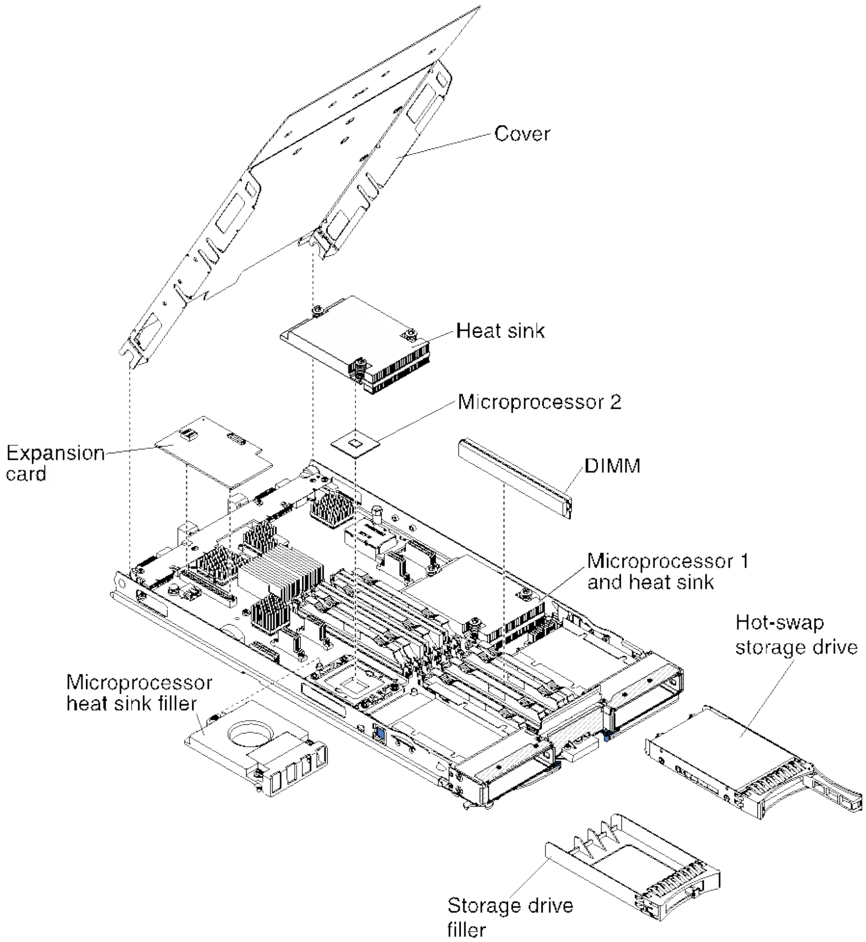

Figure 10. HS-22 Major Components (From IBM, 2011, p. 13)

Microprocessor: Supports up to two dual- or quad-core Intel® Xeon microprocessors.
Note: Use the Configuration/Setup Utility program to determine the type and speed of the microprocessors in your blade server.

Memory:
- Dual-channel DIMMs: 4 DIMM slots
- Type: fully-buffered double-data rate (FB-DDR2), PC2-5300, ECC SDRAM registered x4 (Chipkill) DIMMs
- Supports 512 MB, 1 GB, 2 GB, and 4 GB DIMMs (as of the date of this publication) with up to 16 GB of total memory in the system board
- Additional memory support when an optional IBM BladeCenter Memory and I/O Expansion Blade is installed

Drives: Support for up to two internal small-form-factor Serial Attached SCSI (SAS) drives

Predictive Failure Analysis® (PFA) alerts:
- Microprocessor
- Memory
- Hard disk drives

Electrical input: 12 V dc

Integrated functions:
- Dual Gigabit Ethernet controllers
- Expansion card interface
- Local service processor: Baseboard management controller (BMC) with Intelligent Platform Management Interface (IPMI) firmware
- ATI RN-50 ES1000 video controller
- LSI 1064E Serial Attached SCSI (SAS) controller
- Light path diagnostics
- RS-485 interface for communication with the management module
- Automatic server restart (ASR)
- Serial over LAN (SOL)
- Redundant buses for communication with keyboard, mouse, and removable media drives
- Concurrent keyboard/video/mouse (cKVM) support when optional cKVM feature card is installed

