



**NAVAL
POSTGRADUATE
SCHOOL**

MONTEREY, CALIFORNIA

THESIS

**APPLYING U.S. MILITARY CYBERSECURITY
POLICIES TO CLOUD ARCHITECTURES**

by

Michael K. Atadika

September 2018

Thesis Advisor:
Second Reader:

Karen L. Burke
Neil C. Rowe

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 September 2018	3. REPORT TYPE AND DATES COVERED Master's thesis	
4. TITLE AND SUBTITLE APPLYING U.S. MILITARY CYBERSECURITY POLICIES TO CLOUD ARCHITECTURES			5. FUNDING NUMBERS	
6. AUTHOR(S) Michael K. Atadika				
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.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release. Distribution is unlimited.			12b. DISTRIBUTION CODE A	
13. ABSTRACT (maximum 200 words) The Department of Defense (DoD) has accelerated its adoption of cloud technologies, which come with inherent risks. This thesis investigated four important cybersecurity issues that the DoD must address: customer misconfigurations, data leaks, complications in security controls, and necessary changes to digital forensic incident-response tactics. We examined current U.S. policy documents and found a number of issues that need to be clarified for contracting with cloud service providers. Human misunderstandings largely drive cloud misconfigurations, which eventually become cloud data spills that require a digital forensic incident-response. To prevent misconfigurations, it is essential that DoD staff receive continual in-depth cloud training and that the DoD redefines the roles for virtualized cloud architectures. Fortunately, the selection of the cloud service model can highlight which cloud layers the DoD is responsible for, and therefore which security controls to implement. Federal cloud computing policy, DoD FedRAMP+, specifies the security controls needed based on the sensitivity of the data. However, once a cyber-incident is declared, digital forensics analysts confront a myriad of cloud-specific technological, legal, and boundary challenges. The security vulnerabilities must be considered during a transformational migration from on-premises architectures to cloud technologies. This thesis offers recommendations to address these vexing cybersecurity issues.				
14. SUBJECT TERMS cloud, incident response, misconfigurations, cybersecurity, cloud service providers, DoD cyber incident handling, Cloud Executive Steering Group, CESG, cloud spillage, U.S. Department of Defense, data leaks, DFIR, incident-response tactics, <i>Cloud Computing Security Requirements Guide</i>			15. NUMBER OF PAGES 101	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU	

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release. Distribution is unlimited.

**APPLYING U.S. MILITARY CYBERSECURITY POLICIES TO CLOUD
ARCHITECTURES**

Michael K. Atadika
Civilian, Federal Cyber Corps
BA, University of Virginia, 2001
MBA, Columbia University, 2008

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN COMPUTER SCIENCE

from the

**NAVAL POSTGRADUATE SCHOOL
September 2018**

Approved by: Karen L. Burke
Advisor

Neil C. Rowe
Second Reader

Peter J. Denning
Chair, Department of Computer Science

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

The Department of Defense (DoD) has accelerated its adoption of cloud technologies, which come with inherent risks. This thesis investigated four important cybersecurity issues that the DoD must address: customer misconfigurations, data leaks, complications in security controls, and necessary changes to digital forensic incident-response tactics. We examined current U.S. policy documents and found a number of issues that need to be clarified for contracting with cloud service providers. Human misunderstandings largely drive cloud misconfigurations, which eventually become cloud data spills that require a digital forensic incident-response. To prevent misconfigurations, it is essential that DoD staff receive continual in-depth cloud training and that the DoD redefines the roles for virtualized cloud architectures. Fortunately, the selection of the cloud service model can highlight which cloud layers the DoD is responsible for, and therefore which security controls to implement. Federal cloud computing policy, DoD FedRAMP+, specifies the security controls needed based on the sensitivity of the data. However, once a cyber-incident is declared, digital forensics analysts confront a myriad of cloud-specific technological, legal, and boundary challenges. The security vulnerabilities must be considered during a transformational migration from on-premises architectures to cloud technologies. This thesis offers recommendations to address these vexing cybersecurity issues.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
A.	BACKGROUND	1
B.	BALANCING ACCELERATED CLOUD ADOPTION WITH CRUCIAL RISK MANAGEMENT DECISIONS.....	3
	1. Customer Misconfigurations.....	4
	2. Cloud Leaks.....	5
	3. Complications in the Implementations of Security Controls.....	6
	4. Digital Forensic Incident-Response.....	7
C.	CHAPTER SUMMARY AND THESIS ORGANIZATION	9
II.	CLOUD COMPUTING PRINCIPLES	11
A.	PRELIMINARY ORIENTATION.....	11
B.	DEFINITIONS	13
	1. Cloud Computing.....	13
	2. Cloud Security	17
C.	THE CIA TRIAD AND THE INFORMATION SECURITY EQUATION.....	18
D.	DISTINGUISHING CLOUD-SPECIFIC VULNERABILITIES FROM ON-PREMISES SYSTEMS.....	19
E.	COMPARING CLOUD-SPECIFIC THREATS FROM ON- PREMISES SYSTEM THREATS	21
F.	DISTINGUISHING CLOUD-SPECIFIC COUNTERMEASURES	25
III.	TECHNICAL ISSUES IN CLOUD COMPUTING	27
A.	SERVICE MODEL SELECTION: COMPARATIVE RISK FACTORS	27
B.	SHARED RESPONSIBILITY MODEL.....	32
	1. Why Lift and Shift Does Not Work.....	34
	2. Cloud Logical Layer Commonalities	35
C.	CLOUD DIGITAL FORENSIC INCIDENT-RESPONSE CHALLENGES.....	37
	1. Technical Challenges	43
	2. Legal Challenges	53
	3. Boundary Challenges.....	55

IV.	ADAPTING DoD REQUIREMENTS FOR CLOUD ARCHITECTURES.....	57
A.	FedRAMP	57
B.	DoD CLOUD COMPUTING SECURITY REQUIREMENTS GUIDE	59
1.	FedRAMP+.....	59
2.	FedRAMP+ Security Controls/Enhancements	60
3.	Impact Levels	61
4.	Data Spills.....	63
5.	Incident Response	64
V.	RISK MANAGEMENT ALTERNATIVES: TRANSFORMATIONAL MIGRATION	67
A.	COLLOCATING DATA LOCATION.....	67
B.	RECOMMENDATIONS ADDRESSING CUSTOMER MISCONFIGURATIONS.....	68
C.	RECOMMENDATIONS ADDRESSING CLOUD LEAKS	69
1.	Migrating Security Perimeter	70
2.	Migrating Workflow Processes.....	70
3.	Retraining Workforces for Cloud Computing.....	71
D.	RECOMMENDATIONS ADDRESSING COMPLICATIONS IN THE IMPLEMENTATION OF SECURITY CONTROLS.....	72
E.	RECOMMENDATIONS ADDRESSING DIGITAL FORENSIC INCIDENT-RESPONSE	72
VI.	CONCLUSION	75
	LIST OF REFERENCES.....	77
	INITIAL DISTRIBUTION LIST	83

LIST OF FIGURES

Figure 1.	Examples of Cloud Consumer Services. Source: Liu et al. (2011).	12
Figure 2.	NIST Cloud Computing Model. Source: CSA (2017).....	13
Figure 3.	Public versus Private Cloud. Source: Odell et al. (2015).	15
Figure 4.	Consumer Service Model Security Burden Continuum. Source: Cloud Security Alliance (CSA) (2017).....	28
Figure 5.	Cloud Reference Service Model Stack. Source: CSA (2017).....	28
Figure 6.	Cloud Security Responsibility Matrix (Low Differentiation). Source: Gordon (2016).....	33
Figure 7.	Cloud Security Responsibility Matrix (On-premises Application). Source: DISA (2017).	35
Figure 8.	Cloud Security Alliance Logical Model. Source: CSA (2017).....	36
Figure 9.	Sixty-Five Cloud Challenges Aggregated into Nine Primary Categories and Relevant Sub-categories. Source: NCC FSWG (2014).....	39
Figure 10.	Primary Categories Intersection (Related Categories). Source: NCC FSWG (2014).....	42
Figure 11.	Cloud Security Responsibility Matrix (On-premises Application). Source: DISA (2017)	44
Figure 12.	Requirements to Create an AWS Account. Source: Amazon Web Services (2015).	46
Figure 13.	Impact Level Comparison. Source: DISA (2017).	62

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

Table 1.	Treacherous 12 Threats Summary. Adapted from CSA (2016).	21
Table 2.	Sixty-Five Challenges by Primary Categories. Adapted from NCC FSWG (2014).....	38
Table 3.	Sixty-Five Challenges by Three Broad Classes. Adapted from NCC FSWG (2014).....	41
Table 4.	Technological Challenges Dominated by Data Collection and Architecture Primary Categories. Source: NCC FSWG (2014).	43
Table 5.	DoD FedRAMP+ Incident Response Security Controls. Adapted from Metheny (2013).....	65

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

APT	advanced persistent threat
API	application programming interface
AWS	Amazon Web Services
CAO	chief acquisition officer
CIA	Central Intelligence Agency
CIO	chief information officer
CNSS	Committee on National Security Systems
CSA	Cloud Security Alliance
DDos	distributed denial of service
DHS	Department of Homeland Security
DISA	Defense Information Systems Agency
DoD	Department of Defense
ENISA	European Network and Information Security Agency
FedRAMP	Federal Risk and Authorization Management Program
FISMA	Federal Information Security Management Act
IEEE	Institute of Electrical and Electronics Engineers
IT	information technology
NCC FSWG	NIST Cloud Computing Forensic Science Working Group
NIST	National Institute of Standards and Technology
NSA	National Intelligence Agency
SP	Special Publication (NIST)
VM	virtual machine

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

I want to begin by giving thanks to God; by His grace, all things are possible.

To my wife, Salama, for her tireless love, support, and perseverance: Thank you for allotting me 27 months across the country to study computer science while you shouldered all the responsibilities—most importantly, raising our son. To my son, Roman, I am grateful for your inspiration and for the new ways you discover how to make me laugh each and every day. To my mother, for demonstrating unconditional love and instilling faith and hope in my DNA. To my father, for his personal sacrifices and for modeling a deliberate life of daily self-discipline. To my sisters, Selassie and Novisi, for setting the sibling bar so darn high.

To my friends and extended family, especially Sophronia Spicer and Michael Nirschl, I am grateful for your love, prayers, and relationships. To Roman's godfather, Jonathan Levirne, what can I say? You *are* the Godfather. Thank you. To the Harris family, thank you for providing cover and my home away from home in Livermore. To my mentors Kevin Meenan, Kevin Chavers, and Bryan Lopez, thank you for your graciousness and wisdom.

To my thesis advisor, Professor Burke, thank you for your consistent guidance, beginning with navigating the selection of my thesis topic through final submission. Your recommendations and responsiveness were invaluable. To my second reader, Dr. Rowe, thank you for taking me on as an advisee and for your insistence on rigorous examination and clear explanations. Also, thank you for breaking out the Caymus. To Graduate Writing Center Coaches Alison and Cheryldee, thank you for helping me organize, revise, and express my ideas. Your availability and personal investment were legendary. To my thesis processor, Michele, thank you for providing all your edits and reducing a daunting task right down to size. To the National Science Foundation, for the opportunity to pursue my master's in computer science through the Scholarship for Service. To the Naval Postgraduate School, my professors, and my peers, especially Johnathan Mooring, for creating an enriching environment and my most rewarding educational experience. Special thanks to Chris Eagle, Chris Frenzen, Ty Waller, and Liz Wanic.

To all, I am grateful.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

A. BACKGROUND

Cloud computing, while still in its nascence, has resoundingly demonstrated the benefits associated with its disruptive, transformative, and economic characteristics. Cloud computing is attractive because it offers organizations the flexibility to reduce hardware acquisition costs as well as the opportunity to save money with its pay-per-use pricing scheme. However, cloud computing is not without risks. The confidentiality of data stored in cloud environments remains a major concern, especially because cloud environments are typically multi-tenant, meaning they are shared by multiple clients. Cloud data is also stored remotely, via an internet connection, and this separation warrants additional assurances about data security. The Department of Defense (DoD) is exploring an aggressive transition strategy to best exploit cloud computing benefits while managing its risks.

The primary catalyst for cloud migration, for an overwhelming majority of organizations, is cost savings. ISACA, a global nonprofit thought-leader focused on information system and information technology (IT) governance, clearly lays out the drivers and detractors of cloud adoption, and the prominent risk management tradeoffs between costs and benefits:

As with any emerging technology, cloud computing offers the possibility of high reward in terms of containment of costs and features such as agility and provisioning speed. However, it also brings the potential for unknown and potentially high risk. Cloud computing introduces a level of abstraction between the physical infrastructure and the owner of the information being stored and processed. Traditionally, the data owner has had direct or indirect control of the physical environment affecting his/her data. In the cloud, this is no longer the case. Due to this abstraction, there is already a widespread demand for greater transparency and a robust assurance approach of the cloud computing supplier's security and control environment. (ISACA, 2009, p. 4)

The most pertinent considerations in this description are costs, abstraction, risk, control, environment, transparency, assurance, and security.

To illuminate the significant differences between a cloud computing system and a pre-cloud computer system, we must first define a few terms. For example, an enterprise information system is one that provides functionality (e.g., email, security, file systems) that is integrated for use by an entire organization. In contrast, a specialized system that is used only by a small enclave of organizational users is not considered an enterprise system. Migration is the term used when an organization transitions applications from an on-premises enterprise information system to one with cloud computing capabilities. An on-premises system is one that predates cloud computing and is defined by (i) applications that have not migrated or may never migrate to transmit data between the pre-cloud computing information system and the cloud, (ii) the distinctive features it lacks in comparison to those exhibited by cloud computing. On-premises systems only have access to their own organizational resources; they cannot access a pool of computing resources that are shared by users external to the organization, as a cloud computing system can. Therefore, on-premises systems are unable to unilaterally increase or decrease use of the shared, pooled resources, as cloud computing can. If an on-premises system needs additional resources, the organization would have to purchase the resource (e.g., a server), and then wait for delivery and the subsequent configuration to its system. An on-premises system also does not give the organization the ability to pay only for the resources actually consumed, like how utilities (e.g., water, electricity) are billed. Cloud computing provides organizations the ability to unilaterally increase/decrease resources (on-demand self-service), share common resources (resource pooling), change or eliminate resources based on need instantaneously (rapid elasticity), pay only for the services consumed (measured service), and access the wealth of services through an internet connection (broad network access). Hybrid cloud information systems use a blend of on-premises architectures and cloud computing resources. Incident response is the process of detecting, containing, eradicating, and ultimately recovering from cyber events (Joint Chiefs of Staff, 2012). All information systems require incident response capabilities, whether enterprise, enclave, on-premises, hybrid, cloud computing, or other. Digital forensics is the scientific methodology of repeatedly retrieving and reconstructing computing events from digital artifacts (NCC FSWG, 2014).

The goal of this thesis is to identify the detailed deviations between digital forensic incident-response in on-premises and cloud computing architectures and offer risk management recommendations for migration. Methodologically, the thesis uses a systematic comparative approach to assess incident response cybersecurity risks, both inherent and amplified, in the migration from traditional on-premises enterprise information system architectures to cloud computing architectures. The scope for comparative analysis between on-premises and cloud computing is confined to digital forensic incident-response capabilities. The selection of incident response was deliberate; the forensic processes for incident response differ greatly between traditional (non-cloud) architectures and cloud architectures, which creates a high-pressure situation. The incident remediation process deficiencies are magnified by situational urgency during this crucial period.

The target audience for this thesis is DoD decision makers who are planning policy for migration from traditional data centers to cloud architectures. The thesis examines the related policies authorized by the DoD cyber strategy, benefitting the Cloud Executive Steering Group (CESG), Joint Staff, combatant commands, services, defense agencies, DoD field activities, joint and combatant activities, U.S. Cyber Command, U.S. Strategic Command, and other federal agencies.

B. BALANCING ACCELERATED CLOUD ADOPTION WITH CRUCIAL RISK MANAGEMENT DECISIONS

This thesis is primarily motivated by the impact of the DoD's recent accelerated adoption of cloud technologies (Shanahan, 2017). Cloud technologies inherently bear four issues—or conditions—that represent crucial risk management decisions for the DoD, and they each deserve fresh inspection in relation to information security: customer misconfigurations, cloud leaks, complications in the implementation of security controls, and digital forensic incident-response challenges.

