REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
04/03/2016	Master's of Military Studies	SEP 2015 - APR 2016

4. TITLE AND SUBTITLE
In Search of a Single Switch Multilevel Security Solution: Efficient Network Security and Data Sharing in a Mission Partner Environment

5a. CONTRACT NUMBER
N/A

5b. GRANT NUMBER
N/A

5c. PROGRAM ELEMENT NUMBER
N/A

6. AUTHOR(S)
Adams, Agur, S, Major, USMC

5d. PROJECT NUMBER
N/A

5e. TASK NUMBER
N/A

5f. WORK UNIT NUMBER
N/A

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
USMC Command and Staff College
Marine Corps University
2076 South Street
Quantico, VA 22134-5068

8. PERFORMING ORGANIZATION REPORT NUMBER
N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Mr. Jeff Castro
Capabilities & IT Portfolio Management
HQMC CD&I, C2 & CEW Integration Division
3300 Russell Road
Quantico, VA 22134

10. SPONSOR/MONITOR'S ACRONYM(S)
MMS Mentor's Name

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
N/A

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
This paper examines technical solutions to sharing data between multiple data link layer enclaves, merged physically but separated logically, in a mission partner environment to determine how and where efficiencies can be gained without sacrificing security. This paper proposes a novel network architecture and network switch design using hardware virtualization and software defined networking. Adoption of the proposed design should be weighed against the financial, material, and labor investment necessary to certify and accredit a new cross domain solution. Through hardware virtualization and software defined networking, warfighters can truly achieve maneuver in the information environment.

15. SUBJECT TERMS
Mission Partner Environment; Hardware Virtualization; Software Defined Networking; Cross Domain Solutions; Multilevel Security

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			USMC Command and Staff College
					19b. TELEPHONE NUMBER *(Include area code)*
Unclass	Unclass	Unclass	UU	37	(703) 784-3330 (Admin Office)

United States Marine Corps
Command and Staff College
Marine Corps University
2076 South Street
Marine Corps Combat Development Command
Quantico, Virginia 22134-5068

MASTER OF MILITARY STUDIES

TITLE:

In Search of a Single Switch Multilevel Security Solution:
Efficient Network Security and Data Sharing in a Mission Partner Environment

SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF MILITARY STUDIES

AUTHOR:

Major Agur Adams, United States Marine Corps

AY 15-16

Mentor and Oral Defense Committee Member: Professor Matthew Flynn
Approved: MATTHEW FLYNN
Date: 3/30/16

Oral Defense Committee Member: Mr. Jeffrey Castro, Capabilities Officer/IT Portfolio Manager, HQMC CD&I, C2 and CEW ID
Approved: J. W. Gordon
Date: 3/30/16

DISCLAIMER

THE OPINIONS AND CONCLUSIONS EXPRESSED HEREIN ARE THOSE OF THE INDIVIDUAL STUDENT AUTHOR AND DO NOT NECESSARILY REPRESENT THE VIEWS OF EITHER THE MARINE CORPS COMMAND AND STAFF COLLEGE OR ANY OTHER GOVERNMENTAL AGENCY. REFERENCES TO THIS STUDY SHOULD INCLUDE THE FOREGOING STATEMENT.

Executive Summary

Title: In Search of a Single Switch Multilevel Security Solution: Efficient Network Security and Data Sharing in a Mission Partner Environment

Author: Major Agur Adams, United States Marine Corps

Thesis: Hardware virtualization and software defined networking provide the means to segregate and secure multiple data link layer enclaves on a single switching platform.

Discussion: Present and future military operations demand the integration of international partners, government agencies, nongovernmental organizations, and private entities to achieve lasting national security solutions. Such close integration necessitates information sharing, but current command and control structures prohibit efficient and secure information flow between organizations due to the US military's and international community's dependence on isolated national classified computer networks (e.g., the Secure Internet Protocol Router Network). Currently, these computer networks remain secure, blocking unauthorized access through physically separate but functionally identical information systems. However, modern computer and networking technologies, such as hardware virtualization and software defined networking, can provide integrated, secure networking solutions without physically redundant information systems. Thus, the accepted practice of physical segregation to achieve security wastes money, space, and power while restricting the free flow of information to direct material exchanges and liaisons. This paper examines technical solutions to sharing data between multiple data link layer enclaves, merged physically but separated logically, in a mission partner environment to determine how and where efficiencies can be gained without sacrificing security. This paper proposes a novel network architecture and network switch design but adoption should be weighed against the financial, material, and labor investment necessary to certify and accredit a new cross domain solution.

Conclusion: The proposed solution demonstrates security and efficiency in computer networking can complement each other rather than work at odds as has been the struggle of the past. Through hardware virtualization and software defined networking, warfighters can truly achieve maneuver in the information environment.

Illustrations

Page

Figure 1. The Afghanistan Mission Network 3

Figure 2. Client-Side Separation Virtual Machine Monitor 12

Figure 3. Server-Side Virtualization in Cross-Domain Solutions 14

Figure 4. Proposed Network Architecture Design.....19

Figure 5. Proposed Network Architecture Design with Software Defined Networking Switch 22

Table of Contents

Page

DISCLAIMER i

EXECUTIVE SUMMARY ii

LIST OF ILLUSTRATIONS iii

ACKNOWLEDGEMENTS v

INTRODUCTION 1

BACKGROUND 3

LITERATURE REVIEW 7

TECHNICAL OVERVIEW: HARDWARE VIRTUALIZATION 10

TECHNICAL OVERVIEW: SOFTWARE DEFINED NETWORKING 15

PROPOSED SOLUTION 18

IMPLICATIONS 23

END NOTES 25

BIBLIOGRAPHY 28

Acknowledgements

I gratefully acknowledge Professor Matthew Flynn and Professor Gary Brown of Marine Corps University for their invaluable instruction and contributions to this work.

Additionally, I thank Mr. Jeffrey Castro of the Command and Control / Cyber and Electronic Warfare Integration Division, Capabilities Development Directorate, Combat Development and Integration, Headquarters Marine Corps for proposing this research topic, offering technical guidance, and providing commercial and government points of contact.

I also show gratitude to the commercial information technology providers and government research scientists who graciously offered their time to support this research. Mr. Aaron Brace, Colonel Keith Maresca (retired) and Mr. Jason Lozada of Dell advised technically on multilevel security systems and demonstrated the *SecureView* system. Mr. Anthony Kennedy of Air Force Research Laboratory provided supplemental documentation on the *SecureView* project. Dr. Myong Kang and Ms. Swati Shah of the Naval Research Laboratory introduced the *Network Pump II*^{TM} and *Xenon* cross domain solutions. Ms. Kelly O'Connell and Mr. Ahmed Ali of VMware instructed on incorporating commercial software defined networking products into cross domain solutions. Captain Benjamin Pimentel of the Office of Naval Research helped identify past and current research interests in cross domain solutions in the Navy and Marine Corps.

I offer thanks to my wife, Krista, for her loving support, patience, commitment, and encouragement as well as to my local church and small group community for their prayers, thoughts, and support throughout my time at Marine Corps University.

Introduction

Historically, multilevel security (MLS) networks comprise isolated but identical information systems (IS) that are physically separated to prevent unauthorized data transfer between networks (e.g., the Nonclassified Internet Protocol Router Network and the Secure Internet Protocol Router Network). However, because Department of Defense (DoD) and Intelligence Community (IC) users need continuous and simultaneous access to information at different security levels, duplicate ISs are emplaced from the core to the edge of MLS networks to provide universal access to a single user. Such security restrictions and physical redundancies hamper communication at the operational and national strategic level when trying to share information between coalition partners in a mission partner environment (MPE). Additionally, the large number of network connectivity requirements, excessive information assurance procedures, and command and control structures prohibit the free flow of information between organizations. These problems occur despite the fact that new computer and networking technologies make physical consolidation of redundant ISs possible while preserving security. However, current DoD policies make it difficult to change despite the fact that research has shown physical consolidation of redundant ISs would drastically reduce cost, space, weight, and power consumption while also being more user-friendly.[1] Thus, current DoD and IC practices of duplicating ISs in MLS networks wastes physical resources, creates excess administrative burdens, and introduces unnecessary attack vectors. Once corrected, the gains from this increased functionality will be the freedom to maneuver in the information environment while enhancing security.

As early as 1972, computer security researchers have sought various cross-domain solutions (CDS) (i.e., ISs capable of storing, processing, or viewing information at multiple levels of security) to consolidate redundant physical components while preserving information security

between separate domains.[2] One material solution is to integrate MLS networks at the network switch, the immediate device that links two or more computers together into a Local Area Network (LAN). Dr. Myong Kang and Dr. Ira Moskowitz originally proposed the first network switch multilevel security solution with the introduction of the *Network Pump*^{TM} in 1993.[3] However, due to limitations in technology at the time, their device did not enable the true physical consolidation of networking components. Their device only allowed acknowledgement of information transferred from a lower classified network to a higher classified network by reducing the capacity of a covert timing channel between networks.[4]

Recently, two new technologies make physical consolidation possible while preserving flat data link layer networks: hardware virtualization and software defined networking (SDN). Over the last sixteen years, the maturity, flexibility, and ubiquity of these technologies in commercial markets transformed the commercial computer networking industry and attracted attention from many DoD and IC researchers. The work of these researchers, in conjunction with information technology (IT) commercial companies, has led to numerous point-to-point CDSs operating within branches of the DoD and IC. With this breakthrough, the private sector in cooperation with government research laboratories and agencies, has overcome this longstanding problem.