Environment (non-NEBS):
- Air temperature:
  – Blade server on: 10° to 35° C (50° to 95° F). Altitude: 0 to 914 m (0 to 3000 ft)
  – Blade server on: 10° to 32° C (50° to 90° F). Altitude: 914 to 2134 m (3000 to 7000 ft)
  – Blade server off: -40° to 60° C (-40° to 140° F)
- Humidity:
  – Blade server on: 8% to 80%
  – Blade server off: 5% to 80%

Environment (NEBS):
- Air temperature:
  – Blade server on: 5° to 40°C (41° to 104°F). Altitude: -60 to 1800 m (-197 to 5905 ft)
  – Blade server on (short term): -5° to 55°C (23° to 131°F). Altitude: -60 to 1800 m (-197 to 5905 ft)
  – Blade server on: 5° to 30°C (41° to 86°F). Altitude: 1800 to 4000 m (5905 to 13 123 ft)
  – Blade server on (short term): -5° to 45°C (23° to 113°F). Altitude: 1800 to 4000 m (5905 to 13 123 ft)
  – Blade server off: -40° to 60°C (-40° to 140°F)
- Humidity:
  – Blade server on: 8% to 85%
  – Blade server on (short term): 5% to 90% but not to exceed 0.024 kg water/kg of dry air
  – Blade server off: uncontrolled

Note: "Short term" refers to a period of not more than 96 consecutive hours and a total of not more than 15 days in 1 year. (This refers to a total of 360 hours in any year, but no more than 15 occurrences during that 1-year period.)

Size:
- Height: 24.5 cm (9.7 inches)
- Depth: 44.6 cm (17.6 inches)
- Width: 2.9 cm (1.14 inches)
- Maximum weight: 5.4 kg (12 lb)

Figure 11. HS-21 Blade Server Specifications (IBM, 2007, p. 6)

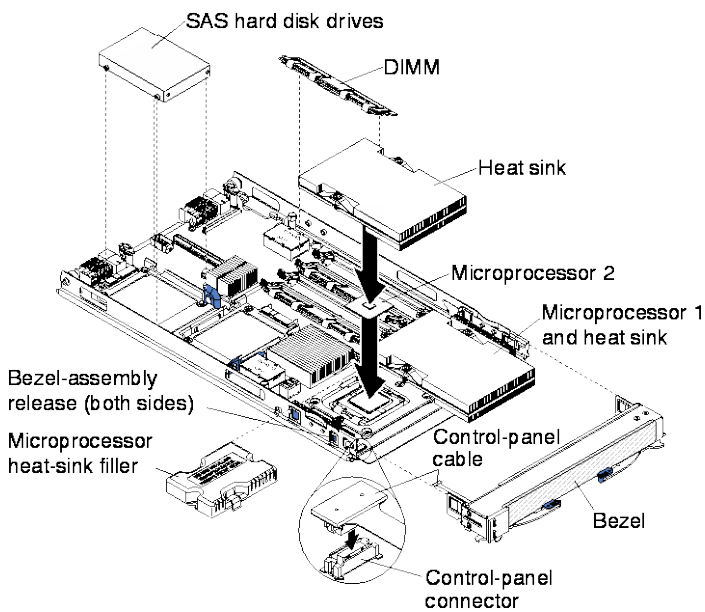

Figure 12. HS-21 Blade Server Major Components (From IBM, 2007, p. 10)

The blade servers were loaded with VMware vSphere ESX 4.1 and View 4.6, software products that allow virtual infrastructures, desktops, and their pertinent applications. As an additional enhancement to the tested NGT VDI, a Fusion-IO PCI card was connected providing storage support to the blade servers as illustrated in Figure 13 (FIRE, 2011).