The head of global security programs for the largest cloud service provider (CSP), Amazon Web Services (AWS), stated that persistent customer misconfigurations continue to garner sustained attention by senior leadership (Clarke, 2015). Several recent cloud

leaks, generated by DoD vendors' unforced errors, have publicly exposed intelligence community data without the aid of malicious actors (Gregg, 2017). By introducing a cloud service provider relationship into digital forensic incident-response (DFIR), the DoD has dramatically altered the coordination of incident response tactics (Cloud Security Alliance [CSA], 2017). This relationship may also complicate the implementation of some of the DoD's security controls (SANS Institute, 2016). At the the National Institute of Standards Technology (NIST), the NIST Cloud Computing Forensic Science Working Group (NCC FSWG) identified 65 digital forensics challenges that either alter or make more difficult to perform established digital forensic incident response processes.

1. Customer Misconfigurations

Cloud computing is a transformative technology; thus, the DoD has committed to expedite cloud adoption within its existing infrastructure. The Deputy Secretary of Defense conveys the DoD's sense of urgency toward cloud adoption and technological modernization in fulfilling the DoD's mission (Shanahan, 2017). In 2017, the secretary announced the initiation of a Cloud Executive Steering Group (CESG) to research, execute, and implement commercial solutions for the DoD's cloud strategy. He recommended that the aggressive cloud initiative take place over two phases: The first phase will focus on developing a customized acquisition strategy to transform the military's enterprise cloud capabilities so that it can handle unclassified, secret, and top-secret information (Shanahan, 2017). This first phase must also include a technical analysis of the processes needed to support the migration and the subsequent training for successful execution. For the second phase, the Cloud Executive Steering Group will quickly integrate cloud security and machine learning into day-to-day practices of selected DoD components.

The secretary's comments, however, only address the DoD's requirements as a cloud client; from the perspective of a cloud service provider, there must also be an emphasis on the client's role in achieving cloud security. In a 2015 interview, the head of AWS' Global Security Programs articulated the dividing line that separates the security duties of the cloud service provider from the duties of the client. The cloud service provider is the organization or vendor that supplies the cloud computing pool of resources deployed

in software as a service, platform as a service, and infrastructure as a service (these concepts are discussed in more detail in Chapter II; Defense Information Systems Agency [DISA], 2017). From the cloud service provider’s perspective, the client has responsibilities—which span from the running applications down to the guest operating system—and the provider is responsible for the host operating system down to the physical data center (Clarke, 2015). This type of cooperative security is commonly referred to as a shared responsibility model in the cloud. In his interview, the AWS representative amplified the grave consequences of misconstruing the roles of the client and the cloud service provider in the shared responsibility model; he expressed particular concern over the possibility that a client might misconfigure applications, which could make the organization vulnerable to data leaks (Clarke, 2015).

This misconfiguration is the first inherent issue of operating securely in the cloud. Consumers often execute a cloud migration incorrectly, assuming that they can simply port their entire traditional IT architecture to the cloud without any modification (often referred to as lift and shift; CSA, 2017). Cloud Security Alliance (CSA), a leading cloud security industry group, recommends embracing the clean slate approach to cloud migration. Given the ineffectiveness of lift and shift, the CSA recommends that, prior to cloud migration, the client should re-evaluate all of the information potentially destined for the cloud.

2. Cloud Leaks

A data spill is any event involving the unauthorized transfer of confidential data from an accredited information system to one that is not accredited (DISA, 2017). A cloud leak is a type of data spill, specifically originating from a cloud environment. As early as 2013, the government had investigated data spillage specific to the cloud, documented in a Department of Homeland Security (DHS) presentation on February 14, 2013, called “Spillage and Cloud Computing.” Clearly, agencies involved in national security matters must effectively reduce cloud leaks; however, the cloud leaks problem has not been solved.

For example, on June 1, 2017, *The Washington Post* reported that

An unnamed employee of federal contracting giant Booz Allen Hamilton temporarily left sensitive government passwords exposed online last week,

raising questions about the McLean company's cybersecurity practices after drawing scrutiny for the way top secret data was mishandled in two earlier, high-profile cases. (Gregg, 2017, para. 1)

The report later confirmed that the password leak was on an AWS server for contracted work with the National Geospatial-Intelligence Agency. The report described a prevailing shift in government agencies pivoting to the cloud:

That Amazon's cloud server was being used to service a contract with a U.S. intelligence agency is indicative of a broader shift happening across the government, as data and applications move off individual computers and internal networks and into less costly and more adaptable cloud-based systems. (Gregg, 2017, para. 14)

This type of cloud leak is not an isolated occurrence. Less than six months later, CNN reported that Pentagon data collection for U.S. Central Command (CENTCOM) and U.S. Specific Command (PACOM) by Vendor X had been exposed on September 6, 2017. The CNN report cited a response from the DoD confirming the data leak: “‘We determined that the data was accessed via unauthorized means by employing methods to circumvent security protocols,’ said Maj. Josh Jacques, a spokesperson for U.S. Central Command. ‘Once alerted to the unauthorized access, Centcom [sic] implemented additional security measures to prevent unauthorized access’” (Larson, 2017). This demonstrates yet another incident of insecure DoD data leaked from within the cloud. The DoD needs to fully understand the likelihood of similar leaks from commercial cloud infrastructures—and the necessary modifications to security procedures—moving forward to reduce the probability of a leak to an acceptable level of risk.

3. Complications in the Implementations of Security Controls

To date, agencies involved in national security matters have not effectively reduced cloud leaks. Use of the cloud dramatically limits proactive governance and replication of some traditional security controls. In some instances, traditional enterprise information system security practices have not been replicated identically in the cloud. The SANS Institute, in a 2016 white paper titled *Implementing the Critical Security Controls in the Cloud*, identified this phenomenon as “cloud negative controls,”—when implementation

of security is more difficult or cost-prohibitive in the cloud. The difference in or complexity of implementing standard controls can lead to less security in the cloud.

When an organization transitions to the cloud, it loses governance, specifically of forensic incident-response. Cloud computing digital forensics attempts to augment digital-forensic scientific methods to accommodate for cloud computing's unique environment. Incident response teams are responsible for investigating and filtering cyber incidents from commonly occurring events. All cyber incidents are events but not all events qualify as cyber incidents; an event is any observable episode on an information system, while cyber incidents are more threatening and either pose a legitimate or potential adverse consequence to an information system (Joint Chiefs of Staff, 2012). In incident response, digital forensics professionals are only needed after the organization has determined that the threshold for an escalated cyber incident has been met, at which point digital forensics professionals initiate a process for identifying, collecting, preserving, examining, and reporting on the digital evidence (NCC FSWG, 2014). But the organization has lost governance because, in cloud computing, the organization no longer owns those resources; it rents them. The computing resources are also remote to the organization, and under the control of the cloud service provider.

The 2016 SANS paper indicates that it is imperative for any organization's security architect to have the ability to discern how on-premises networks differ from virtualized architecture. The SANS paper categorizes security controls into cloud-positive, cloud-negative, and cloud-neutral controls. The three tiers correspond to the ease of application within the cloud. The SANS recommendation, based upon this awareness, allows the security architect to direct greater attention to the cloud-negative controls. The paper specifically identifies logging, boundary defense, and incident response management as cloud-negative controls.

4. Digital Forensic Incident-Response

The European Network and Information Security Agency (ENISA; 2009) advises that the cloud consumer must learn the transparency of incident reporting, what constitutes an incident, the time frame between incident detection and client notification, and the level

of involvement of the consumer. Because cloud service providers have additional restrictions on vulnerability assessment and penetration testing, the consumer must understand the processes performed by the cloud service provider to fulfill those duties. The consumer also needs to determine his or her satisfaction with the quality and quantity of the provider's incident mitigation practices. ENISA also suggests that the consumer should learn if a forensic image of the virtual asset will be made available after an incident. In addition to technical safeguards, ENISA indicates that the consumer should understand the cloud service provider's comprehensive defense against social engineering in causing breaches (ENISA, 2009).

The SANS Institute (2016) comments that fewer options are available in the cloud for intrusion detection systems or intrusion prevention systems than for traditional enterprise information system infrastructure security. If the available options are not up to the necessary standards before cloud migration, the cloud will require alternative solutions. The cloud also interjects the cloud service provider relationship into the customer's workflow. The SANS paper underscores that roles have to be clearly defined to accommodate the introduction of the cloud service provider. Further exacerbating the situation is that the cloud contract is the sole vehicle to arbitrate and memorialize cloud governance (CSA, 2017).

In 2014, the NIST Cloud Computing Forensic Science Working Group (NCC FSWG) published research focused specifically on the challenges of retrieving artifacts from the cloud. The working group observed a bifurcation in the challenges: either they were substantively altering or made it more difficult to perform established digital forensic incident-response processes. Both are direct consequences of the impact of cloud-specific technologies. A universal condition consistent across every permutation of service and deployment model is diminished access to forensic data (NCC FSWG, 2014). Access to data is essential for incident response; restricted access is at the root of most of the working group's itemized challenges. Digital forensic scientists have only a rudimentary understanding of the cloud service provider operating environment. The consumer, in turn, is faced with limited transparency because the cloud remotely delivers abstracted resources dispensed by the cloud service provider. Therefore, the consumer has inadequate

environmental details regarding such important cloud service provider resources as architecture, hardware, software, file system type, and how that particular resource works (NCC FSWG, 2014). Therefore, in addition to consumers being unaware of the interworkings of the cloud service provider environment, they are also uncertain about the indeterminate boundaries between themselves and other clients (NCC FSWG, 2014). These indeterminate boundaries, especially for cloud service providers, which do not provide vertical isolation for the consumer's data, warrant consumers' skepticism; as NIST explains, can the cloud service provider deliver high-assurance data integrity or preservation extended over multiple parties, computers, and locations? Complicating matters further, the consumer has accepted the risks of operating in a multi-tenant environment. The lack of transparency also begs the question: Is the consumer genuinely capable of "collecting, accurate, complete, traceable, audible [sic] and forensically sound evidence" (NCC FSWG, 2014, p. 25)?

Ultimately, there is a viable solution for the challenges facing the DoD in transitioning to a robust, secure cloud environment. However, that solution will require the DoD to reorganize people and processes to minimize the existing gaps between how traditional applications operate and how cloud computing applications are configured. The solution will also require the DoD to incorporate broad uses of encryption, two-factor authentication, digital forensic incidence-response processes tailored to cloud architectures, practicable workarounds that address cloud-negative security controls, and substantially more mandatory cloud training. If any of these are absent, the DoD will find itself in unfamiliar territory and will face limitations for proactively addressing cyber incidents.

C. CHAPTER SUMMARY AND THESIS ORGANIZATION

Chapter I has detailed the DoD's imminent use of cloud capabilities, and has reviewed considerations for managing the associated risk. The DoD actively seeks greater cloud adoption, cloud service providers are concerned about greater client cloud misconfigurations, DoD vendors are publicly leaking confidential data, and incident response manifests itself in more complicated machinations in the cloud. This thesis posits

that the velocity of cloud adoption—multiplied by the immaturity of cloud incident response protocols—warrants a rigorous investigation of the net balance of digital-forensic incident-response risk management capabilities.

Chapter II provides an informal, baseline introduction to cloud computing principles, and defines technical cloud computing terms. When weighing the relative information assurance issues, definitions are fundamental to informed analysis. Therefore, judiciously, this thesis prioritized the technological leadership of NIST as an authoritative source for taxonomy because of NIST’s foundational, vendor-neutral, widely circulated special publications. Organizationally, crucial terms are defined early in Chapter II.

The analysis in Chapter III further refines the deductions established in Chapter II by ordering the thesis’ findings along three themes: technical, legal, and boundary digital forensic incident-response challenges. The related processes of on-premises systems are juxtaposed against the cloud for significant security implications. Chapter IV introduces the federal cloud computing policy, FedRAMP, and contrasts it with DoD FedRAMP+ with respect to cloud-based security controls. Chapter V introduces the concept of transformational migration and offers specific recommendations to the four pressing cloud risk management issues facing the DoD. Chapter VI concludes with a summary and suggestions for tangential future research.

II. CLOUD COMPUTING PRINCIPLES

A. PRELIMINARY ORIENTATION

This chapter introduces crucial cloud computing and cybersecurity definitions used throughout the thesis in evaluating the critical gaps between on-premises and cloud computing capabilities. Before introducing formal cloud definitions, we must first demystify the term “cloud”; this chapter provides an informal heuristic to introduce a more accessible definition of the cloud, which will enable the deconstruction of common misunderstandings about its structures and responsibilities. Substitute the term cloud with a concept instead: that locally operated, on-premises resources will access a larger pool of remotely shared resources over the internet. Cloud service providers own, operate, and make available distributed data centers with massive quantities of fundamental computing power (e.g., random access memory [RAM], central processing units [CPUs], storage, and network connectivity).

NIST’s formal definition for cloud computing is, “a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or cloud provider interaction” (Mell & Grance, 2011, p. 2). The cloud service provider’s product offerings include on-demand operating systems, servers, development tools, hard drives, applications, and other services. Additionally, cloud service providers constantly innovate and deliver new capabilities that are commercially supportable, as illustrated in Figure 1.

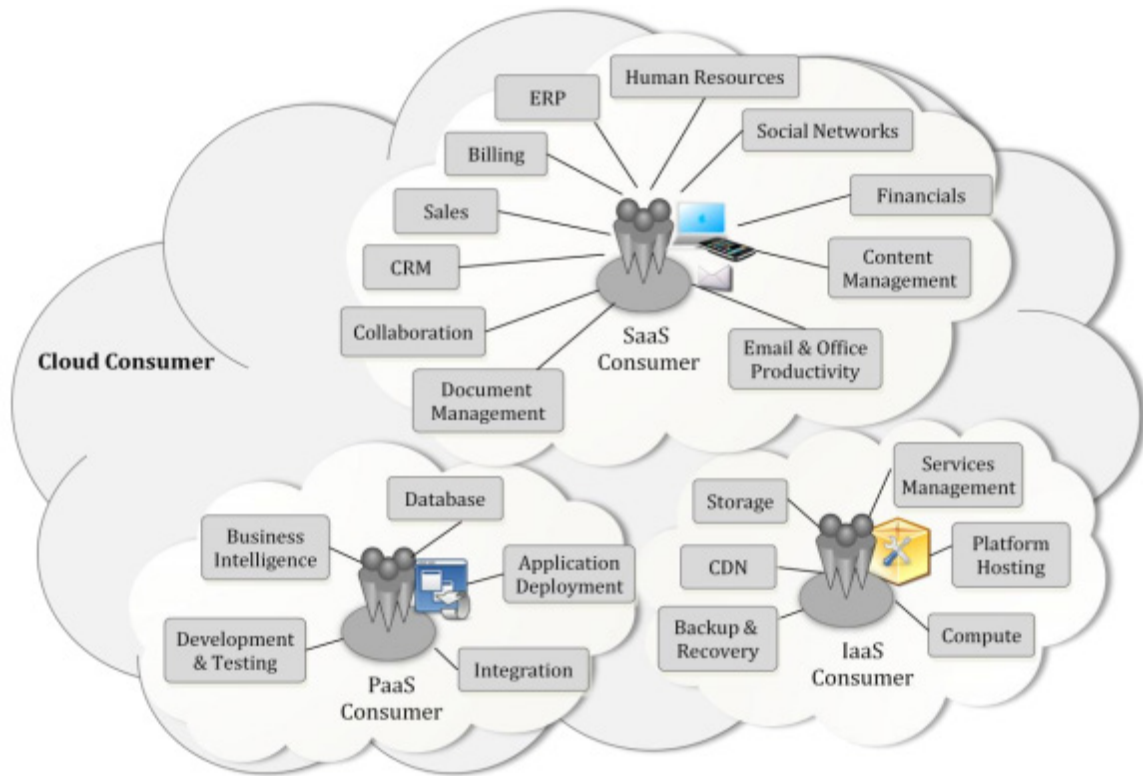


Figure 1. Examples of Cloud Consumer Services. Source: Liu et al. (2011).

The cloud service provider can be seen as a big box retailer that charges a unitized cost for the work/cycle performed by each instance of utilization. This arrangement of unitization enables the cloud client to increase and decrease use, as well as “swap and drop” products with minimal costs or risks. The cloud is predicated on a self-checkout service model with itemized/metered billing in which clients can consume as much as they wish with little involvement from the cloud service provider. While the near-inexhaustible pooled computing resources are readily available to the client, it is the chief responsibility of the cloud service provider to ensure the secure provisioning of currently unavailable resources. The cloud client can gain significant and always compelling additional savings in a reduced IT headcount with fewer on-premises resources to manage.

B. DEFINITIONS

This section divides definitions into two main categories, cloud computing (provided primarily by NIST) and cloud security (provided by the Committee on National Security Systems [CNSS]).

1. Cloud Computing

This section defines deployment and service models, along with other pertinent terms. The paramount takeaway is that any meaningful discussion about cloud security will not refer to the ubiquitous cloud but will instead reference a specific selected architecture instantiation, reflecting committed organizational choices.

Directly, NIST Special Publication (SP) 800-145 has set forth the following definitions as inherent qualities of cloud computing, which fit within the cloud computing model shown in Figure 2.