Despite the success of these CDSs and the advantages of consolidating ISs, the DoD and IC has not comprehensively adopted these newer technical systems due to the security challenges and risks associated with CDSs. Further, when considering employment of cross-domain (CD) technology in an MPE framework of multiple international partners, senior Joint Force communication leaders see no justification to physically interconnect national classified networks with mission partner networks when physical separation can provide adequate security.[5]

Thus, the challenges associated with implementing CDSs in an MPE framework are: 1) developing a network architecture and information assurance policies that ensure the secure use of CD devices; 2) eliminating covert communication (i.e., the illicit transfer of information against a system's security policy by manipulating the system's design properties in a way they were not intended) between security domains; 3) reducing complexity in CD device design for certification and accreditation; and 4) proving the benefits of CD technology outweigh the risks. Given these challenges, this paper examines technical solutions towards achieving the goal of hosting multiple virtual data link layer enclaves on a single network switch.[6] This paper proposes a novel network architecture and network switch design using hardware virtualization and SDN as a solution to the challenges and risks associated with CDSs.

Background

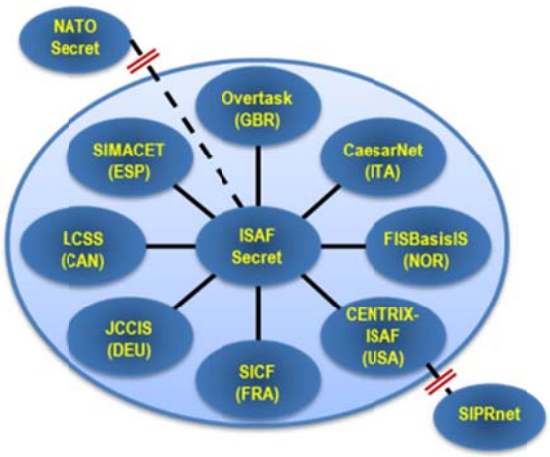

Figure 1. The Afghanistan Mission Network[7]

The North Atlantic Treaty Organization (NATO) was the first to utilize the concept of a MPE, although not named as such at the time, in Operation Enduring Freedom using the Afghanistan Mission Network (AMN). Originally created in 2009, the AMN is a federation of multiple networks using existing commercial ISs (see Figure 1). The International Security Assistance Force (ISAF) used the AMN for coalition command and control communication at the ISAF

SECRET classification.[8] By May, 2011, forty-eight nations, both NATO and non-NATO members, relied on the AMN as the sole means to share information between forces.[9] Each nation maintained access to their own physically separate classified and unclassified national networks, but, per the direction of the ISAF Commander, the majority of communication occurred over the AMN to achieve unity of effort.[10] Troop contributing nations had to maintain six core services on the AMN to facilitate ISAF command and control: 1) web browsing; 2) chat; 3) Voice Over Internet Protocol telephony; 4) Video Tele-Conferencing over Internet Protocol; 5) electronic mail with attachments; and 6) global address list sharing.[11] ISAF communication staff also developed Joint membership and exit instructions (JMEI) to integrate new national militaries using their existing networking hardware and software while adhering to the information assurance policies of the AMN.[12]

The AMN worked well in unifying command and control, but came at the expense of building an entirely separate physical network with additional administrative overhead. This was on top of each nation already maintaining multiple separate national classified and unclassified networks. As an example of the administrative burden, Joint communication staff members recorded 453 issues of compliance with the JMEI when a similar mission partner network was stood up again with various coalition members during exercise COMBINED ENDEAVOR in 2013. In COMBINED ENDEAVOR in 2014, 317 compliance issues arose of which 85% were instances where the JMEI instructions were simply not followed.[13]

Regardless, recognizing the value of this integration, the Joint Staff directed the concepts behind the AMN become a core part of future Joint Force operations. Hence, the Joint Staff J6 took control of the MPE initiative, making it part of continued Joint Force practice. Formally, the Joint Staff defined the MPE as “a federated network concept supporting the connection of multi-

ple networks through existing national systems with applications and tools to enable mission partner information sharing within a single environment."[14] However, Martin Westphal, Vice Director for the Joint Staff J6, argues that a MPE is more than a technological means to communicate: a MPE is the trusted interaction between coalition partners to accomplish the commander's objectives. As such, US Joint Forces will deploy with an additional "stack" of ISs dedicated to provide connectivity to a mission partner network, while establishing separate physical connections to national classified networks through redundant clusters of ISs.[15] In this, Westphal's effective and secure solution to the information sharing problem in an MPE framework utilizes existing materials, but results in an exacerbated level of redundancy and administrative overhead for network administrators.

Separate from the MPE initiative, demand by DoD and IC users and IT administrators for certified and accredited ISs capable of processing multiple levels of security has spawned over 800 different CD products within the last four decades.[16] Such a high number indicates not only the relevance of the issue to military and government agencies, but the large number of disjointed efforts to reach a cohesive solution. Thus, on July 10, 2006, the DoD and Office of the Director of National Intelligence (ODNI) jointly created the Unified Cross Domain Management Office (UCDMO) to oversee the isolated, ineffective, and duplicative efforts in both communities. The UCDMO charter contained two principal tasks: 1) establish and maintain the authoritative baseline list of approved CDSs; and 2) oversee the development, testing, and fielding of all CDSs. As the demand increased for CDSs with the evolution of the DoD's Joint Information Environment and IC's Information Technology Enterprise, the UCDMO's role expanded to oversee enterprise-level cross-domain services (ECDS). ECDSs are unique in that they holistically span entire organizations (in contrast to point-to-point solutions) and that system administrators con-

trol them from the top down. With this new focus on services, UCDMO's title officially changed to the Unified Cross Domain Services Management Office (UCDSMO) on March 26, 2014.[17]

The consolidation of CDSs under the UCDSMO has been effective in coordinating the efforts of the Services, but a gap still remains in existing approved solutions that can integrate multiple virtual data link layer enclaves on a single switching platform. Numerous CDSs exist which can connect networks of different security levels at the Internet Protocol (IP) layer, but no solution is capable of preserving a flat data link layer architecture.

Additionally, given that "security is a system property" and not dependent solely on individually certified and accredited ISs, a DoD component or IC agency cannot simply choose a CDS from the UCDSMO's approved baseline list and begin to use it in a MPE.[18] DoD policy dictates that any approved CDS from the baseline list requires site-specific testing and evaluation before granting an authorization to operate.[19] Additionally, any new development of CD technologies must be approved by UCDSMO and may not duplicate a capability already satisfied under the existing baseline list. Further, new CD technologies must undergo a rigorous security control assessment, and the research sponsor must provide lifecycle maintenance, protection, and monitoring of the product's use. UCDSMO designed these constraints to ensure all future CD efforts align with the strategic goals of cost-reduction, systems integration, and information sharing without wasting resources.[20] Hence, future research and development towards new alternative CDS is a difficult, expensive, and time consuming process that provides little incentive for the Services to invest in new ways to break the status quo of building redundant physical networks.

As part of their work, UCDSMO defined three types of CD devices: 1) transfer, 2) access, and 3) multilevel.[21] UCDSMO distinguishes each type of device based on how it uniquely interacts with data of different security classifications. A transfer device allows information to

move from one security domain to another. An access device allows a user to observe information in multiple security domains, but does not permit information to move between security domains. A multilevel device stores and process information of different security levels allowing users to observe only certain data based on their security clearance. The assumption is that efficient use of information gives the DoD and the IC an advantage over their adversaries.

In summary, the DoD seeks to accomplish information sharing and integration through the MPE framework. NATO and its coalition partners created the AMN out of the need to communicate effectively and efficiently without facing the restrictions of national classification and disclosure requirements on a regular basis. However, given the extensive work that DoD researchers and IT commercial corporations have done to show that CD technology can be deployed safely and securely in MLS networks, DoD's current solution of using additional ISs to establish mission partner networks wastes resources. Thus, the DoD needs a new technical solution; one that incorporates CD technology in an MPE framework. This will accelerate information sharing among partner nations and restore maneuver to the information environment.