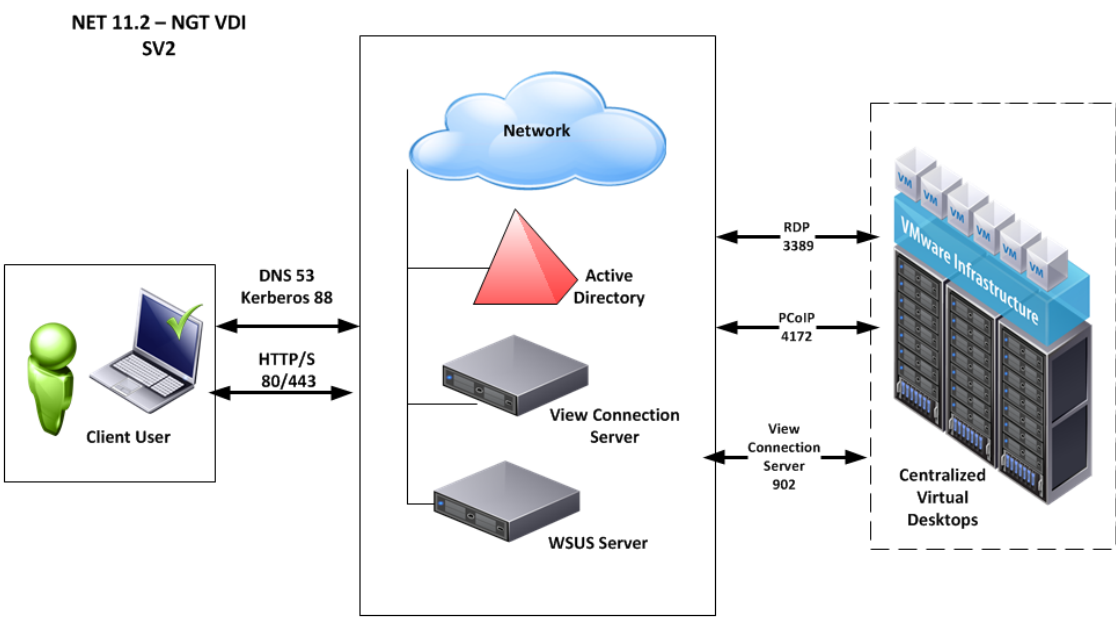

Figure 13. Net 11.2 NGT VDI SV2 (From FIRE, 2011)

Once the physical infrastructure was in place, TW11 Network technicians built a virtualized infrastructure, two groups of virtual desktops associated with physical space locations (WASP's Engineering Log Room and the LFOC), and loaded applications familiar to typical onboard users. Within each assigned physical space, TW NMCI workstations running both Windows XP and 7 were connected to the NGT blade servers through an embarked NGT switch. Each NMCI workstation was loaded with a VMware virtual client that provided the virtualized IT-21 COMPOSE desktop. Functionality tests were conducted covering typical workday application usage and windows updating. WASP crewmembers who were surveyed responded positively to NET 11.01's observed performance as captured in Figure 14.