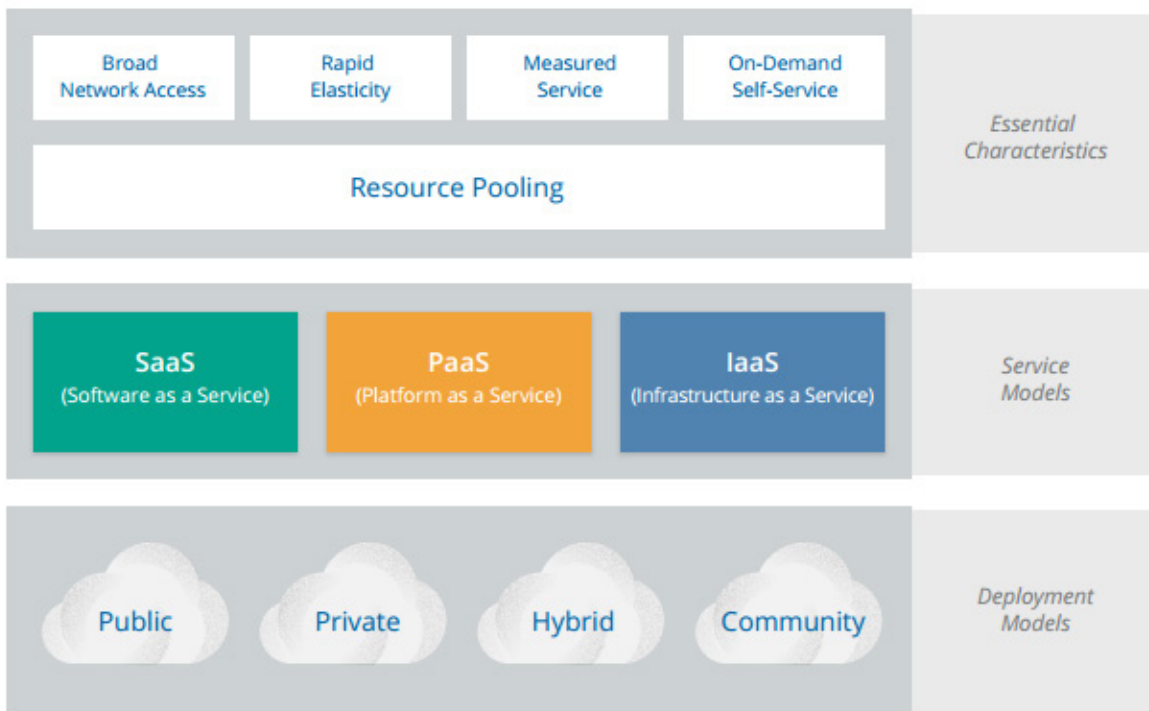


Figure 2. NIST Cloud Computing Model. Source: CSA (2017).

a. Essential Characteristics

In SP 800-145, essential characteristics are described as follows:

On-demand self-service. A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.

Broad network access. Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations).

Resource pooling. The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, and network bandwidth.

Rapid elasticity. Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.

Measured service. Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service. (Mell & Grance, 2011, p. 2)

b. Service Models

SP 800-145 also discusses cloud service models:

Software as a Service (SaaS). The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or

even individual application capabilities, with the possible exception of limited user specific application configuration settings.

Platform as a Service (PaaS). The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.

Infrastructure as a Service (IaaS). The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls). (Mell & Grance, 2011, pp. 2–3)

c. *Deployment Models*

Figure 3 shows an overview of public and private clouds, and the ways in which they differ.

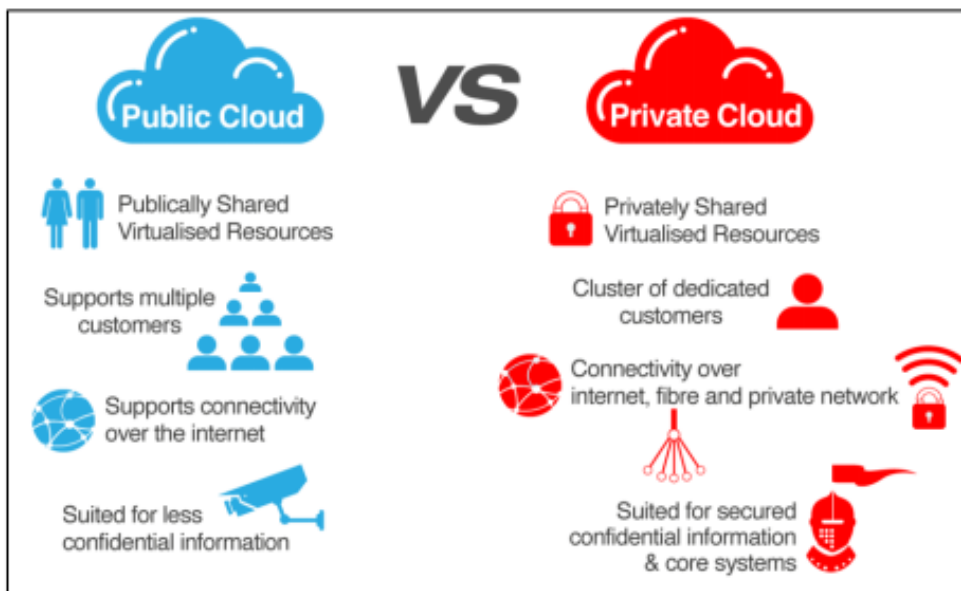


Figure 3. Public versus Private Cloud. Source: Odell et al. (2015).

SP 800-145 defines key terms in this domain as follows.

Private cloud. The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.

Community cloud. The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.

Public cloud. The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.

Hybrid cloud. The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds). (Mell & Grance, 2011, p. 3)

The Official (ISC)² Guide provides another key definition:

Cloud Provisioning. The deployment of a company's cloud computing strategy, which typically first involves selecting which applications and services will reside in the public cloud and which will remain on-site behind the firewall or in the private cloud. Cloud provisioning also entails developing the process for interfacing with the cloud applications and services as well as auditing and monitoring who accesses and utilizes the resources. (Gordon, 2016, p. 9)

And NIST SP 800-125 rounds out the terminology for deployment models with:

Virtualization. The simulation of the software and/or hardware upon which other software runs. This simulated environment is called a virtual machine (VM). (Scarfone, Souppaya, & Hoffman, 2011, p. 2-1)

d. A One-Size-Fits-All Cloud Does Not Exist

Two U.S. intelligence agencies—the Central Intelligence Agency (CIA) and the National Security Agency (NSA)—are relatively recent adopters of cloud computing and, unsurprisingly, demand stringent information security requirements. In a 2014 interview, the CIA’s chief information officer (CIO), Doug Wolfe, confirmed that the two clandestine agencies chose to build out their respective cloud architectures differently (CIA, 2014). Wolfe explained that the CIA cloud was built using commercial cloud products with participation from a commercial cloud service provider, while the NSA cloud was designed in-house, also using commercially available products but without participation from a commercial cloud service provider. The CIA and NSA pursued two different paths in achieving similar cloud computing capabilities. Also in 2014, the NSA CIO, in a separate interview with *Network World*, detailed the cloud decisions the NSA made to integrate a community cloud that could service other intelligence mission partners and meet the security requirements demanded by the NSA for classified data (Smith, 2014).

The NSA CIO interview also revealed that the community cloud was able to enhance performance because it collocates all community data repositories within common data centers (Smith, 2014). Comments from the two CIOs dispel the notion of a one-size-fits-all cloud, even for two seemingly similar technologically sophisticated U.S. intelligence agencies, by contextualizing that each organization has to make choices with respect to cloud deployment and service models.

2. Cloud Security

The Committee on National Security Systems (CNSS) is the U.S. intergovernmental body that develops cybersecurity guidance at the national-level. It is therefore an appropriate forum to consult for information assurance definitions. CNSSI 4009 defines key terms in this domain as follows (CNSS, 2015):

Confidentiality. “Preserving authorized restrictions on information access and disclosure, including means for protecting personal privacy and proprietary information” (p. 30).

Integrity. “Guarding against improper information modification or destruction, and includes ensuring information non-repudiation and authenticity” (p. 68).

Availability. “Timely, reliable access to data and information services for authorized users” (p. 11).

Information System-Related Security Risks. “Risk that arises through the loss of confidentiality, integrity, or availability of information or information systems considering impacts to organizational operations and assets, individuals, other organizations, and the Nation. A subset of information security risk” (p. 66).

Vulnerability. “Weakness in an information system, system security procedures, internal controls, or implementation that could be exploited by a threat source” (p. 131).

Threat. “Any circumstance or event with the potential to adversely impact organizational operations (including mission, functions, image, or reputation), organizational assets, individuals, other organizations, or the Nation through an information system via unauthorized access, destruction, disclosure, modification of information, and/or denial of service” (p. 122).

Impact. The effect on organizational operations, organizational assets, individuals, other organizations, or the nation (including the national security interests of the United States) of a loss of confidentiality, integrity, or availability of information or an information system (p. 60).

Countermeasures. “Actions, devices, procedures, techniques, or other measures that reduce the vulnerability of an information system. Synonymous with security controls and safeguards” (p. 33).

C. THE CIA TRIAD AND THE INFORMATION SECURITY EQUATION

Information security at an elementary level consists of confidentiality, integrity, and availability. Collectively, the three are frequently referred to as the CIA triad. As previously stated, the primary objective of all high-functioning information security programs is preserving the CIA triad. Therefore, a worthy exercise is to establish the attributes that both achieve and violate the CIA triad, in adherence with the definitions published by the CNSS glossary (2015).

The security risk equation in Equation 1, adapted from NIST (n.d.), is not referenced as prevalently as the CIA triad, but it arguably offers comparable utility in concisely and accurately mapping dynamic information assurance interrelationships. The value in the security risk equation is that it encapsulates the entire problem domain for the security practitioner.

$$\text{Security Risk} = \frac{\text{Vulnerability} * \text{Threat} * \text{Impact}}{\text{Countermeasures}} \quad (1)$$

D. DISTINGUISHING CLOUD-SPECIFIC VULNERABILITIES FROM ON-PREMISES SYSTEMS

Reviewing multiple cloud publications—from popular magazines to academic papers—revealed a bifurcated view of extremes regarding cloud security. Groups either claimed that the cloud would solve all traditional on-premises system security problems or could not be trusted at all. This thesis’ research provides greater confidence that the cloud can engender trust and can address traditional on-premises system problems to scale, but with a major caveat—the operating concepts for data, defense perimeter, skillsets, and processes require fundamental restructuring or migration. This thesis implements a more nuanced approach. This section first isolates for general on-premises system vulnerabilities, threats, and countermeasures that need to be considered before determining whether the cloud is more or less secure. Before designating a vulnerability as cloud-borne, it needs to meet a set of criteria—a litmus test to decide if a vulnerability should be assigned as cloud-specific. Determining whether a vulnerability is cloud-borne is helpful in discussions with reluctant managers about the relative risk of the cloud. Published by the Institute of Electrical and Electronics Engineers (IEEE), “Understanding Cloud Computing Vulnerabilities” provides a rubric that helps determine if vulnerabilities are cloud-borne (Grobauer, Walloschek, & Stöcker, 2011). According to the rubric, a vulnerability is cloud-specific if it:

- is intrinsic to or prevalent in a core cloud computing technology

- has its root cause in one of NIST’s essential cloud characteristics
- is caused when cloud innovations make tried-and tested security controls difficult or impossible to implement, or
- is prevalent in established state-of-the-art cloud offerings. (Grobauer et al., 2011, p. 52)

The first bullet refers to web applications, virtualization, and cryptography as the core cloud technologies (Grobauer et al., 2011). Web applications and services are considered central to the cloud because the cloud is transported via internet HTTP protocols—often to a Web browser. Since the internet delivers the cloud, it uses software that emulates hardware, which drives down costs. The delivery mechanism unleashes tremendous economies of scale. Cryptography is intrinsic to core cloud technology; consumers would balk at paying for services lacking in confidentiality, privacy, and integrity. The second bullet alludes to the five essential characteristics attributed to NIST (described in the previous chapter)—on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service. These characteristics are core to the cloud (Mell & Grance, 2011), and are precisely the traits that differentiate the cloud from traditional on-premises systems.

The third bullet identifies instances when on-premises system security practices do not transfer to the cloud—for example, the “cloud-negative controls” identified by SANS (2016). The fourth bullet describes the cloud as pushing present technological boundaries. If a vulnerability is identified in an advanced cloud offering—one that has not been previously identified—then it must be a cloud-specific vulnerability. This thesis agrees that there is some merit to the argument, but ultimately finds it conflated; the IEEE paper includes weak authentication implementations (not exclusive to the cloud) in cloud offerings (Grobauer et al., 2011). Due to the flaw in this interpretation, this fourth indicator can only be seen as partially attributed, or a hybrid cloud-specific vulnerability. This partial attribution is addressed later in this thesis.

E. COMPARING CLOUD-SPECIFIC THREATS FROM ON-PREMISES SYSTEM THREATS

Table 1 shows threats that are attributable to conventional on-premises architectures, and those that can be ascribed to cloud computing architectures. The distribution shows that cloud computing technologies need to account for several additional threats.

Table 1. Treacherous 12 Threats Summary. Adapted from CSA (2016).

Threat in Conventional Architectures	Threat in the Cloud
Data breaches	Data breaches
System and application vulnerabilities	System and application vulnerabilities
Advance Persistent Threats	Advance Persistent Threats
Data loss	Data loss
Denial of service	Denial of service
Weak access management	Weak access management
Account hijacking	Account hijacking
Malicious insiders	Malicious insiders
Insufficient due diligence	Insufficient due diligence
	Insecure APIs
	Nefarious use of cloud services
	Shared technology vulnerabilities

Considering the vulnerabilities borne of cloud architectures, it is important to determine which vulnerabilities could be exploited by cloud-specific threats. In 2016, the CSA released *The Treacherous 12: Cloud Computing Top Threats in 2016*, which it compiled by surveying cloud industry experts. *The Treacherous 12* ranks a dozen security concerns in order of severity:

1. Data Breaches
2. Weak Access Management
3. Insecure APIs
4. System and Application Vulnerabilities

5. Account Hijacking
6. Malicious Insiders
7. Advanced Persistent Threats (APTs)
8. Data Loss
9. Insufficient Due Diligence
10. Abuse and Nefarious Use of Cloud Services
11. Denial of Service
12. Shared Technology Issues
(CSA, 2016, p. 5)

Data breaches are violations of confidentiality that result in unauthorized disclosure. Applying the four conditions from the cloud vulnerability rubric (Grobauer et al., 2011), data breaches are not a cloud-specific threat. They existed prior to widespread cloud deployment, and are a problem for on-premises systems as well. The cloud, however, may present a new threat vector because of multi-tenancy.

Weak access management is a violation of integrity that can lead to destruction, disclosure, and distributed denial-of-service (DDos) consequences (denial of service is discussed later in this section). A poor credentialing scheme can originate from a variety of factors, but the scalability and reuse of resources exacerbates the weakness. Weak access management fits the second condition in the cloud vulnerability rubric—“has its root cause in one of [the] essential cloud characteristics” (Grobauer et al., 2011, p. 52)—because the vulnerability is rooted in the resource pooling essential characteristic. While weak access management is not a cloud-specific threat, the cloud increases the attack surface area based upon its rapid provisioning and de-provisioning capabilities

An insecure application programming interface (API) can lead to violations of confidentiality and integrity, resulting in unattributed unauthorized disclosures. In the cloud vulnerability rubric, insecure APIs exploit the first two conditions: they are specific to the cloud and rooted in essential cloud characteristics. The cloud’s reliance on internet connectivity/Web services clearly exploits the first condition, while the second is exploited because elasticity in the frequent provisioning/de-provisioning of cloud access constitutes

an essential characteristic of the cloud. Additionally, *The Treacherous 12* proposes that risks may increase proportionally with interoperability when third-party interfaces are accepted as trusted relationships even when they should remain untrustworthy (CSA, 2016).

System and application vulnerabilities are simply synonymous with buggy code. Poorly written code can lead to a plethora of undesirable outcomes, including unauthorized disclosure, destruction, and system unavailability. Those with computer experience realize that problematic code written by humans will likely carry inherent flaws. Thus, system and application vulnerabilities are not a cloud-specific threat; however, it is reasonable to argue that the cloud increases the attack surface area because of its growing ubiquity (CSA, 2016).

Account hijacking occurs when a user account is compromised—for example, by way of a phishing scam—and is another threat that can lead to innumerable bad outcomes. This threat does not exploit any of the cloud vulnerability rubric's four criteria specific to the cloud. However, adversaries could use the cloud account to perform more insidious attacks than they can for an on-premises system distributed denial of service attack, leveraging only a breached IP address. Again, while this threat is not cloud-specific, the cloud increases the attack surface area based upon the potential access to the management plane a breach account can gain.