Literature Review

Academic researchers, government scientists, and commercial IT providers have made many efforts to build and verify secure MLS ISs over the last forty-five years.[22] Dr. John Rushby, Program Director at SRI International and a 35 year expert on design and assurance for critical systems, wrote one of the leading papers on this subject entitled *Design and Verification of Secure Systems* in 1981.[23] In this seminal paper, Dr. Rushby proposes a separation kernel as the domain separation mechanism to enforce MLS in either a single or multiprocessor distributed system. Up to this point, researchers and designers advocated for a single security kernel to manage all of the functions of an MLS IS to include security. However, they were unable to verify

the operation of such all-encompassing security kernels for multiple reasons, including the code-base length. Hence, Dr. Rushby's theory of a small, simple separation kernel responsible for only a few well-defined tasks has become one of the foundational concepts in CD design today.

Yet, according to the National Security Agency's Information Assurance Directorate's (NSA's IAD) 2010 report on *Separation Kernels on Commodity Workstations*, researchers and IT providers have still not been able to fully realize Dr. Rusby's separation kernel theory using commodity computer technology.[24] This is due to the fact Computer vendors continually aim to distinguish their products from other competing products. As such, they continuously add proprietary functions to commodity hardware making any assured implementation of a separation kernel unrepeatable. Further, the NSA's IAD realized that assurance of only the separation kernel does not produce a secure system. From supply chain corruption to poorly written applications, verifiers must examine the entire system to give evidence of high assurance.[25] Hence, John McDermott et. al from the Center for High-Assurance Computer systems of the Naval Research Laboratory (NRL), sought to correct this deficiency by relaxing the standard for a separation kernel with the introduction of a separation Virtual Machine Monitor (VMM). With this change in standards, McDermott et al. show in their paper, *Separation Virtual Machine Monitors,* that a small, simple separation VMM can be built and reasonably verified to satisfy the isolation requirements of a separation kernel while using modern commodity hardware.[26]

Separate from the aforementioned initiatives in security research, the commercial IT sector has also made tremendous shifts towards hardware virtualization technologies over the last fifty years. While the International Business Machines Corporation introduced the first working mainframes with hardware virtualization as early as 1968, more recently VMware, Inc. accelerated the transformation of the industry with a series of hardware virtualization products begin-

ning in the late 1990s. The explosion in popularity of hardware virtualization led academics and government researchers to explore hardware virtualization's application to MLS ISs. Two of the latest efforts are the *SecureView* project at the Air Force Research Laboratory (AFRL) and the *Xenon* project at the NRL. These projects capitalize on open source hardware virtualization solutions, such as the *Xen* open source hypervisor project, to create "government of the shelf" options for the Services.[27] Although not published formally in academic literature, the project documentation for *SecureView* helped form the basis of this research and will be discussed throughout this paper.

Lastly, while scholars have written a great deal on SDN itself, current literature shows little focus on applying SDN to MLS networks or MLS network switch design.[28] Xiong Liu, Haiwei Xue, Xiaoping Feng, and Yiqi Dai offer one model in their 2011 paper, *Design of the Multi-Level Security Network Switch System which Restricts Covert Channel*.[29] However, their proposed solution does not provide a comprehensive approach to creating data link layer enclaves through SDN. Instead, their solution focuses on mitigating covert channels through filters as a transfer CDS. Consequently, the scant relevant literature on integrating data link layer enclaves includes VMware's commercial publication, *Micro-segmentation For Dummies*, and the corresponding technical standards from the Internet Engineering Task Force (IETF) on virtualization of data link layer networks over IP layer networks.[30] However, these documents and their proposed solutions were not intended for application to MLS networks and therefore do not address all of the unique considerations of CD design such as the use of encryption as a barrier between security levels. Thus, there remains a need for an overview of a complete technical solution.

Technical Overview: Hardware Virtualization

Hardware virtualization is an integral component of many existing CDSs. To this end, two categories of solutions encompass how hardware virtualization can secure information flow between different levels of security while reducing redundancy: 1) client-side hardware virtualization, and 2) server-side hardware virtualization. As Dr. John Rushby states "the purpose of a secure [cross-domain] system is not to prohibit information flow between different security levels, but to control it."[31] Thus, both client-side and server-side virtualization in CDSs share the common goal of trying to control the flow of information between security levels. Again, for the US military, this functionality is crucial given that it provides an information advantage to the warfighter while increasing efficiency and preserving security.

Hardware virtualization is the emulation of hardware devices in software programs. In general, with hardware virtualization an intermediary mechanism called the hypervisor separates the physical elements of a computer from the software. The hypervisor, similar to an operating system (e.g., Windows 10, Macintosh OS X, etc.), is a piece of software that controls all aspects of the physical computer, but differs from other types of software in that it enables the creation and execution of virtual machines (VM). A VM emulates an entire computer in software permitting other software programs, including operating systems, to execute within the VM as if they were on a physical computer.[32]

The advantages of a VM are that it possesses the attributes of a software file combined with the power of a full computer without the obstacles of a physical machine. Through the hypervisor, VMs (i.e., computers) can be started or stopped, copied or moved, and saved or deleted like any other file on a computer. Multiple instances of VMs can also run on the same physical machine, splitting or sharing the physical machine's resources between them. Further, a system administrator or another software program can move VMs between clusters of computers, such

as within a data center, to ensure high availability. Thus, using hardware virtualization anyone can make a large static physical computing infrastructure into a flexible, scalable platform customized for specific needs.[33]

This was not able to be done in the past because of limits in processing power, lack of computer memory, and the high cost to run emulation machines. Now, microprocessor design and integrated circuit density has scaled to such a level where hardware virtualization technology can be employed on commodity hardware at exceptionally low cost. Additionally, the software ecosystem surrounding virtualization has matured to the point where hypervisors are freely available and ubiquitous. Lastly, the popularity of using VMs has soared with the advent of cloud computing because with VMs companies can dynamically acquire, purchase, sell, and manage computing power at the speed of deploying software. However, although the private sector has gained much from hardware virtualization over the last two decades, there are a number of security concerns that must be addressed before it can be applied safely in a CDS under a military setting. These security concerns begin with the role and function of the hypervisor.[34]

The hypervisor is commonly referred to as a virtual machine monitor (VMM) due to its integral role in managing the associated VMs on one physical machine. The VMM accomplishes its management responsibilities by trapping the VM's "guest" operating system's hardware calls through the emulation process. The VMM translates those system calls into actions on the physical computer and delivers the outcomes back to the VM. Most importantly, as the intermediary, the VMM controls what, if any, access the VM has directly to the computer's physical resources such as memory, the central processing unit, the hard disk drive, and the network interface card.[35]

With client-side virtualization in a CDS, the user's individual physical computer contains a modified type of VMM called a separation VMM.[36] As defined by computer security researchers John McDermott et. al, a separation VMM is advantageous over a conventional VMM in that it will: "1) run on modern commodity hardware; 2) virtualize modern commodity operating systems; 3) be smaller and simpler than a conventional VMM; 4) use fewer and simpler communication paths; and 5) have the highest assurance justifiable for its modern commodity hardware."[37] Typically, a separation VMM is installed directly on the platform hardware of a single laptop or desktop computer. With this, one or more VMs, each having a unique security level, simultaneously execute on top of the separation VMM on a laptop or desktop computer (see Figure 2).

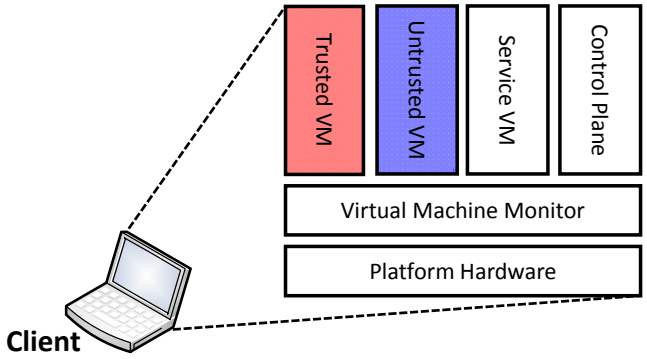

Figure 2. Client-Side Separation Virtual Machine Monitor[38]

The separation VMM isolates each VM from the others in accordance with the security policies of the organization. Separation VMMs typically enforce isolation in three ways: 1) isolating the physical memory regions and hardware devices accessible to the VM, 2) scheduling the VM's access to common resources to prevent interference, and 3) restricting communication between VMs to only authorized channels. Additional VMs, called service VMs, may also be used by the separation VMM to further isolate hardware devices or perform security services for other VMs on the machine. For example, one common service VM is a "logger", which records

any violations of the system's security policies. Existing client-side virtualization CDSs utilize some or all of the aforementioned segregation techniques to achieve isolation.[39]

Client-side virtualization CDSs often also use encryption to achieve isolation. To meet this end, in an effort to promote commercial solutions for classified capabilities (CSfC), the NSA's IAD mandates the use of two layers of encryption to protect classified data at rest and data in transit.[40] The NSA's IAD requires use of approved NSA Suite-B algorithms for sufficient encryption to preserve classified data confidentiality and integrity. Therefore, client-side virtualization CDSs employ NSA Suite-B algorithms in dual combinations of either full disk encryption, platform encryption, or file encryption. Such use of encryption provides separation VMMs with an additional layer of isolation between VMs of different security classifications as each VM can use different encryption keys or encryption techniques.