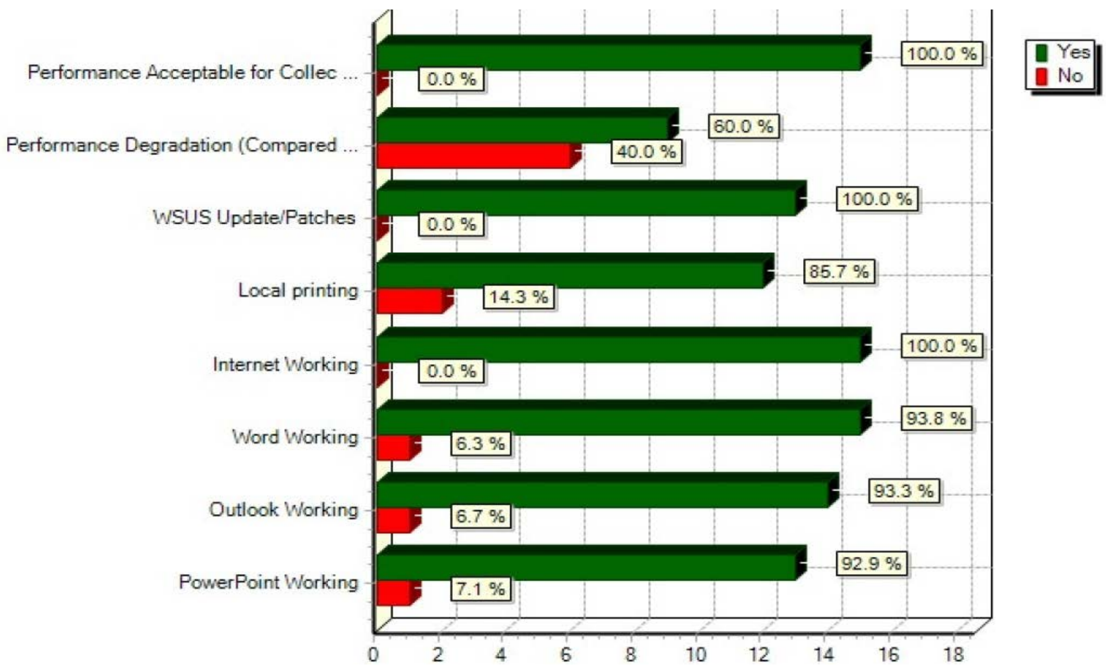

Figure 14. Net 11.01 Trident Warrior Survey Results

Satisfying the nominal afloat workday functionality experiment requirements, the TW11 Networks Focus Area Lead allowed for additional testing to be accomplished. Technical representatives onboard resonated with the challenges of network integration and were open to establishing scenarios geared toward improvementing the observed process. Of the administration and life cycle management benefits VDI presents, the NET 11.01 Final Test Report also recognizes ease of deployment. For instance, the embarked NGT VDI servers and network was installed within a working day! From a pre-existing VDI, ship-hosted perspective, further scenarios demonstrated the capability of creating hundreds of virtual clients in a few minutes.

Final testing for the TW11 underway period consisted of performance comparisons of the blade servers' configuration. Neworks Focus Area idenitified managebility as elemental to VDI's operational functionality. IBM's HS-22 blade servers' performance creating and managing clients was limited to the hardware brought onboard WASP. Scalabilty of the servers to provide for a typical MEU deployment approached costs that seemed unfeasible. Off-the-shelf, HS-22s could provide the services needed for amphibious network agility; however, administrative and user benfits could be negatively influenced by overall performance degradations signature to the servers' spinning disks.

The configuration comparsions centered on modifying the configuration of the HS-22. One represented the COTS standard configuration. The other server had a Fusion-IO, solid-state drive, PCI card connected for storage. The comparision occurred simulating a boot-storm, multiple users logging on at an instant, while observing CPU processing and memory differences. On a smaller scale of five and, later, twenty users logging on, the servers with the Fusion-IO PCI card consistently utilized less CPU processing and accessed more memory faster, nearly instantly as depicted in Table 3.

IBM HS-22 (COTS)				**IBM HS-22 (With Fusion Io and PCI Card)**			
5 VMs				20 VMS			
CPU Usage	**CPU Capacity**	**Memory Usage**	**Memory Capacity**	**CPU Usage**	**CPU Capacity**	**Memory Usage**	**Memory Capacity**
246 MHz	8.532 GHz	3388 MB	16383.05 MB	1463 MHz	224 GHz	35469 MB	49140.05 MB
% CPU Usage		**% Memory Usage**		**% CPU Usage**		**% Memory Usage**	
0.0288		0.2067		0.00653		0.7217	

Table 3. Comparison of 5 VMs Utilizing IBM HS-22 Configuration vs. 20 VMs Utilizing IBM HS-22 with Fusion Io and PCI Card

VI. RE-ENGINEERING AMPHIBIOUS INTEGRATION

A. TECHNOLOGY PROPOSAL

NPS IS4220 and IS4250 course materials recommend when incorporating higher levels of process integration, a technology, portal and middleware, is needed. Previous research, including field testing during Trident Warrior 2011, led to a technology insertion designed to support a re-engineered network integration process. Virtual Desktop Infrastructures (VDIs) could allow ship-hosted desktop capabilities for inorganic/embarking workstations regardless of their hardware and software incompatibilities to all DoD and subsequent policies. A VDI could provide the virtual portal for a virtual client, effectively acting as middleware. As long as the local servers and network infrastructure could accommodate a requested number of "virtual desktops," they would be securely pushed to any terminal connected. VDI technology insertion could directly address the areas needed for improvement introduced while defining the current view.