Malicious insiders are employees who do not work in accordance with an organization's goals; they present a more urgent concern than external adversaries because they already have access to the information system. While malicious insiders are not a cloud-exclusive threat, by engaging in a necessary relationship with a cloud service provider, the organization potentially takes on additional insiders, which increases the odds of encountering bad actors. This thesis expands the definition of insiders to include cloud service provider employees who will have account access privileges equivalent to a system administrator's access. Therefore, cloud service provider personnel should be vetted according to the standards of the organization's hiring practices. Like we have seen with other threats, malicious insiders are also not a cloud-specific threat; however, the cloud again increases the attack surface area.

Advanced persistent threats (APTs) are adversaries who have the motivation and resources to perpetually attempt to breach an information system and, upon a breach, laterally maneuver to exploit to their own design. APTs do not care whether their target is on-premises or cloud-based, and therefore this thesis does not consider APTs to be a cloud-specific threat. Actually, this thesis contends that the additional resources and vigilance of the cloud service provider increase the available resources dedicated to counter APTs, and thus increase the net countermeasures to combat them.

Data loss is exactly what it sounds like: when data is permanently, unintentionally unavailable. It should be clear that this threat is also agnostic to the type of information system, whether cloud or on-premises. Applying Grobauer et al.'s four conditions, the thesis determines with little controversy that data loss is not a cloud-specific. Furthermore, this thesis proposes that the cloud provides for inexpensive redundancy at multiple locations and can therefore, to an extent, reduce certain variants of data loss—for example, via data provenance technologies.

Insufficient due diligence is when an organization has a low threshold of due care and verification for vetting a contracted technological service. Applying the cloud vulnerability rubric, it is not a singularly a cloud-specific threat; however, this thesis holds that insufficient due diligence in the cloud can both amplify and accelerate calamity because the consumer's architecture is no longer remote. As a counterbalance, though, the aforementioned misconfiguration concerns mean that a cloud service provider that passes a rigorous due diligence process can at times offer additional tools to help protect the consumer. Ultimately, insufficient due diligence is a partial cloud-specific threat.

Nefarious use of cloud services is the unauthorized leveraging of cloud resources for unintended purposes. This threat, per the rubric, is necessarily cloud-specific. It exploits all four conditions because the threat presupposes a cloud infrastructure.

Denial of service occurs when an actor purposefully inhibits the availability of data, applications, or both. Denial of service is not a cloud-specific threat; its existence preceded cloud services.

Shared technology vulnerabilities are a cloud-specific threat, as they are a byproduct of multi-tenancy. Applying the rubric conditions, multi-tenancy is both core to cloud computing technology and has root causes from its essential characteristics of resource pooling and rapid elasticity. The prior operational questions address inquiring about the policies of isolating machine images. The extent of the shared technology vulnerability will depend on how the cloud service provider implements its machine isolation and provisioning/de-provisioning.

In summary, when validated against Grobauer et al.'s rubric for cloud-borne vulnerabilities, of the *Treacherous 12* threats, three are fully cloud-specific threats (APIs, nefarious use of cloud services, and shared information technology vulnerabilities). Four are partially cloud-specific, meaning they do not specifically exploit a cloud-borne vulnerability but they either increase the attack surface area or potentially could inflict greater harm (weak access management, account hijacking, malicious insiders, and insufficient due diligence). And the remaining five are not cloud-specific threats. Additionally, two of the threats actually have potentially more effective mitigations within a cloud environment: APIs and data loss.

The takeaway from this analysis is that, of the 12 greatest estimated threats that experts say emanate from the cloud, only three point to truly cloud-specific vulnerabilities. Insecure APIs, nefarious use of cloud services, and shared technology vulnerabilities are the cloud-specific threats that merit additional in-depth defense security measures. Additionally, weak access management, account hijacking, malicious insiders, and insufficient due diligence are the next tier of cloud-specific threats to address.

F. DISTINGUISHING CLOUD-SPECIFIC COUNTERMEASURES

Countermeasures and security controls have the potential to detect, reduce, or even eliminate threats that exploit vulnerabilities. The goals of each specific cloud project, service model, and cloud service provider platform are the critical inputs in determining the additional countermeasures the project should integrate. Countermeasures and security controls are risk management tools. This section on countermeasures is brief, as Chapter V is dedicated to recommending specific cloud computing risk management options.

The first step in arriving at countermeasures is for the organization to conduct a risk assessment based on its use case(s). In further detail, the logical steps are an organization should generate its requirements, map its architecture, and conclude by diagnosing and then prioritizing the remaining security gaps of the cloud service provider (CSA, 2017). There are many sources that offer large catalogs of security controls categorized by needs, such as access control or incident response. The cataloged controls are further refined by impact level for more cost-effective pairings between data impact level and requisite controls. NIST's SP-800-53, the International Organization for Standardization International's and Electrotechnical Commission's ISO/IEC 27001:2013, and the framework known as Control Objectives for Information and Related Technologies (COBIT) 5.0 are three widely adopted security control catalogs. The Cloud Controls Matrix, published by the CSA, is a rational catalog to begin with because it maps its controls side-by-side with many other control catalogs for easy comparison.

Chapter III inspects the litany of technical challenges faced by digital forensic incident-response professionals.

III. TECHNICAL ISSUES IN CLOUD COMPUTING

This chapter calls attention to the specific security implications that come with the different service models an organization can choose to adopt: infrastructure as a service, platform as a service, or software as a service. The service model an organization selects determines the level of involvement it must have in application development within the cloud service provider’s environment. The service model selection in itself also determines a particular set of consumer challenges balanced against greater autonomy in managing cloud-specific settings, configurations, and controls. The majority of the digital forensic challenges this chapter identifies are cloud-intrinsic and, consequently, predominately service model–neutral. So why does distinguishing between service models even matter? The initial benefit is that doing so raises the awareness of additional cloud security pain points, enabling the consumer to abate these issues through a combination of policy changes or contracts with additional security services. The sustaining, long-term benefit is that doing so provides an organization with a foundation for “transformational migration.” Transformational migration, a term coined in this thesis and discussed with the recommendations in Chapter V, is an operational strategy—essentially a blueprint—to help organizations interface with cloud service providers to best optimize cloud capabilities and minimize security risks.

A. SERVICE MODEL SELECTION: COMPARATIVE RISK FACTORS

The Cloud Security Alliance’s 2017 white paper provides a high-level overview of security responsibilities based on service models. Compared in Figure 4, infrastructure as a service and software as a service reflect the greatest and least consumer security responsibilities, respectively (CSA, 2017). The first step in assessing the significance of the discrete technological, legal, and boundary challenges the cloud presents is to use a generic-abstracted model to trace how the fundamental cloud building blocks interact. As depicted in Figure 5, an intuitive perspective views each cloud service as a series of discrete service functionalities stacked on top of the preceding service model.



Figure 4. Consumer Service Model Security Burden Continuum. Source: Cloud Security Alliance (CSA) (2017).

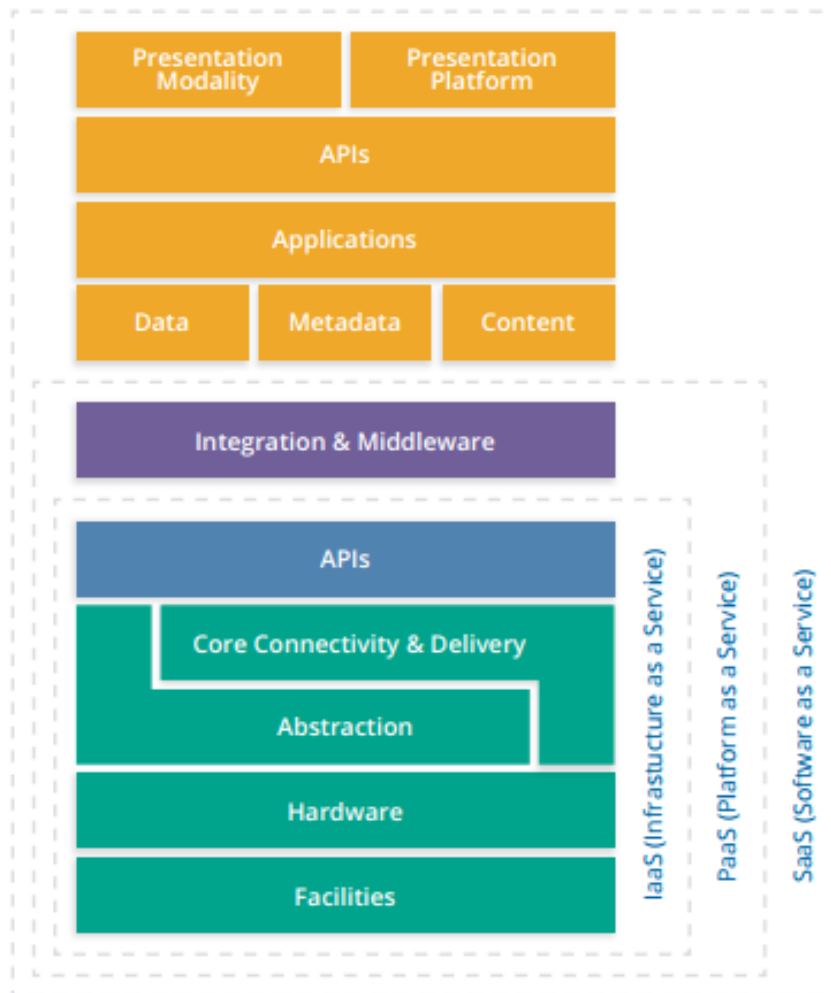


Figure 5. Cloud Reference Service Model Stack. Source: CSA (2017).

(1) Infrastructure as a Service

- Infrastructure as a service allows the consumer discretion over raw computing resources for the purposes of creating any application of its choosing.
- In an infrastructure as a service deployment, the consumer only uses the infrastructure, not the platform or the applications.
- This structure allows the most freedom in tailoring a cloud environment, but creates the greatest consumer security burden in exchange.
- From a consumer perspective, infrastructure as a service is the base structure; it is more similar than the others to an on-premises information system (Mell & Grance, 2011).

(2) Platform as a Service

- Platform as a service builds on top of the infrastructure as a service abstraction; it contributes middleware to the cloud stack.
- Consumers can build their own applications using the cloud service provider's tools (Mell & Grance, 2011).
- Platform as a service comes with a suite of tools, including—almost exhaustively—any in-demand programming language.
- Middleware provides a state-of-the-art development environment platform to develop applications, including software as a service for external clients.
- From a consumer perspective, platform as a service enables the consumer to develop without concern for acquiring and maintaining many of the traditionally associated high-budget hardware costs (e.g., servers).

- According to the Cloud Security Alliance (2017), because platform as a service deployment is accessed through the platform provided by the cloud service provider, the cloud service provider is responsible for securing the platform, with all other security responsibilities assigned to the consumer (e.g. built applications and data).

(3) Software as a Service

- At the top of the service model stack in Figure 5, software as a service represents a full-scale cloud application.
- Consumers can only run applications pre-built by the cloud service provider (Mell & Grance, 2011).
- From a consumer perspective, familiar software as a service applications include Microsoft Office 365 and Dropbox; users access the application based on a pay-as-you-go subscription agreement.
- While software as a service can be developed with a combination of infrastructure as a service or platform as a service—even from different cloud service providers—the user will be one of many clients; this gives rise to multi-tenancy concerns.
- As noted in Figure 4, the consumer has the least onerous direct security responsibilities with software as a service.
- However, indirect security concerns remain. Because the delivery may entail a potential combination of infrastructure as a service and platform as a service development, the consumer must rely upon cumulative trust of each and every service model present in development of the software as a service product.
- The consumer’s administrative authority is limited to users who can access objects and associated privilege levels (CSA, 2017).

(4) Service Model Stack Components

Infrastructure as a Service:

- Facilities are the physical data center.
- Hardware consists of the inventory of computing physical assets (e.g., servers), networking, and storage.
- Abstraction decouples the physical resources through a process called virtualization (defined in Chapter II.C.1), which allows resource pooling.
- Core connectivity and delivery is the process of integrating and automating the abstracted resources into pools that can deliver metered services to multiple customers.
- APIs are protocols and tools the cloud service provider makes available for the customer to communicate with cloud services, such as configuration settings (CSA, 2017).

Platform as a Service:

- Integration and middleware is an additional layer built on infrastructure as a service that provides the customer with access to programming tools to build applications accessed through platform as a service APIs. Middleware allows the customer to leverage the infrastructure as a service abstraction without managing the hidden resources (e.g., patching networks or load balancing servers; CSA, 2017).

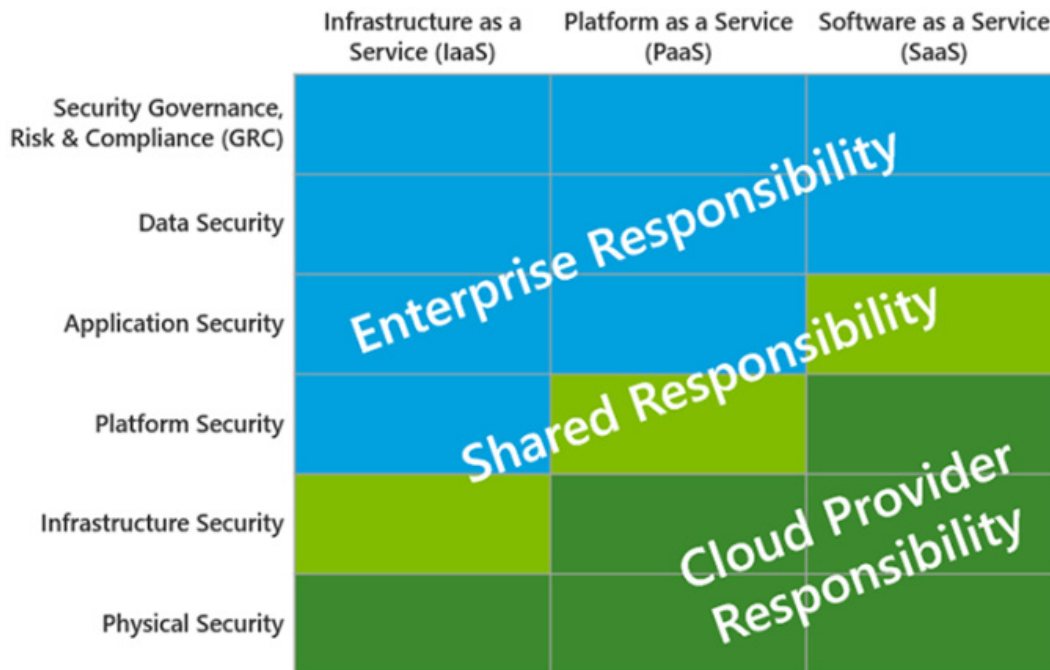
Software as a Service:

- Data is an information source that is either stored or processed.
- Metadata is data about the data, such as when a data file was created or which user created the data file.

- Content is data that has been processed for end users in a form different from the original unprocessed data.
- The majority of cloud software as a service applications are not created in isolation but in combination with contributions from other infrastructure as a service and/or platform as a service.
- APIs are protocols and tools the cloud service provider makes available for the customer to communicate with cloud services, such as configuration settings (CSA, 2017).
- Presentation modality refers to the security features for an application, often determined by whether the end user is on a consumer platform (e.g., social media) or an enterprise platform.
- Presentation platform differs based on the end-point device (e.g., tablet, mobile phone, specialty medical device) to deliver an optimal user experience based on end-point device constraints (CSA, 2011).

B. SHARED RESPONSIBILITY MODEL

The consumer and the cloud service provider both share responsibilities in the cloud relationship. As previously mentioned, confusion or uncertainty about the division of responsibilities often leads to customer misconfiguration. Figure 6 shows the general, if inexact, division of responsibility between the consumer and the cloud service provider, based solely on the service model selection.



As noted in the text, this matrix lacks essential granularity; the matrix provided in Figure 7, in the next subsection, is more complete.