With server-side virtualization in CDS, the separation VMM resides on a computer server set apart from the client machine by one or more networks. The server may store VMs and their data in separate locations from both the client's physical machine and the server. Similar to client-side virtualization, the hypervisor manages the server's physical resources partitioning them among one or more VMs using the aforementioned secure segregation techniques. However, when an authenticated user seeks access to his or her VMs, the server delivers either an image, whole, or part of the VMs across the network(s) through a series of protocols (see Figure 3). The client's machine, using the same protocols, manipulates either the image, whole, or part of the VMs remotely on the server through their physical machine. Hence, with server-side virtualization in CDSs the user must maintain some form of network connectivity with the server to access their VMs.[41]

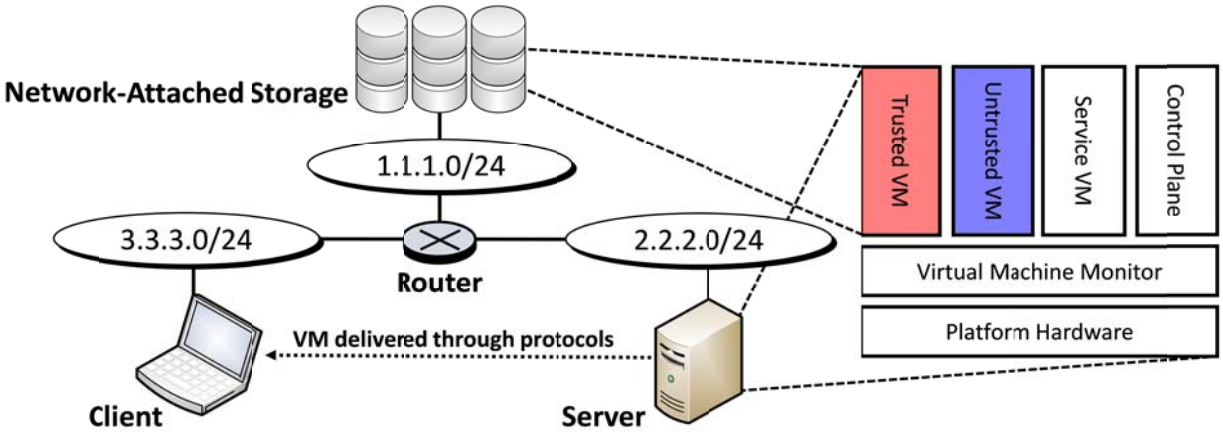

Figure 3. Server-Side Virtualization in Cross-Domain Solutions

In summary, three important factors separate client-side virtualization and server-side virtualization when applying hardware virtualization to handle multiple levels of classified data: 1) the location of the separation VMM; 2) the location of the VMs and their data; and 3) the level of connectivity between the client and server computers. Under client-side virtualization, because the separation VMM is local to the user's physical computer, there is no requirement to connect to a network. A user has the advantage of protection by the separation VMM on their device. VMs respond quickly to commands with local execution. System administrators can also add new users in direct correlation with the number of devices available. Most importantly, a decentralized architecture creates resiliency; that is, disabling one or more physical computers will not disrupt access for the majority of users. However, server-side virtualization CDSs have the advantage of central management. System administrators can control access to data, enforce security policies, and update VM guest software without requiring access to the user's computer. Users need only minimal computing requirements to receive either a desktop image, whole, or part of their VMs (i.e., 'thin' or 'zero' clients). Lastly, system administrators control how VMs interact with each other within the data center.

Technical Overview: Software Defined Networking

Similar to hardware virtualization, SDN abstracts the physical components of a network into a flexible software-based architecture using a combination of programmatic functionality, centralized control, and packet encapsulation.[42] These three features help separate the data plane (i.e., the means by which information moves through the network) from the control plane (i.e., how decisions are made to transfer information) so the network can be reshaped dynamically by the administrator.[43] For the US military, this functionality is crucial in that it provides a way to protect classified data in transit using one physical network while being flexible enough to meet the evolving requirements of a MPE. Further, it removes the burden of physical separation to enforce security because network administrators can logically separate multiple networks at different security levels and thereby reduce the total cost of communication equipment.

In traditional networks, administrators distribute control of how information flows across the many routers and switches of the network, where each device makes independent decisions about where to send and receive packets without regard for the overall topology. To do this, system administrators must program every router and switch to meet the design objectives of their overall network architecture, because no single device has a holistic perspective of how the network should function. While seemingly a simple task, administrators may introduce mistakes when programming individual network devices leading to network outages or introducing security risks in the network's implementation. Further, the network infrastructure can only scale at the speed to which system administrators can acquire, configure, and deploy new equipment. Hence, network administrators are reluctant to change their networks once built because of the sheer complexity of large-scale enterprise networks, which often include multiple firewalls, intrusion detection systems, and many other devices or features. Commercial IT providers and academic

researchers created SDN in response to these factors to allow network administrators greater flexibility and the ability to rapidly provision new networks and services.[44]

With SDN, a central software controller controls all of the network devices within the physical network, similar to the role of the hypervisor in hardware virtualization. The controller accomplishes this by embedding signaling information into the network devices using an SDN protocol, such as OpenFlow, to instruct the network devices in how they should route packets across the network. The network administrator programs the controller to create the desired virtual network topology and the controller implements this design by crafting the necessary signaling information to match the administrator's instructions. The network administrator can use or develop any number of software applications to implement new features into the network through the SDN controller's network operating system. Hence, SDN permits a virtual network topology built at the speed of software over top of the physical network. This functionality matters to the military because changes occur regularly but can implemented much faster through software vice changing the physical infrastructure underneath. Additionally, with any changes, administrators do not need to recreate security policies because they can be carried over by the SDN controller onto any new virtual networks.[45]

To further enhance the integration of hardware virtualization and SDN in data centers, commercial IT providers also proposed various packet encapsulation techniques.[46] As more organizations shifted to cloud computing, cloud infrastructure providers needed to host large numbers of tenants in one or more data centers. In the past, providers accommodated multiple tenants over one physical network using virtual local area networks (VLANs), which segregated tenants into different logical networks using a 12-bit VLAN identifier in each data frame. However, providers can only build 4094 VLANs due to the length of the 12-bit VLAN identifier, which caps

the number of tenants possible. Compounding the issue of hosting thousands of tenants over one physical network, many tenants were also using VMs on their servers which significantly expanded the medium access control (MAC) address tables on "top-of-rack" switches introducing connection delays.[47] Thus, cloud providers created a new framework using a 24-bit identifier, capable of supporting over 16.5 million tenants across a single physical network. Further, these new virtual networks can be spread across multiple physical networks by encapsulating a data frame into a new packet and frame in accordance with the directions of the SDN controller.[48]

The advantage of SDN in CDSs is the increased granularity in control over security. In order for the SDN controller to implement each virtual network, the controller communicates with the hypervisor on every server to coordinate the addition and removal of the server's VMs to the virtual networks. Thus, the SDN controller has direct influence over the network activities of every VM. With this granular control, the network administrator can apply security policies specific to each VM regardless of where that VM is stored or moved within the physical network. This is not possible in traditional networks without excessive cost. In the past network administrators have tied security apparatuses to the physical infrastructure resulting in a strong network perimeter but only limited protections inside. The DoD's policy of physical separation is the most literal interpretation of this philosophy where an air gap between networks of different security classifications segregates information at different security levels. However, once an attacker penetrates the perimeter defenses, there are often little protections which prohibit lateral movement.[49]

In summary, both hardware virtualization and SDN offer the tools to enforce the level of segregation required to implement MLS for both data at rest and data in transit. The private commercial sector and academic researchers have matured these technologies to such a degree

where ignoring them may lead to missed opportunities for the DoD and the IC to reduce cost and build more efficient and secure networks. However, the integration of hardware virtualization and SDN into a single network switch MLS solution has not been explored to a significant degree. Hence, a technical solution is still needed which can encompass both features.

Proposed Solution

Hardware virtualization and SDN provide the means to segregate and secure multiple data link layer enclaves on a single switching platform. Hardware virtualization provides the separation and enforcement for classified data at rest or in execution on the physical server while SDN provides the construction of virtual data link layer networks and separation of classified data in transit. Given the evolution of CD design and the recommendations of the NSA's IAD for CSfC capabilities, the use of encryption in hardware virtualization and SDN is central to the proposed solution as the necessary logical boundary between security levels.[50] However, while encryption has been robustly implemented in CDSs with hardware virtualization to establish multiple levels of security, it has not been developed to same degree with SDN.[51] Hence, the major roadblock impeding a single switch solution is the integration of encryption into SDN to enable authentication, preserve data integrity, and assure data confidentiality between networks. Fortunately, the needed types of encryption protocols to apply SDN in this use case exist but have not yet been integrated into SDN commercial products.[52] Thus, implementation of the proposed solution would require the integration of these encryption protocols into any new certified and accredited SDN product. The following section outlines the general network architecture of the proposed solution, the appropriate use of encryption in the solution, and the design of a network switch to incorporate non-virtualized network devices into the solution's SDN framework.