1. Revised Enterprise Architecture

Observations from the existing integration process revealed a multitude of supporting activities contributing overall inefficiency and low operational agility. Large durations of time are spent waiting for workstations to be separately brought into compliance. When in possession of the embarking computers, functional workers are utilized beyond rational limits. The VDI technology proposal enables a ship's IT crew to no longer take custody of inorganic equipment. As long as the capacity exists on the servers and switches, fully compliant virtual clients can be created and pushed to drops within seconds. Any updates requisite of O/S, ship image, or security patch is completed at the server upon an enterprise virtual client that is replicated and shared.

2. Re-Engineering Goals

Upon presenting the results of our current process, we were instructed to meet the following Business Process Re-Engineering goals:

a) Reduce the total duration to less than or equal to 7 days.
b) Reduce utilization for each to less than 70%.
c) Reduce overall wait-time by 90%.

In order to meet our goals, nearly every aspect of the current process stood in question. Sub-processes and their functional personnel had to be scrutinized for their necessity and contribution to the overall desired output. Technology solutions were sought to alleviate potentially needless activity. The Process Re-Engineering section provides detail to the solutions we applied approach the assigned goals.

3 Process Re-Engineering

Previous research, including field testing during Trident Warrior 2011, led to technology insertion to improve the re-engineered network integration process. Virtual Desktop Infrastructures (VDIs) could allow ship-hosted desktop capabilities to inorganic/embarking workstations regardless of their hardware and software incompatibilities to all DoD and subsequent policies. As long as the local servers and network infrastructure could accommodate a requested number of "virtual desktops," they would be securely pushed to any terminal connected. With VDI capability, the new process emerged:

e) JNOC receives a request to integrate a quantified number of workstations onto the ship's network. Servers are checked to verify if the capacity exists to host the requested virtual clients (with nominal processor, RAM, and storage performance) and manage the connected workstations. Overall count is determined and DSR validates if the network infrastructure can support the clients/workstations at the requested/assigned ship space.

f) b) Either remotely or locally on the ship's servers, ADP and IA verify the necessary O/S updates and security measures are current on the enterprise virtual client, respectively.

g) JNOC builds the required virtual clients to support the embarking organization. Simultaneously, DSR assigns and enables ports associated with the drops within arranged ship spaces/compartments.

h) The embarking IT personnel/customer delivers an operable workstation (laptop, desktop, or thin client device) to a directed ship space. JNOC assists with the workstation set up and coordinates with DSR to ensure the workstation successfully connects to the network and receives the virtual client.

i) When LAN connection and virtual client boots, JNOC members turn over the operation of the workstation to the embarkee for a courtesy quality-assurance check.

4. Generalizations for Re-Engineering

The introduction of a VDI and efforts to meet assigned goals forced generalizations to be made for the TO-BE model, BPR of shipboard network integration. They are provided here:

<u>a. VDI technology insertion</u>: Desktop virtualization alleviates the ship's IT crew from previous sub-processes, and their resulting bottlenecks, dealing with workstation transport/hand-over and the requisite hours for hardware/software compliance. Initial workstation operability, inventories and accountability remain with the embarking IT organization's personnel. Trouble-shooting, slicking, re-imaging, updating, and patching approved hardware is no longer necessary. Nearly any network-capable device can receive a virtual desktop.

<u>b. Sub-Process Duration</u>: VDI insertion into the TO-BE model process not only removed former processes, it also gave a greater agility to resulting activities. Rather than ADP and IA individually installing updates and security patches, both need only to ensure their enterprise versions are current, which is done at the server or remotely in minutes. With JNOC and DSR freed from deliveries and trouble-shooting, their time can be concentrated and better utilized toward the refined activities. Maintaining 300 instances within a scenario and reducing the overall duration called for additional personnel within each functional group.