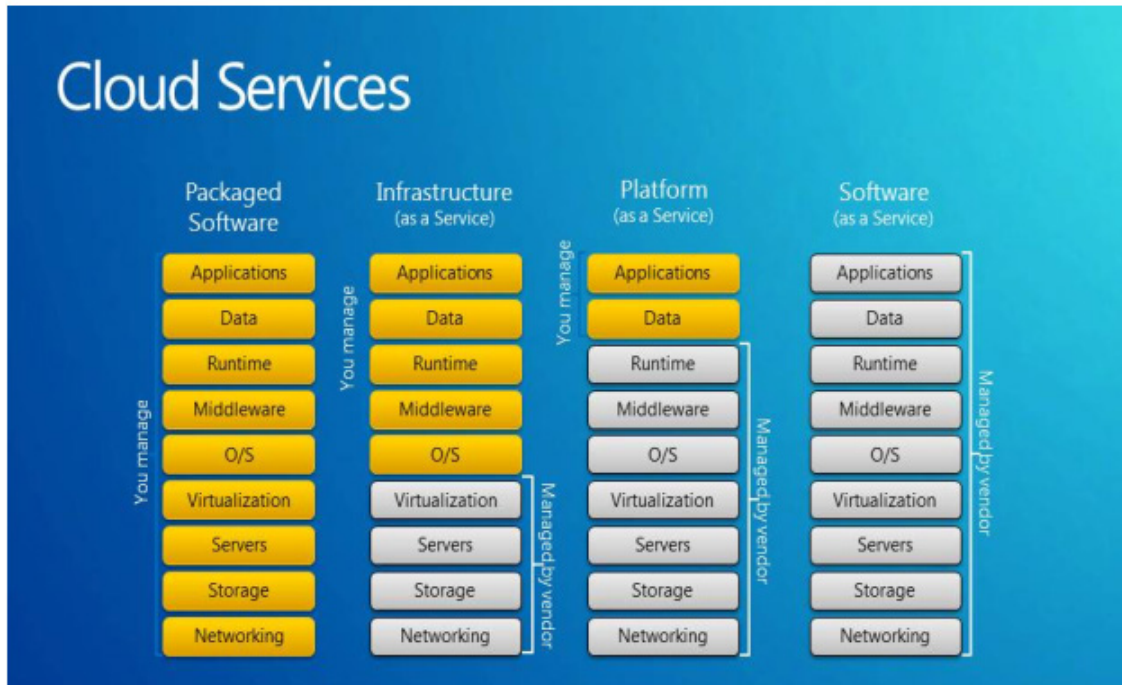
Figure 6. Cloud Security Responsibility Matrix (Low Differentiation).
Source: Gordon (2016).

However, the breadth of Figure 6 is inoperable; when compared to the granularity of the matrix in Figure 7 (presented in the next subsection), relying on the information in Figure 6 alone could lead to misconfigurations. As we begin to understand how the cloud works, Figure 6 is an adequate model because it effectively captures that the three service models have varying levels of consumer participation and joint responsibility. But its effectiveness ends there. The matrix is imprecise about the division of labor and does not usefully compare the three service models to an on-premises application, as the matrix in Figure 7 does. This matrix also inaccurately ascribes the responsibility of data and application security to the consumer in a software as a service model, which shows a misunderstanding about cloud operations. The cloud service provider is managing the entire service model stack in software as a service. This illustrates that customer misconfigurations occur when there is confusion about which entity is responsible for which logical layer of the cloud stack, depending on the service model.

1. Why Lift and Shift Does Not Work

Mistakenly, when some consumers contemplate cloud migration, they assume that they can simply port their entire traditional IT architecture to the cloud without any modification—as previously mentioned, this is known as lift and shift (CSA, 2017). Fallaciously, it is not the responsibility of the cloud service provider to make a lift and shift migration work, because this cloud transition “strategy” is orthogonal to cloud architectures. The lift and shift assumption is largely responsible for the prevalence of consumer misconfigurations. Chapter V.B.1 addresses recommendations for reducing customer misconfigurations.

For a more factual representation of the shared responsibility model, see the leftmost column in Figure 7, which represents an on-premises application. Prior to the cloud, on-premises system applications were highly customized and expected to operate within a standalone data center, and application data was structured for minimal to no interaction with other applications (Bommadevara, Del Miglio, & Jansen, 2018). The remaining three columns in Figure 7 represent how applications interact within each of the three cloud service models. In Figure 7, the yellow highlighting indicates that security for the corresponding layer is a consumer responsibility and the gray highlighting indicates it is the responsibility of the cloud service provider. Cloud applications, juxtaposed against on-premises applications, are highly agile and are expected to operate in multiple data centers; permissioned cloud data is available on-demand for maximal interaction with other applications. This underscores why on-premises applications require changes at multiple logical layers to properly function in a cloud service model, as depicted in the cloud stack in Figure 7. The high customization of on-premises system applications created two deficiencies relative to scalability when compared to cloud systems: it inhibited applications from leveraging data from other applications, and it limited administrator knowledge to a small subset of applications, creating pockets of specialization. This means that applications that could leverage data from another application are inhibited from doing so in the cloud, and application administrator knowledge bases are limited to a small subset of applications (Bommadevara et al., 2018).



Repeated again as Figure 11 for reader convenience.

Figure 7. Cloud Security Responsibility Matrix (On-premises Application).
Source: DISA (2017).

2. Cloud Logical Layer Commonalities

Infrastructure as a service, platform as a service, and software as a service models all share interoperability commonalities that permit communication with a diverse array of endpoint devices. While the service model dictates the aggregate resources of each layer under the customer’s control (depicted by the varying swaths of yellow in Figure 7), only standardized structures and exacting protocols below the surface (diagramed in Figure 8) permit a logical layer to communicate with another layer.

Data are able to transmit via the cloud because, succinctly stated, “Abstraction is separation of interface from internals, of specification from implementation” (Saltzer & Kaashoek, 2009, Chapter 1.3.2, para. 2). This implies that the logical interconnected structure of all cloud service models simply requires the proper configuration of interfaces of the endpoint devices. Communications proceed for all compatible interfaces that meet the service model’s precise specifications for the transmitting logical layers.




Figure 8. Cloud Security Alliance Logical Model. Source: CSA (2017).

Infostructure is data either to be stored or processed by computing processes. Data security concerns will directly address the cloud's infostructure. Applistructure comprises applications services used in building applications or the resultant cloud-deployed application itself (CSA, 2017). Application security concerns will directly address the cloud's applistructure. Infrastructure comprises the enormity of the cloud service provider's pooled core computing, networking, and storage resources. Infrastructure security will directly address the cloud's scalable and elastic infrastructure. Metastructure enables communication interoperability protocols between the various layer interfaces to function cohesively; critical configuration and management settings are embedded in metastructure signals. Configuration, management, and administrative security concerns will directly address the cloud's metastructure. The merits of this logical model maps responsibility to service model selection.

It is this lower-level understanding of the uniform interconnection logic that should assist in preventing customer misconfigurations caused by ignorance of standard cloud communication functions at different layers. The oversimplification of consumer and cloud service provider responsibilities, as displayed in Figure 6, promotes ignorance and ultimately leads to misconfigurations. While determining which service model is best suited to the organization's needs is the obvious first step, the consumer must continue to assess responsibilities throughout the more detailed structures within this logical model. The user experience of an on-premises application may be nearly identical to the new cloud

service; however, the wiring, depending on the service model, introduces several permutations that require different customer configurations. This is why understanding how service model selection dictates the likelihood of security misconfigurations is crucial to reducing them.

C. CLOUD DIGITAL FORENSIC INCIDENT-RESPONSE CHALLENGES

Cloud digital forensic incident-response challenges are characterized as a last-resort concern, only needed when simpler tools and precautions did not stop the penetrating threat. However, when digital forensic incident-response is required, it is also an admission that the organization has experienced a significant security breach worthy of digital forensics specialists. Often, the foremost cloud migration challenge is gaining comfort that the cloud service provider will preserve confidentiality for the sensitive data on its information system that is readily accessible by internet connection (Singh & Singh, 2017). Chapter III.C.1.a touches on the encryption complexity used in the cloud and the difficulty it poses just in identifying users—let alone reading their data. Also, cryptographic erasure (also known as crypto-erase or cryptographic erase) is a technique that can quickly address misconfiguration confidentiality violations. This thesis focuses on the higher-impact but lower-probability scenarios compelling digital forensics and the numerous difficulties forensic scientists confront with cloud computing.

The NIST Cloud Computing Forensic Science Working Group (NCC FSWG) identifies the three main categories of cloud challenges:

The cloud exacerbates many technological, organizational, and legal challenges already faced by digital forensics examiners. Several of these challenges, such as those associated with data replication, location transparency, and multi-tenancy are somewhat unique to cloud computing forensics. (NCC FSWG, 2014, p. 1)

In the future, the NCC FSWG's mission dictates that it will develop mitigation strategies not currently addressed by existing cloud forensic science. The working group's research findings tabulated 65 unique cloud forensic challenges. From this analysis, the group then distributed the 65 challenges into nine primary groupings: architecture, data collection, analysis, anti-forensics, incident first responders, role management, legal,

standards, and training challenges (NCC FSWG, 2014). Table 2 shows how many of the 65 challenges fit into each of the nine primary groupings.

Table 2. Sixty-Five Challenges by Primary Categories. Adapted from NCC FSWG (2014).

Primary Category	Count
Data Collection	19
Architecture	18
Legal	13
Analysis	6
Role Management	4
Standards	2
Training	2
Incident First Responders	1
Anti-forensics	-
Total	65

The working group further established additional subcategories, because not all of the primary categories in practice effectively resolved into a single grouping of the nine primary categories; for example, “multi-tenancy” and “data segregation” are two prominent architecture-related subcategories that emerged. Figure 9 expands the nine groupings to show both the nine primary categories and related subcategories.

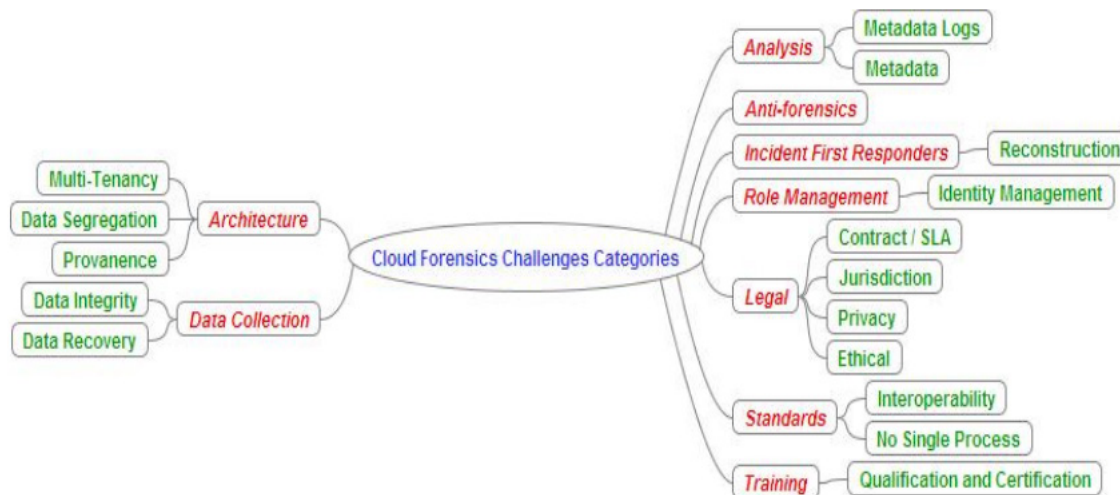


Figure 9. Sixty-Five Cloud Challenges Aggregated into Nine Primary Categories and Relevant Sub-categories. Source: NCC FSWG (2014).

The following list provides an abbreviated description of the nine primary categories, each resulting in a mixture of technical, legal, and organizational complications set forth by the NCC FSWG’s report (2014, pp. 6–7).

- Architecture cloud challenges: challenges that involve the heterogeneity of cloud architectures deployed by cloud service providers; the lack of precedent pertaining to court-admissible standards in the preservation of cloud chain-of-custody artifacts; and the deficit of support services to enable seizure without impacting other tenants.
- Data collection cloud challenges: challenges that arise from the shared and distributed architecture, which can make locating data—which can be in variable locations—difficult. These also include challenges that impede the collection aggregation activities due to the likely incidence of breaching the confidentiality of collocated tenants.
- Analysis cloud challenges: challenges that involve reproducing events from a forensic image virtual asset or highly dynamic storage; validating recovered metadata; synchronizing log file timestamps; and correlating artifacts both inter- and intra-cloud service providers.

- Anti-forensics cloud challenges: challenges that involve tools and techniques specifically purposed for deception and for undermining the integrity of retrieved evidence, or from malicious code that defeats virtual machine isolation safeguards.
- Incident first responder cloud challenges: challenges that arise from a conflict of interest between the cloud service provider’s approach—as an incident first responder and agent—and the approach that the principal data owner wants to employ.
- Role management cloud challenges: challenges that stem from the relative ease of creating fraudulent accounts to conceal identity, or from difficulty establishing non-repudiation for cloud users because of the minimal requirements necessary to open a cloud account.
- Legal cloud challenges: challenges that arise from the laws and regulations of the jurisdiction(s) governing legal access to data; the legal ramifications based upon the location of the cloud service provider, servers, user, and the “effect of applicable treaties or other determinants” (CSA, 2017, p. 37).
- Standards cloud challenges: challenges that stem from inadequate standards, interoperability conventions, policies, procedures, accepted practices, and tools.
- Training cloud challenges: similar standards challenges, training challenges arise due to deficiency in training materials, instructors, and, consequently, trained personnel and naive attempts to retrofit established digital training methodology to the cloud.

The NCC FSWG concluded that all 65 challenges, either grouped or separate, organically divided into three broad sets of challenges: technological, organizational, and legal. Table 3 shows the distribution of the nine primary categories between the three broader classes.

Table 3. Sixty-Five Challenges by Three Broad Classes. Adapted from NCC FSWG (2014).

Technological (47)	Organizational (5)	Legal (13)
Data Collection (19)	Standards (2)	Legal (13)
Architecture (18)	Training (2)	
Analysis (6)	Incident First Responders (1)	
Role Management (4)		
Anti-forensics (-)		

Technological challenges are the driver, as the technological capabilities and limitations dictate the realities that cloud service providers must integrate. The remote delivery of cloud services and the cloud service provider’s capacity as an intermediary give rise to organizational boundary challenges. The multi-geographical operations of cloud service providers create the majority of the legal challenges, as consumers might fall under regulations in multiple jurisdictions.

Cloud forensic investigation hurdles require an interdisciplinary approach because the complications “cannot be solved by technology, law, or organizational principles alone. Many of the challenges need solutions in all three areas” (NCC FSWG, 2014, p. 3). The observed overlapping is noteworthy and extends from the top of the cloud forensics challenges hierarchy into the primary categories, primary subcategories, and related categories. The related category challenges are intersections of the nine primary categories and exist higher in the hierarchy. Figure 10 shows the overlapping of primary category challenges with related categories challenges. For example, in Figure 10, the primary category “legal” identifies “privacy” as a subcategory, and “privacy” overlaps with “challenge #40: Cryptographic key management,” which is definitively technical in nature (NCC FSWG, 2014, p. 51).

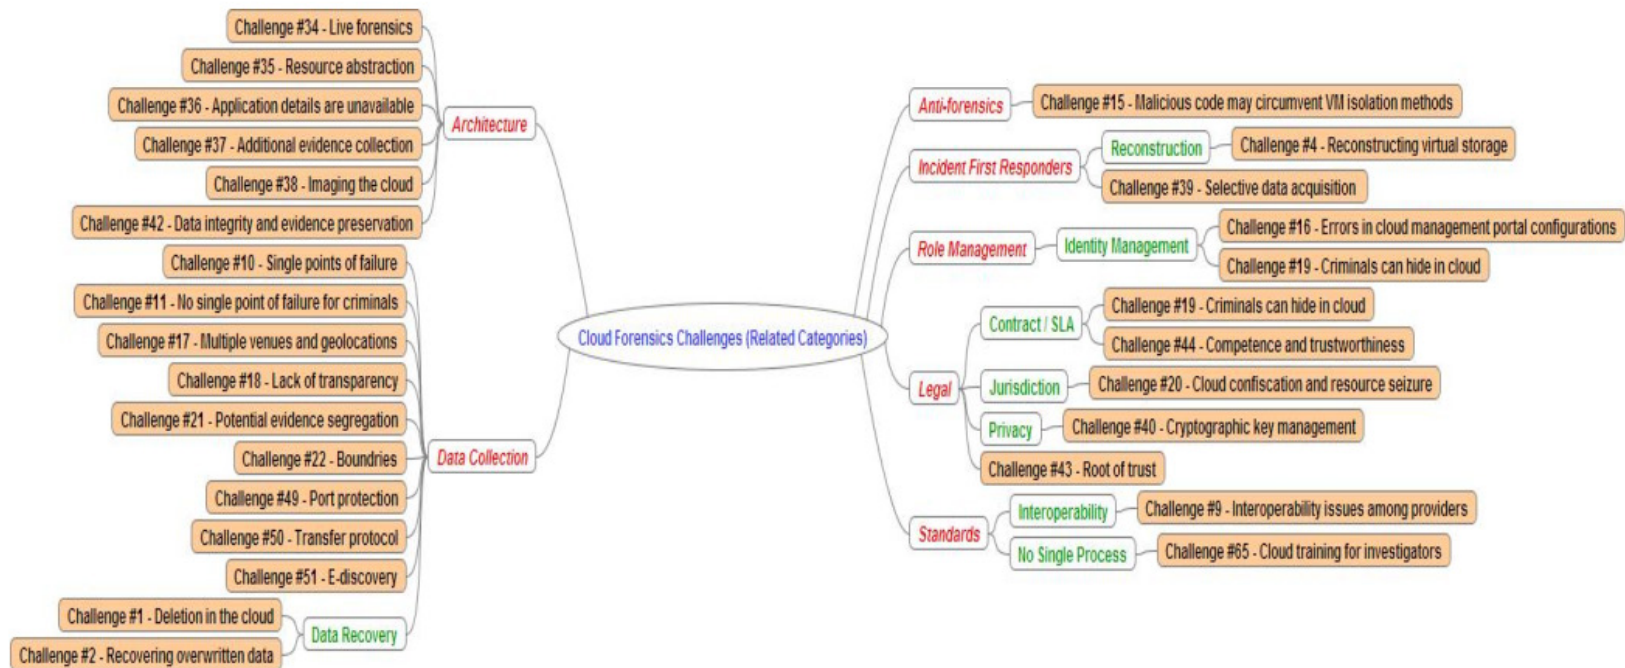


Figure 10. Primary Categories Intersection (Related Categories). Source: NCC FSWG (2014).