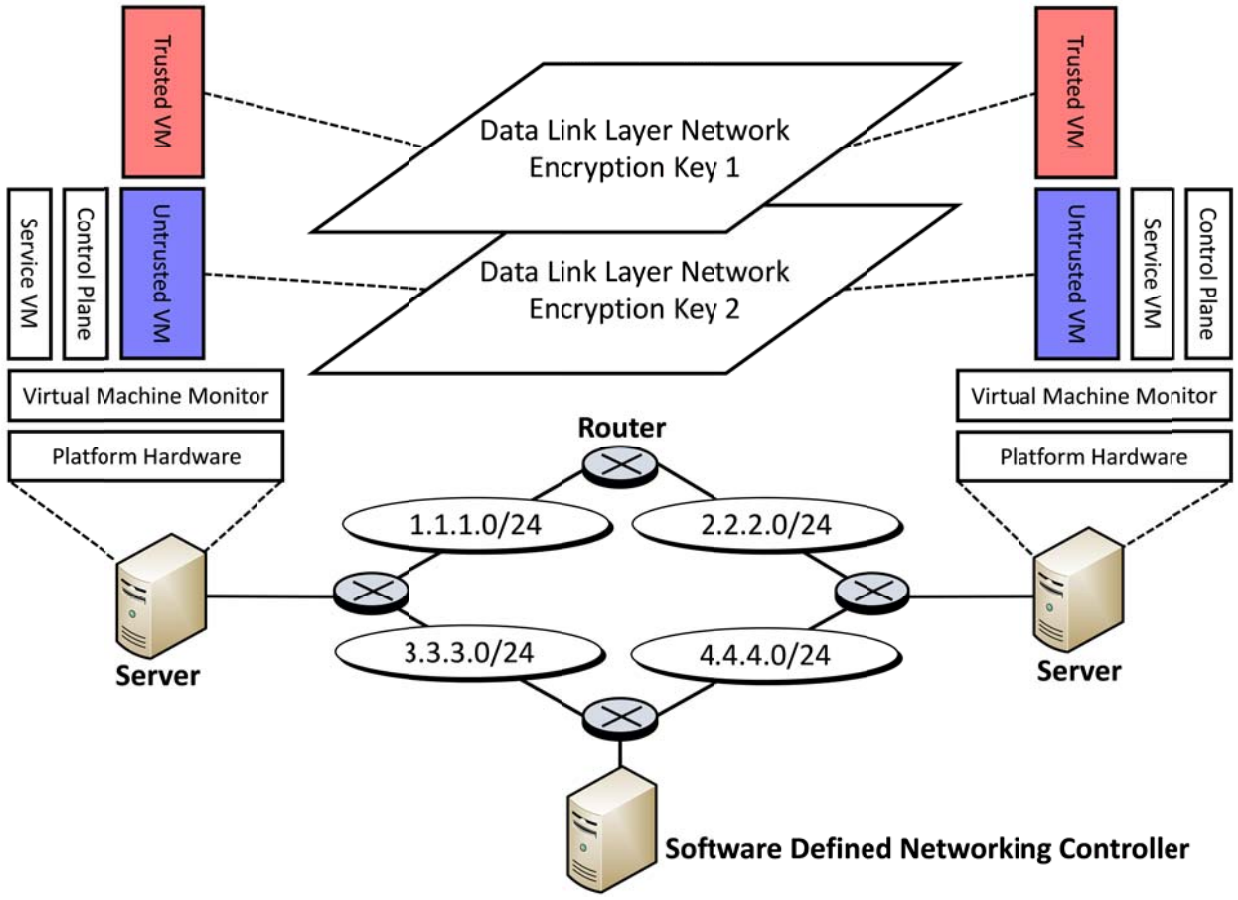

Figure 4. Proposed Network Architecture Design

The general network architecture of the proposed solution rests on building virtual data link layer networks over Internet Protocol (IP) layer networks through a packet encapsulation protocol such as network virtualization using generic routing encapsulation (NVGRE) or the virtual extensible local area network (VXLAN) protocols (see Figure 4).[53] The first element of this architecture is the SDN controller, which resides on the physical network infrastructure and has authority over the control planes of all physical network devices. Through programmatic functions established by the network administrator, the controller instructs the routers and switches through an SDN protocol (e.g., OpenFlow) to build the data link layer networks of different security levels. The controller accomplishes this by assigning a 24-bit virtual network identifica-

tion number to each network as well as establishing the necessary security associations to enable encryption between the designated end points of each network.

Second, the separation VMMs of the connected servers act as the end points of the encrypted tunnels and are responsible for both packet encapsulation and encryption of the VMs's network traffic. The separation VMMs, as the end points, use a shared symmetric encryption key created during the setup of the security associations by the SDN controller to generate the appropriate encrypted tunnels. Finally, with the encrypted tunnels and virtual networks in place, the VMs generate network traffic as normal to communicate with other hosts at the same security level. The VMs are not aware their packets have been encrypted and encapsulated by their separation VMM or that they are traversing the physical IP networks. To the VMs, other hosts appear on the same LAN segment as if they were connected by a network switch.

The principal setback to achieving fully encrypted communication in the proposed network architecture is the number of encryption keys required and their efficient distribution to network end points. Under the NVGRE and VXLAN protocols in order for broadcast traffic to be received by those VMs on the same virtual LAN segment, the separation VMMs use multicast IP routing in the physical network to distribute packets (e.g., Address Resolution Protocol requests).[54] The separation VMMs use multicast IP routing, the process of sending one packet to many but not all users, for broadcast traffic because not every endpoint may have VMs participating in the virtual data link layer network. However, the IETF did not include a method for secure multicast IP routing with encryption as part of their original proposal for the security architecture for the Internet Protocol in Request For Comment (RFC) 2401.[55] RFC 2401 instead focused on securing only unicast communication, the process of sending packets from one host to only one other, with encryption. Consequently, RFC 2401 can apply when one VM communi-

cates directly with another VM through the virtual data link layer network, but it does not suffice when considering VM broadcast traffic. Multicast IP routing requires a separate set of security properties (e.g., Group membership, group secrecy, source authentication, etc.) from secure unicast communication for encryption to work.[56]

Further, the potential combination of using both unicast and multicast encryption escalates the number of encryption keys and management requirements across the network as a whole. Fortunately, Internet researchers have proposed various multicast encryption protocols such as the multicast group security architecture or the multicast encryption security protocol, which can abate key expansion.[57] Most notably, the IETF endorsed the Group Domain of Interpretation (GDOI) in RFC 3547 as an official standard protocol for the distribution of encryption keys to support secure group communications.[58] Commercial IT providers have also successful implemented GDOI in their networking devices as early 2006, which indicates it is feasible for incorporation into SDN products in the future.[59]

For those devices that do not have virtualization technology built into them, a network switch can serve as the gateway into the proposed SDN framework (see Figure 5).[60] The switch would be responsible for the same requirements as the separation VMM, that is: 1) encapsulating and encrypting packets from the non-virtualized device; 2) maintaining a translation table between the address schemes of the physical network and the virtualized data link layer networks; and 3) implementing the commands of the SDN controller in the switch's control plane. Of concern, with the network switch acting as the end point in this specific use case, encryption would not be applied in the link between the switch and the non-virtualized device. To remedy this potential vulnerability, the switch can utilize link layer encryption between the non-virtualized de-

vice and itself, such as the Institute of Electrical and Electronics Engineers' 802.1AE standard for MAC security.[61]