<u>c. Decision Probabilities</u>: Decision points and their probabilities were required to be readdressed. Server and network infrastructure capacity decisions were still necessary to meet requests of various arrangements of embarking personnel and workstations; however VDI hardware components allow for greater capacities from the ship's hosting server. Also, final JNOC operational testing of the workstations is conducted after the initial check, from embarking personnel have proven it functional. Also, all decisions made through DSR swim-lanes remain consistent with the AS-IS model.

d. Associated Costs: Where the current process assumed only two personnel at any time during a workday are focused upon an integration activity, goals in overall duration, utilization, and wait-time required more support. AS-IS pay-grades did remain the same; however, a more detailed study could distinguish a more realistic dispersion of enlisted pay-grades. Overall numbers of supporting IT personnel do stay true to shipboard and MEU manning.

VII. CONCLUSION

Currently, LHD/LHAs are being assigned a broad range of missions where formerly disparate organizations must quickly and seamlessly integrate to begin addressing the task ahead. Although used for lack of alternatives, AS-IS core processes' strict focus on standardization interferes with critical timeliness and full integration. Field testing VDI insertion, like efforts within Trident Warrior 2011, peer into a new architecture where mission planning expectations can be met. No other technology shows as much promise to allow supported processes be in accord with an organizing logic of unification. Field tests lead to full enterprise deployment; however, only through a well-thought Enterprise Architecture, can technology determine success.

With nearly limitless combinations of options, we maintained a focus on our goals and created a tailored, yet practical, LHD/LHA VDI capable of hosting any embarking organization. The conducted BPR transformed the integration process and provided insight to manning levels, as well. In order to meet the prescribed requirements, each functional workplace had to increase their supporting personnel. We found the optimal blend to be: 12 members within the embarking organization, 11 within JNOC, 7 supporting DSR, and 4 each in ADP and IA. Any deviation from this arrangement led to missing our goals or drifting into unrealistic manning levels. Addressing each goal:

a) Reduce the total duration to less than or equal to 7 days:

TO-BE model resulted in less than 10 hours of overall duration. Nearly full commitment of ship's IT personnel could integrate a MEU's full UNCLAS network integration in one extended working day. If done after normal working hours, the integration could have very little impact on overall operations.

b) Reduce utilization for each to less than 70%:

Utilization percentages of JNOC and "Embarkee QA" personnel were only slightly above threshold, 72% and 70%, respectively. The TO-BE model required unrealistic numbers within these two groups in order further reduce their utilization.

Being well below 80%, their workload through the revised integration process is acceptable. Utilization for DSR, ADP, and IA were able to fall below our instructed ceiling.

c) Reduce overall wait-time by 90%:

Another drastic improvement from the current integration process was the reduction in wait-time. The TO-BE model produced a 93% reduction in overall waiting time from 22,548 to 1656 and is illustrated in Table 4.

Admittedly, our BPR is limited to our experiences, data collected, and necessary assumptions. Overall, the current personnel, processes, and technology regarding workstation network integration for a specific platform has been accurately captured. Without exception, all input received addressing each area in the AS-IS architecture has shown much room for improvement. Our research reveals that addressing solely personnel, processes, technology, or a combination of only two, will not yield a sufficient solution to current enterprise practices and resources. It is in the best interest for the U.S. Navy to pursue multiple means for interoperability.

Simulation Results					
Duration	9:45:20	Time			

Process Time And Cost					
Process	Scenario	Instances	Total Cost	Waiting Time (Time)	Total Time (Time)
Matthias_Becker_Gruber_ToBe	(default)	300	3795.43	1656:11:20	1402:45:00

Performers Queue Length and Utilization					
Name	Average	Min	Max	Utilized(%)	Idle(%)
Any member of JNOC	39.19	0	174	72.83	27.17
Any member of EMBARKEE QA	35.13	0	122	70.9	29.1
Any member of ADP	16.31	0	55	56.59	43.41
Any member of DSR	86.76	0	230	68.34	31.66
Any member of IA	16.31	0	55	56.59	43.41

Table 4. TO-BE Amphibious LAN Integration Metrics

A. FOLLOW ON RESEARCH RECOMMENDATIONS

The findings from this research have led to further questions and areas of interest that could be tapped for future research. The view of this research was from a broad view that maintained a wide scope which left open the opportunity for further analysis and in-depth research in the realms of technology, personnel management and costs with regards to VDI technology onboard Naval vessels.