There is prevalent overlapping lower in the cloud forensics challenges hierarchy as well. Table 3 showed that the technological category is the largest of the three broad categories, with 47 challenges. In building on the initial analysis of the working group, Table 4 further dissects the technological challenges and reveals a dominant statistic: that over 60 percent of all the forensic challenges are attributed to data collection, architecture, or a combination of both, or either of the two paired with any of the other seven primary categories. This dominant concentration of attributable challenges directed the focus of the technological challenges in this thesis on the specific subset that featured at least data collection or architecture as one of the primary or related categories.

Table 4. Technological Challenges Dominated by Data Collection and Architecture Primary Categories. Source: NCC FSWG (2014).

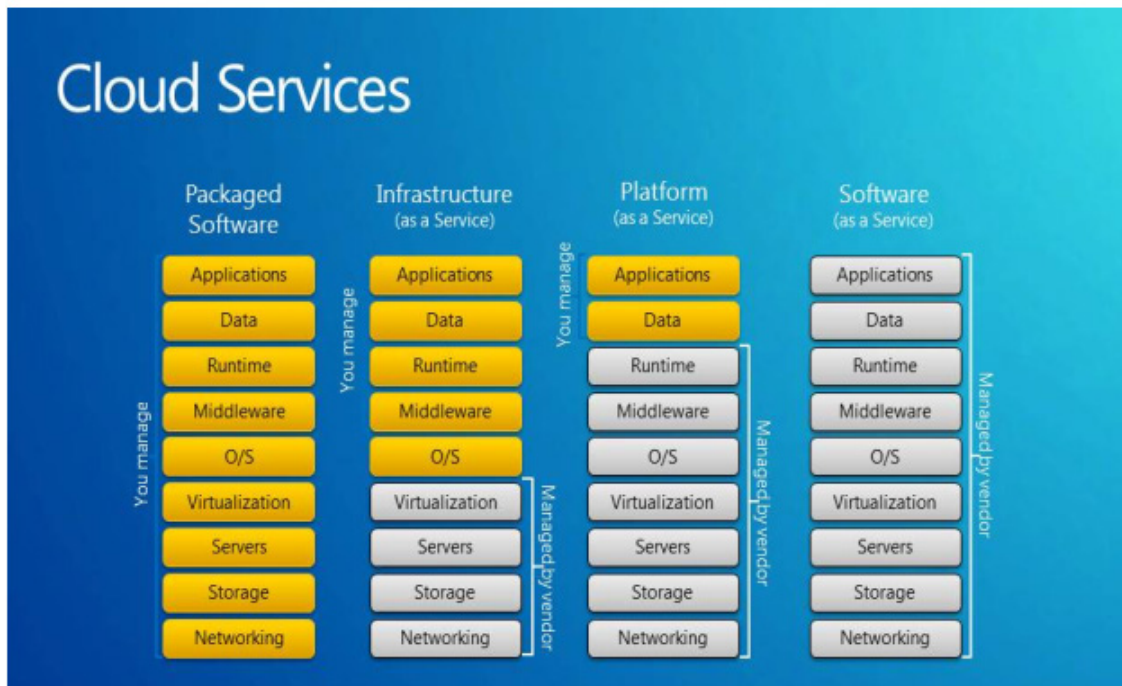
Primary Category	Count	Percentage
Data Collection	10	15.38
Architecture	5	7.69
Data Collection <i>and</i> Architecture	14	21.54
Other	25	38.46
Data Collection/Architecture <i>and/or</i> Other	11	16.92

1. Technical Challenges

a. *Big Differences*

A major difference between the cloud and on-premises systems, in the event of a cyber-incident, is the efficiency of locating stored media, which includes artifacts, log files, and other evidentiary traces (NCC FSWG, 2014). Unlike in on-premises systems, virtualized computer instances in the cloud are not physically attached to local persistent storage; they are attached to storage temporarily, but not after abstraction to enable pooling and dynamic customer provisioning. This decoupling was raised as a concern in Chapter III.A.3 in discussing abstraction as part of detailing Figure 5, the cloud reference service model stack. Comparing the leftmost column in Figure 11, which represents an on-

premises application, against each of the cloud service models in the other columns shows that in each of the service model deployments, storage is designated as a cloud service provider responsibility. The NCC FSWG characterized the separation of a virtual machine from local persistent storage: “Thus, the operational security model of the application, which assumes a secure local log file store, is now broken when moved into a cloud environment” (2014, p. 36). Locating storage in the cloud is more burdensome and represents the first order of significance of breaking the operational security model exemplified by self-contained on-premises system applications.



Repeated previously from Figure 7 for reader convenience.

Figure 11. Cloud Security Responsibility Matrix (On-premises Application).
Source: DISA (2017)

Identifying digital forensic evidence in the cloud is more taxing because the generation of evidence sources in the cloud dramatically departs from on-premises systems and represents the second-order consequence of breaking the application operational security model (NCC FSWG, 2014). The NCCFSWG observed that in the cloud, “User-

based login and controls are typically in the application rather than in the operating system” (2014, p. 37). Cloud applications have the ability to log activity, but the application developers must specifically include this functionality; this is a significant departure from the assurance—with on-premises systems—that the operating system will log activity without fail. Infrastructure as a service and platform as a service cloud applications are developed by the organization that likely will be using them, and can therefore enforce high-assurance security features in development. But organizations do not have such enforcement rights over externally deployed software as a service applications. Similarly, cloud application developers determine the generation, availability, controls, documentation, and storage management capabilities of their applications (NCC FSWG, 2014). In on-premises systems, the operating systems dependably and centrally manage the consistent generation and storage of valuable evidence traces, and the information is well documented. Previously, major software development firms such as Microsoft and Oracle, with large teams solely dedicated to documentation, developed the large market share of applications. Cloud computing empowers smaller organizations by providing a development environment and tools that rival the major software development firms. However, to date, there is a development gap between conventions that were standard in on-premises applications and conventions that are needed for cloud architectures. As a result, acquiring evidentiary traces has become potentially more burdensome in all cloud service models; additionally, related documentation has diminished materially.

As non-repudiation is not as rigorously enacted for actors external to an organization, identity management represents another new challenge in the cloud. Specifically, identification and authentication pose unique challenges for cloud environments; cloud technologies and policies do not sufficiently enforce unique identities (NCC FSWG, 2014). As an example, Figure 12 shows the few steps required to begin using the AWS platform.

How do I create and activate a new Amazon Web Services account?

Issue

I'm getting started with Amazon Web Services. How do I create and activate a new AWS account?

Resolution

Here's an overview of the process:

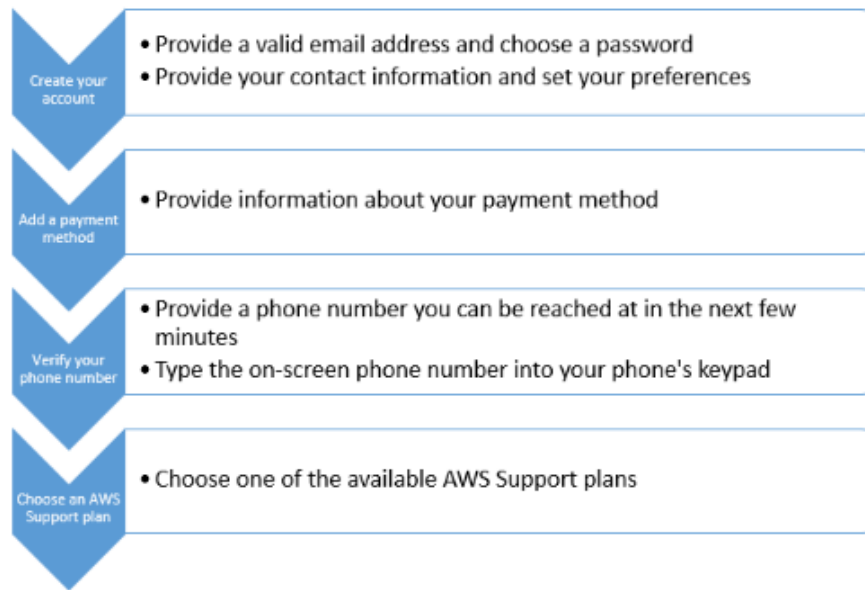


Figure 12. Requirements to Create an AWS Account.
Source: Amazon Web Services (2015).

A related aggravating factor is that cloud technologies decouple user identification credentials from a corresponding physical workstation (NCC FSWG, 2014). The binding of identification credentials with a local physical object is fundamental in the forensic collection of network metadata (NCC FSWG, 2014). The ability to ascribe metadata is in turn integral in establishing non-repudiation, and the working group asserts that, “there is no mandatory non-repudiation methods implemented in the cloud” (NCC FSWG, 2014, p. 40). What this means is non-repudiation methods are available but not standard developers would need to include the non-repudiation methods in development. Additionally, cloud technologies further impede the acquisition of network metadata by concealing cloud

identities, obfuscating network proxy services and complex encryption schemes (NCC FSWG, 2014).

b. Exacerbated Challenges

The forensic investigator faces challenges in correlating evidence due to the sheer enormity of the cloud. For example, unifying log formats in on-premises systems is consistently a pain point. In cloud computing, unified logs allow for “a single, efficient, performant API for capturing messaging across all levels of the system... The system implements global settings that govern logging behavior and persistence...” (Apple, n.d.). The significance of unified logs is their ability to remotely access a single repository with system-wide settings that persist. Criticisms of on-premises systems have centered on volume and the heterogeneity of log formats. The cloud has introduced additional proprietary log formats, which have proven a challenge to convert or unify because of the forensic investigator’s concerns of omitting applicable data between each of the relevant sources (NCC FSWG, 2014). A unified logging system would overcome the proprietary formats and volume through a customer application programming interface. The NCC FSWG perceived the decentralization of cloud resources as a sprawl of “multiple physical machines that are spread across multiple geographical regions, between the cloud infrastructure and remote web clients including numerous end points” (2014, p. 21). In other words, the cloud requires potential data collection over a more expansive technological footprint. Ultimately, the influx of devices spread over this technological footprint has also affected forensic data collection by increasing the probability of inconsistent log timestamps between devices (NCC FSWG, 2014).

For digital forensics, the rewriting phenomena of cloud metadata becomes even more confounding than the straightforward issue of correlating timestamps. Investigations often key on establishing timelines and vetting the alibis of suspects. The cloud makes establishing the forensic timelines an impossibility. In on-premises systems, forensic investigations seized on “MAC times.” MAC is a mnemonic for the operating system timestamps: modification time, access time, and creation time. The timestamp names are self-explanatory, but for clarity, the modification time is the last time a file changed, access

time is the last time an access function was performed on a file (e.g., open, print, etc.), and creation time is either the time a file was created or the time it was copied onto a file system. Timestamp metadata is fundamentally different in the cloud; the NCC FSWG identifies that metadata “may be changed as the data is migrated to and within the cloud. Metadata may also be changed during the collection process, giving rise to both authentication challenges and spoliation worries” (2014, p. 22). For on-premises systems, there is already a problem known as time-stomping—when bad actors change the date of a timestamp. However, the cloud arbitrarily rewriting metadata is more problematic than on-premises time-stomping; the NCC FSWG’s findings warn that the preservation of cloud metadata is dubious due to the common processing of data transference within the cloud. Therefore, correlating evidence in the cloud is irresolvable because of the metadata inclusivity, log format incongruity, and timestamp ambiguity predicated on the proliferation of end devices

c. Concentration of Data Collection and Architecture Challenges

Cloud service providers are encountering the current technological limits of rapid provisioning and straightforward evidence retrieval. Collecting evidential traces in the cloud is difficult for a number of reasons. One of the early challenges in locating evidence is navigating the cloud ecosystem, which proves to be a web of dependencies (NCC FSWG, 2014). The NCC FSWG shares an example of the enmeshed relationships a forensic investigator may have to unwind: “A cloud Provider that provides an email application (SaaS) may depend on a third-party provider to host log files (i.e., PaaS), which in turn may rely on a partner who provides the infrastructure to store log files (IaaS)” (p. 28). A likely next challenge is in locating the data, which is, itself, a moving target; data movement could be initiated by the cloud service provider, an enmeshed cloud service provider partner, or the consumer (NCC FSWG, 2014). If the data is located it may relocate again, and the forensic investigator’s lack of understanding of the native cloud environment may raise uncertainty about the completeness of the evidence (NCC FSWG, 2014). Essentially, the forensic investigators must navigate a moving crime scene.

Imaging (saving) evidence from the cloud faces the added challenge that related data is so voluminous that it not possible to image all available evidence; but a less-than-

complete data collection poses ramifications to chain of custody/legal admissibility (NCC FSWG, 2014). This volume is not driven by more data but surprisingly by overwhelming amounts of associated metadata and by deciding which metadata to discard without jeopardizing the chain of custody. The NCC FSWG reported a second quandary specific to forensic scientists attempting to reconstruct virtual storage: “Imaging of media has an added level of complexity in some cloud environments which could cause damage to the original media and add the risk of being sued” (2014, p. 21). The possibility that imaging can cause damage to cloud service provider equipment does not lend itself to cloud service provider cooperation. In contrast, when resources are confiscated from on-premises systems, the seizure will not negatively affect multiple tenants, and on-premises imaging has not been reported to damage equipment. Ultimately, managing the scope of cloud investigations is unpredictable given a lack of insight into the cloud service provider’s proprietary protocols, the unknown geo-location of the data, the multiplicity of data centers in which data can reside, and the possibility that data could migrate after initial location—all of which makes an opportunity for additional collection unlikely (NCC FSWG, 2014).

d. Deletion and Data Remanence

Cloud deletion and data remanence challenges are more rightfully framed as uncertainties. These uncertainties stem from the previous discussion about how the cloud breaks the operational security model of applications. The challenges raise four questions: 1) Is it possible to attribute deleted data to a unique user? 2) Is data recoverable from a deallocated¹ virtual machine (VM) instance? 3) Is overwritten data recoverable from a deallocated VM? 4) How does the implementation of dynamically allocated storage amplify the conclusions of the preceding three questions? The first open question is directed at scrutinizing the mechanics of the VM deallocation process. When a VM is deallocated, what is known is that the node pointing to the VM’s storage is deleted (NCC FSWG, 2014). The deleted node was previously pointing to a remote physical storage device, which was formerly storing data for the deallocated VM. What remains unknown is whether access to and recoverability of the uncoupled data is possible. If data of a

¹ This section uses the terms *deleted* and *deallocated* interchangeably.

previously deallocated VM is attributable to specific users, this may reveal difficulties to attribute deletions in platform as a service and infrastructure as a service tools (NCC FSWG, 2014).

With respect to the second question, the NCC FSWG asserted that “no research has been conducted on determining what data is associated with removed VM instances” (2014, p. 29). Thus, it is unclear whether any evidential traces or even the complete VM would be retrievable. The third question scrutinizes the effects of the VM deallocation process and seeks to determine if an analogous recoverability is technically feasible in the cloud’s multi-tenant architecture when a deallocated VM is reallocated.