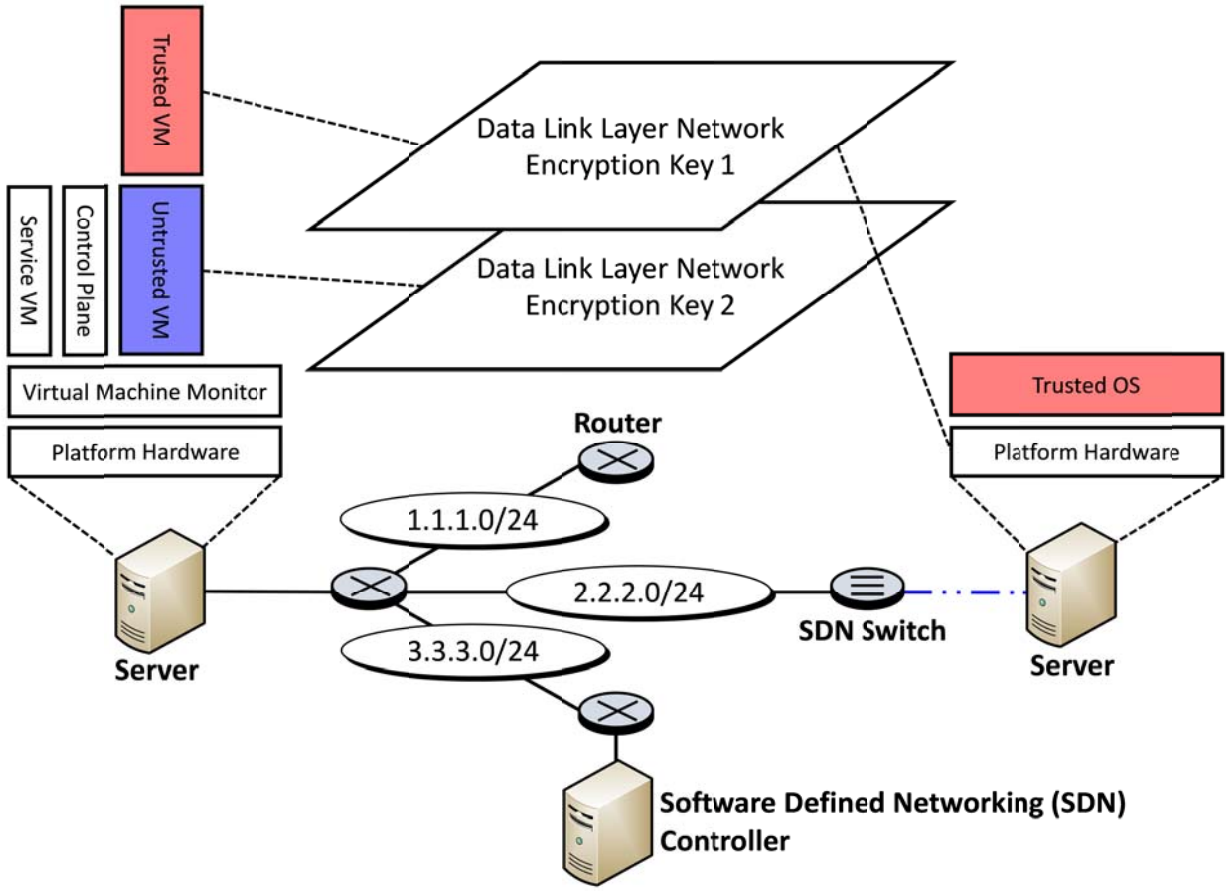

Figure 5. Proposed Network Architecture Design with Software Defined Networking Switch

When done well, this solution means that military network administrators can enjoy the same flexibility as private sector network administrators by rapidly deploying and reconfiguring networks at the speed of software through SDN. Second, this solution enhances security because every separation VMM can apply the security policies of the organization to each individual VM, which is impractical to achieve with physical devices alone because of the overwhelming administrative and financial investment required to match this granularity.[62] Third, the total cost of ownership is drastically reduced through consolidation of the redundant physical networking

components into one data back plane. These benefits hinge on the flexibility of introducing software into the physical network to logical separate multiple levels of security through the power of encryption.

Implications

There are a number of implications associated with the introduction of the proposed solution. First, despite the proposed solution's advantages with regards to flexibility, efficiency, and security, the DoD and IC must not ignore the human element in the advancement of IT. As security researcher Morrie Gasser makes clear in his canonical text, *Building a Secure Computer System*, "the problem [of security] appears to be solely one of people, but it is exacerbated by a technical deficiency of the system."[63] That is, whether building a separate redundant network or using a new type of CDS technology, security policies and procedures must also address the human side of the problem. IT professionals cannot view security strictly as a technical problem as even secure ISs may also contain flaws. Thus, security is fundamentally a human problem and cannot be overcome solely with technical solutions.

Second, while it is possible to consolidate MLS networks at the network switch using hardware virtualization and SDN, adopters should weigh the steep financial and labor investments necessary to implement these new technologies, the required lengthy certification and accreditation process, and the burden of lifecycle support for new devices. Any military service that has not yet made the investment into CD technologies should consider the range of mature CDSs already available. For example, the *SecureView* system, a CDS designed by the AFRL, with VPN concentrators can be deployed immediately to have IP networks with multiple levels of security. *SecureView* also has a robust lifecycle support system already in place, and is shown to reduce to the total cost of ownership relative to alternative CDS.[64]

Third, when considering the state of CD technologies as a whole, the NSA's IAD does not consider the development of new CD access technologies as the most pressing need. Instead, according to the NSA's IAD, future research and development should focus instead on next generation CDS technologies and standards, such as: 1) advanced filtering technologies; 2) standardized remote management and monitoring; 3) standardized data flow configuration; and 4) integrity measurement and attestation.[65]

Once the US military is able to conduct joint operations with the proposed network architecture, the gains will be significant. For one, commanders and analysts will have simultaneous and continuous access to information at multiple levels of security to make informed decisions in a timely manner and they can share information securely and efficiently with coalition partners. For another, network administrators will gain the flexibility and control over their networks unlike ever before, building new networks and deploying services at the speed of software. When taken together, this network architecture means that security and efficiency can complement each other rather than work at odds as has been the struggle of the past. Through hardware virtualization and SDN, warfighters can truly achieve maneuver in the information environment.

End Notes

[1] Ryan Durante and John Woodruff, "SecureView: Government/Industry Collaboration Delivers Improved Levels of Security, Performance, and Cost Savings for Mission-Critical Applications" (Air Force Research Laboratory, December 17, 2012), 1.
[2] John P. L. Woodward, "Applications for Multi Level Secure Operating Systems," in *Managing Requirements Knowledge, International Workshop on*, vol. 0 (Los Alamitos, CA, USA: IEEE Computer Society, 1979), 320.
[3] Myong H. Kang and Ira S. Moskowitz, "A Pump for Rapid, Reliable, Secure Communication," in *Proceedings of the 1st ACM Conference on Computer and Communications Security* (ACM, 1993), 119–29, http://dl.acm.org/citation.cfm?id=168604.
[4] M.H. Kang, I.S. Moskowitz, and S. Chincheck, "The Pump: A Decade of Covert Fun," in *Computer Security Applications Conference, 21st Annual*, 2005, 360–66.
[5] Martin Westphal and Thomas Lang, "Conducting Operations in a Mission Partner Environment," *Joint Forces Quarterly* 74, no. 3rd Quarter (July 1, 2014): 47–48.
[6] National Computer Security Center, *A Guide to Understanding Covert Channel Analysis of Trusted Systems* (DIANE Publishing, 1994), 5.
[7] Joint Staff J6 Deputy Director Cyber and C4 Integration, Interoperability and Integration Division, "Mission Partner Environment (U.S. Contribution to FMN): Multi-National Maritime Information Services Interoperability (M2I2) Board 15-2," (PowerPoint presentation, September 7, 2015), 23.
[8] John Allen, "Directive to Utilize the Afghanistan Mission Network (AMN)" (Headquarters International Security Assistance Force, March 2012), 1.
[9] Westphal and Lang, "Conducting Operations in a Mission Partner Environment," 46.
[10] Allen, "Directive to Utilize the Afghanistan Mission Network (AMN)."
[11] Westphal and Lang, "Conducting Operations in a Mission Partner Environment," 47.
[12] Joint Staff J6 Deputy Director Cyber and C4 Integration, Interoperability and Integration Division, "Mission Partner Environment (U.S. Contribution to FMN): Multi-National Maritime Information Services Interoperability (M2I2) Board 15-2," 33–44.
[13] Joint Staff J6 Deputy Director Cyber and C4 Integration, Interoperability and Integration Division, "Mission Partner Environment (U.S. Contribution to FMN): Multi-National Maritime Information Services Interoperability (M2I2) Board 15-2," 42–43.
[14] Westphal and Lang, "Conducting Operations in a Mission Partner Environment," 47.
[15] Westphal and Lang, "Conducting Operations in a Mission Partner Environment," 47.
[16] Unified Cross Domain Management Office, "The Unified Cross Domain Management Office: Bridging Security Domains and Cultures" (Washington, DC: OASD(IIA) Unified Cross Domain (CD) Management Office (UCDMO), July 2008), 21, http://www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA487191; Woodward, "Applications for Multi Level Secure Operating Systems," 319.
[17] Teresa Takai and Al Tarasiuk, "Establishment of the Unified Cross Domain Services Management Office as the Cross Domain Requirements and Engineering Services Manager" (Department of Defense, Office of the Director of National Intelligence, March 26, 2014).
[18] John Rushby and Rance DeLong, "MILS Integration Protection Profile" (PDF presentation, Menlo Park, CA, n.d.), 2, http://www.csl.sri.com/users/rushby/slides/mipp-jan07.pdf.
[19] Rushby and DeLong, "MILS Integration Protection Profile," 2; Terry Halvorsen, "Department of Defense Instruction 8540.01: Cross Domain (CD) Policy" (Department of Defense, May 8, 2015), 30–43, http://www.dtic.mil/whs/directives/corres/pdf/854001p.pdf.
[20] Unified Cross Domain Management Office, "The Unified Cross Domain Management Office: Bridging Security Domains and Cultures," 21, 23.
[21] Unified Cross Domain Management Office, "The Unified Cross Domain Management Office: Bridging Security Domains and Cultures," 21.
[22] Woodward, "Applications for Multi Level Secure Operating Systems."
[23] J. M. Rushby, "Design and Verification of Secure Systems," in *Proceedings of the Eighth ACM Symposium on Operating Systems Principles*, SOSP '81 (New York, NY, USA: ACM, 1981), 12–21, http://doi.acm.org/10.1145/800216.806586.