This thesis focused on one configuration of VDI that was considered to be the most easily achievable through the use of minimal changes to technology onboard LHA/LHDs while providing the most benefit in terms of performance and reduction in man hours and overall manning. There is obviously much more technology available to the market than what was discussed and these additional technologies may present possibilities for future research. Follow on research may be able to discover better performing technologies or technological methods that push past the goals of this research and open up new possibilities in the endeavor to integrate and jointly operate.

Although only LHA/LHDs were discussed, other platforms would greatly benefit from the ability to be able to embark Marines and personnel from disparate groups. One notable and new platform that would greatly benefit from ability to embark groups of personnel and give them computing capability quickly and efficiently is the Littoral Combat Class platforms. These new platforms are designed to embark and debark personnel, Sailors, Marines and Civilians as part of LCS's mission module systems. The mission modules require additional equipment and personnel to embark upon the ships, requiring Internet access, network access, and computing capabilities for completion of mission tasking and basic administration. Despite the massive size difference between a large deck amphibious vessel and a Littoral Combat Ship, the needs are almost the same. Follow on research could be used to determine if VDI could be a solution to the challenges faced onboard LCS and what technology and methods could be utilized in order to meet needs of embarking personnel and equipment.

Further research could be leveraged onboard LCS and other small platforms in terms of personnel manning. The goals of this thesis were to greatly reduce the amount of time spent on embarkable computing needs and the total amount of manning required to accomplish embarking personnel. It was found that through the use of VDI these objectives could be met and the total time and total amount of personnel could be reduced. The benefits in the reduction of man-hours required and the amount of personnel required would be huge onboard platforms like LCS, which is manned by a total of forty Sailors. With so few Sailors and requirements similar, but at a lesser scale, as the larger vessels, reduction is manning through improved processes and technology is an area of research that could open up LCS, smaller platforms and the Navy as whole to more mission flexibility and speed.

An area not covered but of interest is the monetary cost of VDI systems and the implementation onboard Naval vessels. Further research could be used to determine costs associated with upgrading to and installing VDI systems onboard Naval vessels. Costs through the reduction of personnel and man hours spent could also be tested and researched in order to determine if the costs associated with the upgrades and install could be offset and surpassed through the cost savings from the reduction of man hours and the reduction total manning to support onboard networks.

LIST OF REFERENCES

15th Marine Expeditionary Unit. (2010, September 2). Retrieved from http://www.i-mef.usmc.mil/external/15thmeu/.

22nd Marine Expeditionary Unit. (2011). *What is a MEU?* Retrieved from http://www.marines.mil/unit/22ndMEU/Pages/CE/MEU.aspx

Anderson, S. (2011). Bold Alligator 2011. The blue and green team together again. *CHIPS*. Retrieved from http://www.doncio.navy.mil/CHIPS/ArticleDetails.aspx?ID=2328

Baskerville, R., & Myers, M. (2004). Making IS research relative to practice. *Management of Information Systems Quarterly*. Spec. issue on *Action Research in Information Systems*, *28*(3), 329–336.

Carbone, M., Wenke L., & Zamboni, D. (2008). Taming virtualization. *Security & Privacy, IEEE*, *6*(1), 65–67.

Chahal, S., & Glasgow, T. (2007). Memory Sizing for Server Virtualization. Retrieved from http://www.intel.com/it/pdf/memory-sizing-for-server-virtualization.pdf.

Chairman, Joint Chiefs of Staff. (2008, February 13). Joint operations. (DoD Joint Publication 3–0). Washington, DC: Author.

CHIPS (2011). *A Few Minutes with Rear Admiral Michelle Howard*. Retrieved from http://www.doncio.navy.mil/CHIPS/ArticleDetails.aspx?ID=2338.

Clark V. (2002). Sea Power 21: Projecting decisive joint capabilities. *United States Naval Institute Proceedings.* Retrieved from http://www.navy.mil/navydata/cno/proceedings.html.

Cohen, F. (2010). The virtualization solution. *Security & Privacy, IEEE, 8*(3), 60–63.