The fourth question, regarding dynamically allocated storage, is focused on how quickly previously deallocated VMs are reallocated. Dynamic storage is a process in which cloud service providers sift through deallocated VMs to optimize performance of reading data by quickly processing VM reallocations (NCC FSWG, 2014). To determine the expediency of dynamic storage algorithms, forensic scientists must determine if it is viable to do any of the following: attribute deallocated VMs to specific users, recover evidential traces from deallocated VMs, or recover overwritten data from deallocated VMs. The answers can also possibly constrain the extent of backups that cloud service providers are able to retain based on storage limitations (NCC FSWG, 2014).

e. Interoperability

When it comes to security, that complete interoperability is not the goal because the security breach of one system could lead to contagion. Another relationship between interoperability and security to consider is the degree to which interoperability is welcome; cloud service providers tend to desire less interoperability and consumers more (Kostoska, Gusev, & Ristov, 2016). The interoperability context for digital forensics is in reducing the proprietary knowledge needed by minimizing wide divergence between cloud service providers. The opaqueness of cloud service provider architectures and operating environments perpetuates the current lack of interoperability standards. The lack of transparency is a significant hurdle in understanding the cloud service provider environment because the interoperability is both abstract and remote. The NCC FSWG

commented on the difficulties of correlating evidence in the cloud, stating, “there is no interoperability between cloud Providers” (2014, p. 21). Cloud interoperability issues are not limited to correlating cloud service provider logs. Interoperability issues are relevant to the development and subsequent documentation of applications; specifically, the NCC FSWG observed that “Private and confidential details of cloud based software/applications used to produce records are typically unavailable to the investigator” (2014, p. 32). This reality leaves investigators ignorant of the application’s full functionality, data structures, and potential vulnerabilities that may have led to the investigation.

The confounding aspect of interoperability is that the cloud integrates multiple sophisticated technologies, cloud service providers, servicing counterparties, logical layers, hardware, and endpoint devices. The NCC FSWG identified that “the trustworthiness and integrity of cloud forensics data is a challenge because of the dependence on the cumulative integrity of multiple layers of abstraction throughout the cloud system” (2014, p. 34). A cloud service provider’s trustworthiness is compromised if any of the multiple parties or technological interchanges are compromised

f. Threat Enhancements

Similar to many legitimate organizations, criminal enterprises are also attracted to the cloud’s pay-as-you-go cost savings. The cloud offers criminals superior computing power for a fraction of the cost. For example, the cloud can rapidly provision enough nodes to perpetrate a botnet attack without having to actually compromise the underlying computers that a botnet operating outside of the cloud would need. The NCC FSWG recognized that the cloud has been used to enhance the threat capabilities of bad actors: in 2009, a botnet used Google’s AppEngine as its command-and-control-network; the NCC FSWG report also mentions the existence of a cloud password cracking service, and how a security researcher used Amazon’s EC2 services to crack Wi-Fi passwords (2014).

Criminals can also remain anonymous in the cloud, inexpensively. As was shown in Figure 12, most cloud service providers only require users to provide a name, address, and credit card to register for an account. Identity theft trends portend that criminals should be able to easily obtain stolen identities to register many fraudulent accounts (NCC FSWG,

2014). The fraudulent accounts are yet another hurdle for forensic investigators to bypass to determine the authentic identity behind a cloud username.

Furthermore, criminals can decentralize their operations by utilizing multiple cloud service providers. The NCC FSWG offered an example of how an enterprise can use decentralization to obscure related criminal activities at separate cloud service providers: “A criminal organization can choose one cloud provider as a storage solution (e.g., Dropbox), obtain compute services from a second cloud provider (e.g., Amazon EC2), and route all of their communications through a third (e.g., Gmail or Pastebin)” (2014, p. 23). For a forensic investigator, the scope and complexity of reconstructing evidential traces distributed among multiple cloud service providers is significantly more demanding than an investigation limited to a single cloud service provider. Forensic incident response teams must update their threat model to incorporate how bad actors are presently leveraging cloud services in their attacks.

g. Timing

In the cloud, real-time investigations to collect volatile forensic data, also known as *live forensics*, are improbable because of the level of collaboration required between the consumer and cloud service provider (NCC FSWG, 2014). Volatile data is any data for short-term memory purposes that is lost without power. Persistent data is considered non-volatile, and is any data exceeding short-term memory purposes, which remains available in storage media even after losing power. The NCC FSWG underscores that volatile data “can only be collected in real time by placing sensors into the real-time environment” (2014, p. 24). Real-time investigations are improbable because targeting a specific cloud account for these sensors confronts the reality that cloud infrastructures are complex and may be constituted by leased lines from a variety of service providers. The end result is that any volatile evidential traces not in storage will be lost when a suspected VM loses power.

In the event that multiple parties were able to successfully coordinate to conduct a live forensics investigation on a running VM, the NCC FSWG warns that “it is impossible for a third party to verify, after acquisition, that the data collected is correct because the

data is no longer the same as at the time of acquisition” (2014, p. 31). As previously mentioned, cloud environments rewrite metadata, and this would challenge the ability of a third party to verify the collected evidential traces. Additionally, because the system is running, just the act of introducing the live forensics collection tools would change the VM environment, biasing verification against a running system unexposed to the live forensics tools. Another issue complicating the use of live forensics tools on a running system is that the system could have anti-forensic functionality to sabotage the data collection (NCC FSWG, 2014).

Live forensics tools will need to evolve to address the highlighted shortcomings, as VMs contain highly valuable log data with evidential traces captured in system, registry, and network logs. Systems, registries, and networks contain a treasure trove of forensic value. Chapter V.E.1 introduces a Linux distribution that can address several shortfalls in accessing valuable system and registry data. On-premises systems are able to image (save) the volatile data and then take a digital fingerprint (hash) for later comparison with the live forensics collected data. On-premises systems do not experience the phenomena of rewriting metadata, which means the collected data can be successfully verified against the hash of the imaged file. This data collection comparison capability is what cloud live forensics tools will need to achieve.

2. Legal Challenges

The majority of legal challenges concern contractual, jurisdictional/law enforcement, or chain of custody questions.

a. Scope

The remainder of this thesis examines cloud incident-response from the perspective of DoD policy. As is discussed in more detail in Chapter IV, the Federal Risk and Authorization Management Program (FedRAMP) establishes federal policy for cloud computing. The DoD follows FedRAMP+, which is tailored guidance for the DoD’s mission. Chapter V explains the drivers for FedRAMP+ and specific regulatory guidelines for federal cloud adoption, but many of these legal challenges are outside the scope of this thesis.

b. Reliance on Contract

The contract is the sole vehicle that arbitrates and memorializes cloud governance, and the consumer needs to obtain everything upfront, especially for forensic terms of agreement. The NCC FSWG (2014) identified missing contractual terms in service-level agreements as a legal challenge. Specifically, the working group spotlighted the omission of clauses that require the cloud service provider to produce relevant evidence within specific time limits. The multi-tenancy inherent in the cloud limits the ability of any single consumer to consummate a highly negotiated custom solution. In stark contrast, the negotiating flexibility is generally much greater in bilateral agreements. Because the cloud is shared by multiple clients, each consumer is also therefore limited in its customization. Notwithstanding, it is expected that software as a service and infrastructure as a service will provide the greatest negotiating space: for software as a service, primarily because the multi-tenancy factor is minimized by the cloud access via applications; and for infrastructure as a service, because it most resembles a green data center. The least tailored solution would be offered by platform as a service.

To maximize service agreements, in 2012 the Chief Information Officer Council and Chief Acquisition Officers (CAO) Council published guidance to improve enforceability. They advise that every agreement must comply with federal laws and regulations and remain consistent with agency governance, and all service-level agreements should aim to include measurable conditions that hold the cloud service provider accountable. Consumers should focus on the definitions used in each agreement. The cloud has matured to a common lexicon; however, wide variability exists among cloud service providers with respect to defined terms and related metrics (Chief Information Officer Council & Chief Acquisition Officers Council [CIO & CAO Councils], 2012). The Councils highlight that many service-level agreements are negligently ambiguous regarding ramifications for underperformance. Lastly, the Councils explain that non-disclosure agreements are required whenever federal data is transmitted.

3. Boundary Challenges

The majority of boundary challenges center on the few certainties the consumer or forensic scientists can establish about cloud service provider environments. In addition to challenges of unfamiliarity with cloud environments, other boundary challenges include loss of governance and decoding the cloud service provider environment.

a. Loss of Governance

“There is no cloud, it’s just someone else’s computer.” This unattributed saying resonates with consumers who experience loss of governance through the constraints presented by cloud service providers. As the cloud is effectively someone else’s computer, cloud service providers may require permissions or limit the scope for routine security measures such as port scanning, penetration testing, and other methods of formally analyzing security vulnerabilities. Furthermore, when consumers transition from on-premises systems, they will find gaps within their existing security policies and how they interplay with the contracted terms and conditions of an executed service-level agreement. The CIO and CAO Councils (2012) warn that the lack of agency will require consumers to seek alternative methods to achieve equal measures or higher of confidentiality, integrity, and availability (the CIA triad) within a cloud service provider environment.

b. Decoding the Cloud Service Provider Environment

This thesis has expressed repeatedly that cloud interoperability standards are, to date, grossly inadequate. Counterparties reliant on cloud service providers have found it difficult to understand how the cloud service provider environment operates because cloud technologies are relatively immature, generally remote to consumers, and use proprietary protocols. All these factors limit the advancement of professional knowledge, standards, training, and experimentation. The NCC FSWG assessed cloud training for investigators in its report, finding, “Most digital forensic training materials ... are not applicable in cloud environments ... and there is an absence of proper tools to effectively investigate the cloud computing environment.... Only few standard operating policies are in place ... making the approach more trial and error than scientific” (2014, p. 41). Acquiring knowledge is problematic because, unlike in on-premises systems, existing cloud record logs are held by

custodians. The data held by custodians usually only becomes public during court testimony (NCC FSWG, 2014). The withheld details would greatly help consumers acquire knowledge, test validation, and conduct training and professional development. Standards development is especially necessary considering the unsatisfactory interoperability progress to date. The NCC FSWG also has concerns regarding the standards development process; “there is no one accepted standard, and the majority of organizations are creating their own SOPs, which may or may not be based on an existing process model” (2014, p. 40). If cloud service providers granted access to academics and researchers, their work could go a long way to address the many outstanding uncertainties of the cloud service provider environment.

IV. ADAPTING DoD REQUIREMENTS FOR CLOUD ARCHITECTURES

This chapter marks a transition from general cloud computing to a much more specific DoD-centric perspective. Chapter I established four security-related issues facing the DoD in its aggressive adoption of cloud computing technologies. Chapter II introduced definitions, vulnerabilities, and threats specific to cloud computing for use in analyzing risk management disparities with on-premises information systems. Chapter III discussed the impact of service model selection and the shared responsibility model, and imparted a general awareness about technical cloud computing digital forensic incident-response challenges. These preceding chapters provided a necessary foundation about how the cloud is structured, how it operates, and which organization is responsible for which layer. We are now oriented to assess the efficacy of FedRAMP's direction in navigating cloud computing security risks.

A. FedRAMP

FedRAMP (the Federal Risk and Authorization Management Program) is a 2011 federal policy that details the minimally required security authorization procedures with which an agency must comply when engaging with a cloud service provider for contracted cloud services. FedRAMP was specifically drafted to direct federal cloud computing acquisitions, and its goal was to accelerate adoption of cloud services and enforce standardized cybersecurity requirements government-wide. A 2011 memo by Steven VanRoekel, the federal chief information officer, was FedRAMP's documented debut. The memo established policy for protecting federal data in cloud architectures, and specifically prescribes seven tasks of compliant agencies:

- i. Use FedRAMP when conducting risk assessments, security authorizations, and granting ATOs [authorization to operate] for all Executive department or agency use of cloud services;
- ii. Use the FedRAMP PMO [project management office] process and the JAB [Joint Authorization Board]-approved FedRAMP security authorization requirements as a baseline when initiating, reviewing, granting and revoking security authorizations for cloud services;

- iii. Ensure applicable contracts appropriately require CSPs [cloud service providers] to comply with FedRAMP security authorization requirements;
- iv. Establish and implement an incident response and mitigation capability for security and privacy incidents for cloud services in accordance with DHS guidance;
- v. Ensure that acquisition requirements address maintaining FedRAMP security authorization requirements and that relevant contract provisions related to contractor reviews and inspections are included for CSPs;
- vi. Consistent with DHS guidance, require that CSPs route their traffic such that the service meets the requirements of the Trusted Internet Connection (TIC) program; and
- vii. Provide to the Federal Chief Information Officer (CIO) annually on April 30, a certification in writing from the Executive department or agency CIO and Chief Financial Officer, a listing of all cloud services that an agency determines cannot meet the FedRAMP security authorization requirements with appropriate rationale and proposed resolutions. (VanRoekel, 2011, p. 5)

After a U.S. federal agency adheres to the specific FedRAMP standards, the agency is deemed the responsible party based on the risk undertaken within the cloud environment. In addition to the procurement requirements, FedRAMP identifies seven critical security policy controls that need to be addressed in every cloud implementation: security-authorization requirements, continuous monitoring, incident response, key escrow, forensics, two-factor authentication with Homeland Security Presidential Directive 12 (HSPD-12), and auditing (CIO & CAO Councils, 2012).

To fulfill the security authorization requirements for cloud computing, all U.S. federal agencies must categorize their data in accordance with Federal Information Processing Standards (FIPS) 199 and 200, and must additionally contract with cloud service providers to implement security countermeasures commensurate to the impact level for the environment that will process the data (CIO & CAO Councils, 2012). The CIO and CAO Councils proclaim that cloud service providers must adhere to the continuous monitoring standards proposed in NIST SP 800-137. The cloud service provider must also follow DHS's directions for documenting changes made that pertain to defending the cloud service provider environment, integrating new Federal Information Security Management

Act (FISMA) requirements, and ensuring there is sufficient time to make efficient, forward-looking risk management decisions (CIO & CAO Councils, 2012).

The contracting U.S. federal agency must oversee incident response cloud computing standards, but the incident response activities are fulfilled by the cloud service provider. The cloud service provider must adopt NIST SP 800-61, *Computer Security Incident Handling Guide*, which explains how to select the appropriate response for information system incidents. Adherence to SP 800-61 guidance allows the agency to cooperate with DHS's United States Computer Emergency Readiness Team (U.S. CERT) when responding to dynamic changes to risk within cloud service provider environments (CIO & CAO Councils, 2012). Clear expectations need to be documented regarding acceptable responses to recover from a security incident. The federal agency needs to ensure that any contractual language clearly defines expectations for satisfactory corrective action by the cloud service provider.

U.S. federal agencies can comply with the key escrow requirement by assessing the cloud service provider's encryption practices against NIST SP 800-152, *A Profile for U.S. Federal Cryptographic Key Management Systems*. Additionally, HSPD-12, *Policy for a Common Identification Standard for Federal Employees and Contractors*, is a required federal standard for two-factor authentication that features the use of compliant personal identity verification cards (CIO & CAO Councils, 2012). It is a best practice for contractual agreements to recognize HSPD-12 guidance.

Federal agencies can comply with the FISMA log preservation requirement by verifying that the cloud service provider's environment is highly similar to the best practices identified in NIST SP 800-92, *Guide to Computer Security Log Management*. FISMA also requires all cloud service provider personnel who log data to have appropriate clearances (CIO & CAO Councils, 2012).

B. DoD CLOUD COMPUTING SECURITY REQUIREMENTS GUIDE

1. FedRAMP+

Cloud requirements for the DoD exceed requirements for other federal government agencies; for that reason, the DoD issued the *Cloud Computing Security Requirements*

Guide (DISA, 2017), which describes FedRAMP+. FedRAMP+ adds DoD-specific security controls to fulfil the DoD’s mission requirements. It is applicable to DoD agencies and dictates that their cloud service providers will be assessed according to the *Cloud Computing Security Requirements Guide*. A cloud service provider can lose its authorization if it is unable to maintain compliance with FedRAMP+ requirements (DISA, 2017). FedRAMP+ determines security controls based on the sensitivity of the data to be processed in the cloud, defined as security objectives, combined with the potential impact in the event of a CIA triad violation, defined as impact levels. Interestingly, FedRAMP+ impact levels are only based on confidentiality and integrity objectives. Availability objectives are unaddressed and it is the mission owner’s responsibility to appraise the cloud service provider’s availability during the service provider selection process (DISA, 2017). Additionally, FedRAMP+ provides an opportunity to further tailor individual security controls with “security controls/enhancements” (DISA, 2017, p. 44). Accordingly, the *Cloud Computing Security Requirements Guide* directs that the cloud contract is the appropriate space to address available security controls/enhancements.