[24] Systems and Network Analysis Center, Information Assurance Directorate, "Separation Kernels on Commodity Workstations" (National Security Agency, March 11, 2010), 3, http://www.niap-ccevs.org/announcements/Separation%20Kernels%20on%20Commodity%20Workstations.pdf.
[25] National Security Agency, "Separation Kernel Protection Profile Sunset Q&A," accessed March 27, 2016, http://www.niap-ccevs.org/announcements/SKPP%20Sunset%20Q&A.pdf.
[26] John McDermott et al., "Separation Virtual Machine Monitors," in *Proceedings of the 28th Annual Computer Security Applications Conference*, ACSAC '12 (New York, NY, USA: ACM, 2012), 419–28, http://doi.acm.org/10.1145/2420950.2421011.
[27] "Xen Project: Open Source Hypervisor, High Performance Clouds, Free Software, Open Source Virtualization, Open Virtualization, Cloud Computing, Desktop Virtualization, Mobile Virtualization, Embedded Virtualization," accessed March 26, 2016, http://www.xenproject.org/.
[28] A. Lara, A. Kolasani, and B. Ramamurthy, "Network Innovation Using OpenFlow: A Survey," *IEEE Communications Surveys Tutorials* 16, no. 1 (February 2014): 493–512.
[29] Xiong Liu et al., "Design of the Multi-Level Security Network Switch System Which Restricts Covert Channel," in *2011 IEEE 3rd International Conference on Communication Software and Networks (ICCSN)*, 2011, 233–37.
[30] Mallik Mahalingam et al., *Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks*, Request for Comments 7348 (RFC Editor, 2014), https://rfc-editor.org/rfc/rfc7348.txt; Thomas Hardjono and Brian Weis, *The Multicast Group Security Architecture*, Request for Comments 3740 (RFC Editor, 2004), https://rfc-editor.org/rfc/rfc3740.txt; Thomas Hardjono, Sheela Rowles, and Brian Weis, *The Group Domain of Interpretation*, Request for Comments 6407 (RFC Editor, 2015), https://rfc-editor.org/rfc/rfc6407.txt; Randall Atkinson and Dr Stephen T. Kent, *Security Architecture for the Internet Protocol*, Request for Comments 4301 (RFC Editor, 2005), https://rfc-editor.org/rfc/rfc2401.txt; Yu-Shun Wang and Pankaj Garg, *NVGRE: Network Virtualization Using Generic Routing Encapsulation*, Request for Comments 7637 (RFC Editor, 2015), https://rfc-editor.org/rfc/rfc7637.txt; Lawrence Miller, *Micro-segmentation For Dummies, VMware Special Edition* (John Wiley & Sons, Inc., 2015), https://horizonworkspace.vmware.com/data/shf/WaZH4E4lQMu5pGtDvp97MQAFMTA0NTUA.
[31] John Rushby, "A Trusted Computing Base for Embedded Systems," in *Proceedings 7th DoD/NBS Computer Security Conference* (Citeseer, 1984), 6, http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.23.5523&rep=rep1&type=pdf.
[32] Robert P. Goldberg, "Architectural Principles for Virtual Computer Systems" (Springfield, Virginia: National Technical Information Service, US Department of Commerce, February 1973), 15–16.
[33] Jim Smith and Ravi Nair, *Virtual Machines: Versatile Platforms for Systems and Processes* (Elsevier, 2005), 12–13.
[34] Cynthia E. Irvine and John Scott Robin, "Analysis of the Intel Pentium's Ability to Support a Secure Virtual Machine Monitor," August 2000, 138–139, https://www.usenix.org/legacy/events/sec00/full_papers/robin/robin.pdf.
[35] Irvine and Robin, "Analysis of the Intel Pentium's Ability to Support a Secure Virtual Machine Monitor," 129–132.
[36] McDermott et al., "Separation Virtual Machine Monitors," 420.
[37] McDermott et al., "Separation Virtual Machine Monitors," 420.
[38] McDermott et al., "Separation Virtual Machine Monitors," 422.
[39] McDermott et al., "Separation Virtual Machine Monitors," 421–422.
[40] National Security Agency Information Assurance Directorate, "Virtual Private Network Capability Package," August 20, 2015, https://www.nsa.gov/ia/_files/VPN_CP_3_2.pdf; National Security Agency Information Assurance Directorate, "Commercial Solutions for Classified Data-at-Rest Capability Package," April 2, 2015, https://www.nsa.gov/ia/_files/DAR_CP_v2.0.pdf.
[41] Thomas D. Nadeau and Ken Gray, *SDN: Software Defined Networks* (O'Reilly Media, Inc., 2013), 157–183; Keir A. Fraser et al., "The Xenoserver Computing Infrastructure" (Technical Report UCAM-CL-TR-552, University of Cambridge, Computer Laboratory, 2003), https://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-552.pdf.
[42] Nadeau and Gray, *SDN*, 1–8.
[43] Nadeau and Gray, *SDN*, 5.
[44] Lara, Kolasani, and Ramamurthy, "Network Innovation Using OpenFlow"; Nick McKeown et al., "OpenFlow: Enabling Innovation in Campus Networks," *SIGCOMM Comput. Commun. Rev.* 38, no. 2 (March 2008): 69–74.
[45] Nadeau and Gray, *SDN*, 117–156.

[46] Mahalingam et al., *Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks*; Wang and Garg, *NVGRE: Network Virtualization Using Generic Routing Encapsulation*.
[47] Mahalingam et al., *Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks*, 6.
[48] Mahalingam et al., *Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks*; Wang and Garg, *NVGRE: Network Virtualization Using Generic Routing Encapsulation*; Nadeau and Gray, *SDN*, 163–205.
[49] Miller, *Micro-segmentation For Dummies, VMware Special Edition*, 20–24.
[50] National Security Agency Information Assurance Directorate, "Virtual Private Network Capability Package"; National Security Agency Information Assurance Directorate, "Commercial Solutions for Classified Data-at-Rest Capability Package."
[51] Durante and Woodruff, "SecureView: Government/Industry Collaboration Delivers Improved Levels of Security, Performance, and Cost Savings for Mission-Critical Applications," 5.
[52] Roger Fortier, "Security for the New Battlefield," *The Network Virtualization Blog*, November 30, 2015, http://blogs.vmware.com/networkvirtualization/2015/11/security-for-the-new-battlefield.html; Marcia Savage, "VMware Expanding NSX Security | Network Computing," accessed April 2, 2016, http://www.networkcomputing.com/networking/vmware-expanding-nsx-security/1030784217.
[53] Mahalingam et al., *Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks*; Wang and Garg, *NVGRE: Network Virtualization Using Generic Routing Encapsulation*.
[54] Mahalingam et al., *Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks*, 7–9; Wang and Garg, *NVGRE: Network Virtualization Using Generic Routing Encapsulation*, 9, 13–14.
[55] Atkinson and Kent, *Security Architecture for the Internet Protocol*, 12.
[56] E. Ali, T. El-fouly, and A. Badr, "MESP: A Modified IPSec for Secure Multicast Communication," in *2006 6th International Conference on ITS Telecommunications Proceedings*, 2006, 812.
[57] Hardjono and Weis, *The Multicast Group Security Architecture*; Ali, El-fouly, and Badr, "MESP."
[58] Hardjono, Rowles, and Weis, *The Group Domain of Interpretation*, 1.
[59] Cisco Systems, "Cisco IOS Secure Multicast," 2006, http://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/ip-multicast/prod_white_paper0900aecd8047191e.pdf.
[60] Mahalingam et al., *Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks*, 18.
[61] Institute of Electrical and Electronics Engineers, "802.1AE - Media Access Control (MAC) Security," 2006, http://standards.ieee.org/findstds/standard/802.1AE-2006.html.
[62] Miller, *Micro-segmentation For Dummies, VMware Special Edition*, 17–22.
[63] Morrie Gasser, *Building a Secure Computer System* (New York, New York: Van Nostrand Reinhold Company New York, 1988), 12, http://iser.ruc.edu.cn/page/team/wshi/readings/CS02.pdf.
[64] Durante and Woodruff, "SecureView: Government/Industry Collaboration Delivers Improved Levels of Security, Performance, and Cost Savings for Mission-Critical Applications."
[65] Boyd Fletcher, "Next Generation CDS Technologies," (PDF presentation, October 9, 2015), 2, 4.