Crump, G. (2010). For data center virtualization, IOPS is more important than air. *Switzerland: Storage Switzerland LLC*. Retrieved from http://www.storage-switzerland.com/Articles/Entries/2010/9/14_For_Data_Center_Virtualization.html.

FFC Trident Warrior Program. (2011). *Trident Warrior 11 (TW11) Orientation Guide*. 2011. Retrieved from https://fire8.nps.navy.mil/.

Fusion io. (2011). *V3 Systems reinvents virtual desktop infrastructure with fusion powered I/O*. Fusion io Case Studies. Retrieved from http://www.fusionio.com/case-studies/v3systems/

Government Accountability Office. (2006). *DoD needs to ensure that Navy Marine Corps intranet program is meeting goals and satisfying customers.* Retrieved from http://www.gao.gov/products/GAO-07–51.

Hammes, T. (2007). The future of warfare. In Mclvor, A. (Ed.), *Rethinking the Principles of War.* Annapolis, MD: Naval Institute Press.

Hewlett-Packard (2011). *A step into the future*. Retrieved from http://h10131.www1.hp.com/public/nmci/about-nmci/.

Holmes, N. (2005). The turning of the wheel. *Computer, 38*(7), 98–99.

IBM (2011). *BladeCenter HS22 Type 7870, 1936, and 1911 installation and user's guide*. (11th ed.). Research Triangle Park: IBM.

IBM (2007). *BladeCenter HS21 Type 8853 installation and user's guide*. (4th ed.). Research Triangle Park: IBM.

Jordan, K. (2006). The NMCI Experience and Lessons Learned. Case Studies in National Security Transformation. 12, 9–10. Retrieved from National Defense University website: http://www.ndu.edu/CTNSP/docUploaded/Case%2012%20%20The%20NMCI%20Experience%20and%20Lessons%20Learned.pdf

Market Intel Group. (2010). The Future of Virtualization, Cloud Computing and Green IT – Global Technologies & Markets Outlook 2011–2016. MIG, 385. Retrieved from http://www.researchandmarkets.com/reports/1402312/the_future_of_virtualization_cloud_computing_and.

McKenna, P. (2011). U.S. navy chief: "I'm on a mission to stop using oil." *New Scientist, 2811*(29). Retrieved from http://www.newscientist.com/article/mg21028110.200-us-navy-chief-im-on-a-mission-to-stop-using-oil.html

Mullen, M. (2011). The national military strategy of the United States of America. 2011 Redefining America's military leadership. Retrieved from http://www.jcs.mil/content/files/2011–02/020811084800_2011_NMS_-_08_FEB_2011.pdf

Paquet, R. (2004). *Why you need a storage department*. Retrieved from http://www.fusionio.com/case-studies/v3systems/.

Perez, R., van Doorn, L., & Sailer, R. (2008). Virtualization and hardware-based security. *Security & Privacy, IEEE, 6*(5), 24–31.

Rosenblum, M., & Garfinkel, T. (2005). Virtual machine monitors: current technology and future trends. *Computer, 38*(5), 39–47.

Siebert, E. (2011). Managing storage for virtual desktops. Retrieved from http://searchstorage.techtarget.com/magazineContent/Managing-storage-for-virtual-desktops?pageNo=1

Vaughan-Nichols, S.J., (2006). New approach to virtualization is a lightweight. *Computer*, *39*(11), 12–14.

Vaughan-Nichols, S.J.; (2008). Virtualization sparks security concerns. *Computer*, *41*(8): 13–15.

VMware (2011). Top 5 reasons to adopt virtualization software. 22d MEU. Retrieved from http://www.vmware.com/virtualization/why-virtualize.html

VMware (2011). What is a virtual machine? Retrieved from http://www.vmware.com/virtualization/virtual-machine.html#c25951.

VMware (2011). History of virtualization. Retrieved from http://www.vmware.com/virtualization/why-virtualize.html.

VMware (2011). Reduce costs with a virtual infrastructure. Retrieved from http://www.vmware.com/virtualization/virtual-infrastructure.html.

Wikipedia (2012). Comparison of SSD with hard disk drives. Retrieved from http://en.wikipedia.org/wiki/Solid-state_drive

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. D.C. Boger
   Naval Postgraduate School
   Monterey, California