Bibliography

Ali, E., T. El-fouly, and A. Badr. “MESP: A Modified IPSec for Secure Multicast Communication.” In *2006 6th International Conference on ITS Telecommunications Proceedings*, 812–16, 2006.

Allen, John. “Directive to Utilize the Afghanistan Mission Network (AMN).” Headquarters International Security Assistance Force, March 2012.

Atkinson, Randall, and Dr Stephen T. Kent. *Security Architecture for the Internet Protocol*. Request for Comments 4301. RFC Editor, 2005. https://rfc-editor.org/rfc/rfc2401.txt.

Cisco Systems. “Cisco IOS Secure Multicast,” 2006. http://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/ip-multicast/prod_white_paper0900aecd8047191e.pdf.

Durante, Ryan, and John Woodruff. “SecureView: Government/Industry Collaboration Delivers Improved Levels of Security, Performance, and Cost Savings for Mission-Critical Applications.” Air Force Research Laboratory, December 17, 2012.

Fletcher, Boyd. “Next Generation CDS Technologies.” PDF presentation, October 9, 2015.

Fortier, Roger. “Security for the New Battlefield.” *The Network Virtualization Blog*, November 30, 2015. http://blogs.vmware.com/networkvirtualization/2015/11/security-for-the-new-battlefield.html.

Fraser, Keir A., Steven M. Hand, Timothy L. Harris, Ian M. Leslie, and Ian A. Pratt. “The Xenoserver Computing Infrastructure.” Technical Report UCAM-CL-TR-552, University of Cambridge, Computer Laboratory, 2003. https://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-552.pdf.

Gasser, Morrie. *Building a Secure Computer System*. New York, New York: Van Nostrand Reinhold Company New York, 1988. http://iser.ruc.edu.cn/page/team/wshi/readings/CS02.pdf.

Goldberg, Robert P. “Architectural Principles for Virtual Computer Systems.” Springfield, Virginia: National Technical Information Service, US Department of Commerce, February 1973.

Halvorsen, Terry. “Department of Defense Instruction 8540.01: Cross Domain (CD) Policy.” Department of Defense, May 8, 2015. http://www.dtic.mil/whs/directives/corres/pdf/854001p.pdf.

Hardjono, Thomas, Sheela Rowles, and Brian Weis. *The Group Domain of Interpretation*. Request for Comments 6407. RFC Editor, 2015. https://rfc-editor.org/rfc/rfc6407.txt.

Hardjono, Thomas, and Brian Weis. *The Multicast Group Security Architecture*. Request for Comments 3740. RFC Editor, 2004. https://rfc-editor.org/rfc/rfc3740.txt.

Institute of Electrical and Electronics Engineers. "802.1AE - Media Access Control (MAC) Security," 2006. http://standards.ieee.org/findstds/standard/802.1AE-2006.html.

Irvine, Cynthia E., and John Scott Robin. "Analysis of the Intel Pentium's Ability to Support a Secure Virtual Machine Monitor," August 2000. https://www.usenix.org /legacy/events/sec00/full_papers/robin/robin.pdf.

Joint Staff J6 Deputy Director Cyber and C4 Integration, Interoperability and Integration Division. "Mission Partner Environment (U.S. Contribution to FMN): Multi-National Maritime Information Services Interoperability (M2I2) Board 15-2." PowerPoint presentation, September 7, 2015.

Kang, M.H., I.S. Moskowitz, and S. Chincheck. "The Pump: A Decade of Covert Fun." In *Computer Security Applications Conference, 21st Annual*, 360–66, 2005.

Kang, Myong H., and Ira S. Moskowitz. "A Pump for Rapid, Reliable, Secure Communication." In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, 119–29. ACM, 1993. http://dl.acm.org/citation.cfm?id=168604.

Lara, A., A. Kolasani, and B. Ramamurthy. "Network Innovation Using OpenFlow: A Survey." *IEEE Communications Surveys Tutorials* 16, no. 1 (February 2014): 493–512.

Liu, Xiong, Haiwei Xue, Xiaoping Feng, and Yiqi Dai. "Design of the Multi-Level Security Network Switch System Which Restricts Covert Channel." In *2011 IEEE 3rd International Conference on Communication Software and Networks (ICCSN)*, 233–37, 2011.

Mahalingam, Mallik, T. Sridhar, Mike Bursell, Lawrence Kreeger, Chris Wright, Kenneth Duda, Puneet Agarwal, and Dinesh Dutt. *Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks*. Request for Comments 7348. RFC Editor, 2014. https://rfc-editor.org/rfc/rfc7348.txt.

McDermott, John, Bruce Montrose, Margery Li, James Kirby, and Myong Kang. "Separation Virtual Machine Monitors." In *Proceedings of the 28th Annual Computer Security Applications Conference*, 419–28. ACSAC '12. New York, NY, USA: ACM, 2012. http://doi.acm.org/10.1145/2420950.2421011.

McKeown, Nick, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. "OpenFlow: Enabling Innovation in Campus Networks." *SIGCOMM Comput. Commun. Rev.* 38, no. 2 (March 2008): 69–74.

Miller, Lawrence. *Micro-segmentation For Dummies, VMware Special Edition*. John Wiley & Sons, Inc., 2015. https://horizonworkspace.vmware.com /data/shf/WaZH4E4lQMu5pGtDvp97MQAFMTA0NTUA.

Nadeau, Thomas D., and Ken Gray. *SDN: Software Defined Networks*. O'Reilly Media, Inc., 2013.

National Computer Security Center. *A Guide to Understanding Covert Channel Analysis of Trusted Systems*. DIANE Publishing, 1994.

National Security Agency. "Separation Kernel Protection Profile Sunset Q&A." Accessed March 27, 2016. http://www.niap-ccevs.org/announcements/SKPP%20Sunset%20Q&A.pdf.

National Security Agency Information Assurance Directorate. "Commercial Solutions for Classified Data-at-Rest Capability Package," April 2, 2015. https://www.nsa.gov/ia/_files/DAR_CP_v2.0.pdf.

———. "Virtual Private Network Capability Package," August 20, 2015. https://www.nsa.gov/ia/_files/VPN_CP_3_2.pdf.

Rushby, J. M. "Design and Verification of Secure Systems." In *Proceedings of the Eighth ACM Symposium on Operating Systems Principles*, 12–21. SOSP '81. New York, NY, USA: ACM, 1981. http://doi.acm.org/10.1145/800216.806586.

Rushby, John. "A Trusted Computing Base for Embedded Systems." In *Proceedings 7th DoD/NBS Computer Security Conference*, 294–311. Citeseer, 1984. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.23.5523&rep=rep1&type=pdf.

Rushby, John, and Rance DeLong. "MILS Integration Protection Profile." PDF presentation, Menlo Park, CA, n.d. http://www.csl.sri.com/users/rushby/slides/mipp-jan07.pdf.

Savage, Marcia. "VMware Expanding NSX Security | Network Computing." Accessed April 2, 2016. http://www.networkcomputing.com/networking/vmware-expanding-nsx-security/1030784217.

Smith, Jim, and Ravi Nair. *Virtual Machines: Versatile Platforms for Systems and Processes*. Elsevier, 2005.

Systems and Network Analysis Center, Information Assurance Directorate. "Separation Kernels on Commodity Workstations." National Security Agency, March 11, 2010. http://www.niap-ccevs.org/announcements/Separation%20Kernels%20on%20Commodity%20Workstations.pdf.

Takai, Teresa, and Al Tarasiuk. "Establishment of the Unified Cross Domain Services Management Office as the Cross Domain Requirements and Engineering Services Manager." Department of Defense, Office of the Director of National Intelligence, March 26, 2014.

Unified Cross Domain Management Office. "The Unified Cross Domain Management Office: Bridging Security Domains and Cultures." Washington, DC: OASD(IIA) Unified Cross Domain (CD) Management Office (UCDMO), July 2008. http://www.dtic.mil

/cgi-bin/GetTRDoc?AD=ADA487191.

Wang, Yu-Shun, and Pankaj Garg. *NVGRE: Network Virtualization Using Generic Routing Encapsulation*. Request for Comments 7637. RFC Editor, 2015. https://rfc-editor.org/rfc/rfc7637.txt.

Westphal, Martin, and Thomas Lang. “Conducting Operations in a Mission Partner Environment.” *Joint Forces Quarterly* 74, no. 3rd Quarter (July 1, 2014): 44–49.

Woodward, John P. L. “Applications for Multi Level Secure Operating Systems.” In *Managing Requirements Knowledge, International Workshop on*, 0:319–28. Los Alamitos, CA, USA: IEEE Computer Society, 1979.

“Xen Project: Open Source Hypervisor, High Performance Clouds, Free Software, Open Source Virtualization, Open Virtualization, Cloud Computing, Desktop Virtualization, Mobile Virtualization, Embedded Virtualization.” Accessed March 26, 2016. http://www.xenproject.org/.