2. FedRAMP+ Security Controls/Enhancements

FedRAMP+ is the cloud computing tailored approach to NIST 800-53 security controls, and it is applicable to all impact levels except Level 2 (described in more detail in the next section). These controls “were selected primarily because they address issues such as the Advanced Persistent Threat (APT) and/or Insider Threat, and because the DoD ... must categorize its systems in accordance with CNSSI 1253, beginning with its baselines, and then tailoring as needed” (DISA, 2017, p. 44). CNSSI 1253 is the Committee on National Security Systems Instruction No. 1253, titled *Security Categorization and Control Selection for National Security Systems* (CNSS, 2014). A comparison of security controls, as seen in Equation 2 (adapted from DISA, 2017), indicates that 32 CNSSI 1253 controls were added to the NIST SP 800-53 moderate baseline and 88 NIST 800-53 moderate controls were subtracted from the CNSSI 1253 moderate baseline (DISA, 2017).

$$\begin{aligned}
 \text{FedRAMP Plus} &= \text{NIST 800-53 Moderate baseline} + 32 \text{ CNSSI 1253 M-M-x C/CEs} \\
 &\text{or} \\
 \text{FedRAMP Plus} &= \text{CNSSI 1253 M-M-x} - 88 \text{ NIST 800-53 Moderate baseline controls}
 \end{aligned}
 \tag{2}$$

The increase of control enhancements correlates directly to rising impact levels (DISA, 2017). For example, Impact Level 6 uses a classified information overlay that prescribes an additional 94 security controls and enhancements (DISA, 2017).

3. Impact Levels

The security control baseline for all four cloud-computing impact levels is moderate confidentiality and moderate integrity. Missions that require systems with higher confidentiality or integrity “must deploy to facilities assessed using CNSSI 1253 high baselines through the DoD RMF (typically a DoD facility) or contract for the added security from a commercial [cloud service provider]” (DISA, 2017, p. 25). The four information impact levels are 2, 4, 5, and 6, as shown in Figure 13. Level 2 includes all data that has been cleared for public release (i.e., non-controlled unclassified information), Level 4 includes controlled unclassified information (CUI), Level 5 includes CUI but can also process unclassified national security systems, and Level 6 processes classified information up to Secret (DISA, 2017). To accommodate Level 5, FedRAMP+ uses an additional security controls/enhancements structure. Level 6 requires the complete dedication of cloud infrastructure to the DoD or federal government community and is therefore not considered a “commercial” service (DISA, 2017). All processed data must be within a legal U.S. jurisdiction, as indicated in column 4 of Figure 13, which resolves potential legal jurisdictional challenges. FedRAMP+ has multi-tenant separation requirements based on the impact level, as indicated in column 6 of Figure 13.

IMPACT LEVEL	INFORMATION SENSITIVITY	SECURITY CONTROLS	LOCATION	OFF-PREMISES CONNECTIVITY	SEPARATION	PERSONNEL REQUIREMENTS
2	PUBLIC or Non-critical Mission Information	FedRAMP v2 Moderate	US / US outlying areas or DoD on-premises	Internet	Virtual / Logical PUBLIC COMMUNITY	National Agency Check and Inquiries (NACI)
4	CUI or Non-CUI Non-Critical Mission Information Non-National Security Systems	Level 2 + CUI-Specific Tailored Set	US / US outlying areas or DoD on-premises	NIPRNet via CAP	Virtual / Logical Limited "Public" Community Strong Virtual Separation Between Tenant Systems & Information	US Persons ADP-1 Single Scope Background Investigation (SSBI)
5	Higher Sensitivity CUI Mission Critical Information National Security Systems	Level 4 + NSS & CUI-Specific Tailored Set	US / US outlying areas or DoD on-premises	NIPRNet via CAP	Virtual / Logical FEDERAL GOV. COMMUNITY Dedicated Multi-Tenant Infrastructure Physically Separate from Non-Federal Systems Strong Virtual Separation Between Tenant Systems & Information	ADP-2 National Agency Check with Law and Credit (NACLC) Non-Disclosure Agreement (NDA)
6	Classified SECRET National Security Systems	Level 5 + Classified Overlay	US / US outlying areas or DoD on-premises CLEARED / CLASSIFIED FACILITIES	SIPRNET DIRECT With DoD SIPRNet Enclave Connection Approval	Virtual / Logical FEDERAL GOV. COMMUNITY Dedicated Multi-Tenant Infrastructure Physically Separate from Non-Federal and Unclassified Systems Strong Virtual Separation Between Tenant Systems & Information	US Citizens w/ Favorably Adjudicated SSBI & SECRET Clearance NDA

Figure 13. Impact Level Comparison. Source: DISA (2017).

4. Data Spills

Human error can lead to cloud security misconfigurations and those misconfigurations can lead to data spills. An example of such a data spill was provided in Chapter 1.B.2, in which a federal contracting employee of Booz Allen Hamilton was responsible for spilling data collected on behalf of the National Geospatial-Intelligence Agency. Misunderstandings and inadequate training contribute to this type of cloud security misconfiguration, despite the robust encryption methods employed by cloud service providers to enforce rigorous confidentiality standards. The *Cloud Computing Security Requirements Guide* recognizes the need to employ different sanitizing methods for data spills in the cloud as compared to on-premises systems:

Cloud environments present a unique challenge for data spill response... [Cloud service provider] use of storage virtualization makes physical data locations difficult to ascertain. This makes physical sanitization methods non-viable for data spill remediation in cloud services. These challenges require a method for mitigating data spill cyber incidents that occur in the cloud. (DISA, 2017, p. 80)

The *Cloud Computing Security Requirements Guide* endorses cryptographic erase as the primary method to address cloud data spills. Cryptographic erase is credited by the DoD as “high-assurance data destruction”— “media sanitization is performed by sanitizing the cryptographic keys used to encrypt the data, as opposed to sanitizing the storage locations on media containing the encrypted data itself” (DISA, 2017, p. 113). Cryptographic erase also accommodates “partial sanitization,” in which a subset of the data is sanitized, but this requires the use of unique keys for each subset (DISA, 2017). Cryptographic erase paired with deleting files is more expedient than physically sanitizing a cloud service provider environment. However, cryptographic erase is only effective for encrypted data. Therefore, the mission owner is responsible for ensuring that all DoD data is encrypted for data-at-rest. Furthermore, the DoD must have exclusive control of both the encryption keys and key management; this facilitates the DoD’s ability to remediate unilaterally, high-assurance data destruction, without any cloud service provider cooperation (DISA, 2017). In the event of a data spill, the mission owner can delete the keys related to the unauthorized data. Additionally, the mission owner must follow the

extra measures of sanitizing any existing data in an unencrypted state (e.g., memory in an allocated virtual machine) and remediate backups and mirrored storage (DISA, 2017).

However, cryptographic erase is not a panacea. This technology is an effective tool to resolve non-malicious misconfiguration-related data spills generated by human error, but it would likely prove ineffective against data spills initiated by malicious code. Cryptographic erase would be unable to contain a running process while data is still in use. Because of the extra sanitizing measures, the mission owner needs to understand for any data in unencrypted states will put that data at risk, open to precisely the types of vulnerabilities that malicious code attempts to exploit. For example, cryptographic erase would be ineffective if a malicious process reads confidential data of another running process while both processes are in the same memory space. Additionally, cryptographic erase is only effective in infrastructure as a service— and some platform as a service— cloud deployments when the mission owner determines exactly how the data is stored. The *Cloud Computing Security Requirements Guide* acknowledges that “the Mission Owner relies on the [cloud service provider] and the security posture of its SaaS offering for the protection of DoD information” (DISA, 2017, p. 100). Each mission owner must perform risk analysis assessments and weigh whether the DoD can potentially accept the risk (DISA, 2017).

5. Incident Response

FedRAMP+ lists nine security controls specific to cloud computing incident response with increasing security controls and enhancements. Table 5 lists the incident response security controls—each parenthetical number indicates the stipulation of additional security controls and enhancements.

Table 5. DoD FedRAMP+ Incident Response Security Controls.
Adapted from Metheny (2013).

Incident Response (IR) Security Controls				
ID	Control Description	Low Impact	Moderate Impact	High Impact
IR-1	Incident Response Policy and Procedures	IR-1	IR-1	IR-1
IR-2	Incident Response Training	IR-2	IR-2	IR-2 (1) (2)
IR-3	Incident Response Testing	Not Selected	IR-3 (2)	IR-3 (2)
IR-4	Incident Handling	IR-4	IR-4 (1)	IR-4 (1) (2) (3) (4) (6) (8)
IR-5	Incident Monitoring	IR-5	IR-5	IR-5 (1)
IR-6	Incident Reporting	IR-6	IR-6 (1)	IR-6 (1)
IR-7	Incident Response Assistance	IR-7	IR-7 (1) (2)	IR-7 (1) (2)
IR-8	Incident Response Plan	IR-8	IR-8	IR-8
IR-9	Information Spillage Response	Not Selected	IR-9 (1) (2) (3) (4)	IR-9 (1) (2) (3) (4)

The book *Federal Cloud Computing* describes many of these controls best (Metheny, 2013):

- Security control IR-1 requires that incident policies are updated at least every three years and procedures annually.
- Security control IR-2 requires the cloud service provider to define the period in which it will provide new user training, and also stipulates annual refresher training.
- Security control IR-3 requires that the cloud service provider to test and document incident response effectiveness at least annually.
- Security control IR-4 requires the cloud service provider to implement an incident handling program consistent with CJCSM 6510.01B, *Chairman of the Joint Chiefs of Staff Manual: Cyber Incident Handling Program*; the program includes: preparation, detection and analysis, containment eradication, and recovery steps.
- Security control IR-5 requires that the cloud service provider collect and analyze cyber incidents.

DISA (2017), however, best describes IR-6:

- Security control IR-6 “requires cloud service providers to report cyber incidents to the Department of Homeland Security (DHS) United States Computer Emergency Readiness Team (U.S. CERT) and the consuming Federal Agencies” (p. 129). Reporting also requires the submission of an initial incident report within one hour of discovery.

And *Federal Cloud Computing* rounds off the descriptions of IR-7 through IR-9 (Metheny, 2013):

- Security control IR-7 requires that the cloud service provider maintain a collaborative relationship with all external stakeholders.
- Security control IR-8 requires that the cloud service provider specify the criteria of reportable incidents, metrics to assess incident response effectiveness, and approved personnel.
- Security control IR-9 requires that the cloud service provider document its procedures in responding to data spills.

The *Cloud Computing Security Requirements Guide* specifically recognizes the need for employing digital forensic methods in the cloud as compared to on-premises systems: “Digital forensics in the cloud has many challenges.... the CC SRG [*Cloud Computing Security Requirements Guide*] provides initial guidance regarding the DoD requirements for enabling and performing Cloud Forensics” (DISA, 2017, p. 133).

V. RISK MANAGEMENT ALTERNATIVES: TRANSFORMATIONAL MIGRATION

At the outset, Chapter I.B.1-4, this thesis intended to address four cloud-specific issues: customer misconfigurations, cloud leaks, complications in the implementation of security controls, and altered digital forensic incident-response challenges. This chapter discusses potential DoD risk management solutions addressing the four cloud-specific conditions and is able to connect many of the recommendations to the proposed concept of transformational migration.

Statements from the former NSA CIO's interview, first discussed in Chapter II.B.1.d, illuminate core tenets of transformational migration:

But we do utilize a variety of security protocols at every layer of the architecture, as well as a robust encryption strategy. The NSA cloud brings together multiple data sets and protects each piece of data through security and enforcement of the authorities that specify its use. We do this by marking each individual piece of data with a set of tags that dictate its security protections and usage. In addition to data markings, security is applied throughout the architecture at multiple layers to protect data, systems, and usage. (Smith, 2014)

This thesis has chosen the term *transformational migration* to represent the requisite coordinated collection of adaptations required for a successful and transformational cloud migration. Transformational migration focuses on the relocation of data, security perimeter, knowledge base, and work processes to align with how the cloud actually functions.

A. COLLOCATING DATA LOCATION

The NSA CIO interview discusses the assembling of data (Smith, 2014). On-premises systems have typified data stranded in an application that prevents or significantly impedes access from other applications. The Institute for Defense Analyses suggests that:

Legacy applications, especially those designed prior to 2005, were likely designed to run on a single server, not on multiple shared and redundant cloud servers, and they may need to be rewritten to account for scenarios such as one of the host servers going offline. (Odell et al., 2015, p. 4-2)

Transformational migration mandates the collocation of relevant data sets or accessibility through secure API calls. Data migration is necessary to achieve machine learning capabilities named in the Cloud Executive Steering Group’s phase two goals, previously referenced in Chapter I.B.1, to quickly integrate cloud security and machine learning (Shanahan, 2017).

B. RECOMMENDATIONS ADDRESSING CUSTOMER MISCONFIGURATIONS

A better understanding of how the service model relates to the intent of the application can reduce the risk of customer misconfigurations. This section summarizes five migration strategies—rehost, refactor, revise, rebuild, and replace—and indicates the service model that maps to that strategy. Additionally, the section briefly introduces service-oriented architecture to increase portability and counter vendor lock-in.

Mapping Cloud Motivations to the Appropriate Service Model Migration

Rehost (lift and shift) applications on infrastructure as a service platforms are not going to benefit from the current and future capabilities of cloud technologies (Woods, 2011). Refactor (backward-compatible) application code with platform as a service will benefit from the current and future capabilities of cloud technologies (Woods, 2011). Revise (lift and refit) application code with infrastructure as a service or platform as a service can retrofit existing code for further rehosting or refactoring migration (Woods, 2011). The revised application with retrofitted code would then benefit from the current and future capabilities of cloud technologies. Rebuild (cloud native) application code with platform as a service abandons legacy code so that the rebuilt application will benefit from the current and future capabilities of cloud technologies (Woods, 2011). Replace (discard) applications with software as a service also abandon legacy code. The commercial software will deliver the benefit from current and future capabilities of cloud technologies. This solution does not allow customer customization and can lead to vendor lock-in by making the porting of data more challenging. As general cloud migration guidance, the Institute for Defense Analyses recommends first incorporating cloud-borne application for non-

mission essential data and then determining the necessary controls for mission-essential applications approved for Impact Level 5 (Odell et al., 2015).

This thesis only mentions service-oriented architecture because it is an abstraction that helps counter cloud service provider vendor lock-in. The Linux Information Project defines vendor lock-in as “the situation in which customers are dependent on a single manufacturer or supplier for some product (i.e., a good or service), or products, and cannot move to another vendor without substantial costs and/or inconvenience” (LINFO, 2006). Service-oriented architecture development produces applications that are treated like “services,” as in “anything as a service.” Once the application can be treated as a service, it should be able to port or “plug” into any cloud service provider seamlessly and temper the fears of having to make large-scale changes to existing code bases for interoperability with proprietary requirements of the new cloud service provider. Service-oriented architecture is easily reconfigurable.

C. RECOMMENDATIONS ADDRESSING CLOUD LEAKS

The DoD urgently needs to address cyberattacks not only on its own data centers but throughout its entire supply chain. This thesis, in Chapter 1.B.2, cited two DoD contractor-related breaches. *The Washington Post* reported another successful, high-profile cyberattack in the summer of 2018:

Chinese government hackers have compromised the computers of a Navy contractor, stealing massive amounts of highly sensitive data related to undersea warfare—including secret plans to develop a supersonic anti-ship missile for use on U.S. submarines by 2020, according to American officials. (Nakashima & Sonne, 2018a)

“Deliver Uncompromised” is a new strategy proposed by MITRE Corporation, which conducts federally funded research on behalf of the government, to address cybersecurity lapses that extend to DoD contractors (Nakashima & Sonne, 2018b). Deliver Uncompromised encourages adding security assessment attainment levels in the awarding of contracts along with traditional cost and performance considerations. The new strategy believes the cloud can contribute to protecting the DoD supply chain by specifically encouraging its contractors “to shift information systems and applications to qualified,

secure cloud service providers. The security outcome for many companies using the cloud will be superior compared to measures taken for on-premises systems” (Nissen, Gronager, Metzger, O’Donnell, & Rishikof, 2018). This thesis is supportive of such a move with the following caveats. The contractors 1) allow transformational migration principles to guide their organization’s transition to the cloud, 2) are aware of the misconceptions that lead to misconfigurations and subsequent cloud leaks, and 3) are aware of—and have a plan to address—the expected complications that arise in the implementation of cloud security controls, including digital forensic incident-response.

1. Migrating Security Perimeter

The NSA CIO’s interview also, as previously discussed, also mentions protecting data at a cellular level (Smith, 2014). On-premises systems have focused almost exclusively on the network security perimeter or internet-facing demilitarized zone. This has proven ineffective and there is growing support within information security to assume that organizational network boundaries are frequently breached. Transformational migration mandates that migration of the security perimeter comport with a more accurate assessment of adversarial threats. Transformational migration also supports extending the perimeter from the network boundary to include the boundary of specific chunks of data. Migrating the perimeter enables the user to leverage metadata tagging to unleash stricter enforcement of file authorizations and legal compliance. The NSA CIO affirmed his personal assessment that current cloud capabilities would have likely prevented U.S. soldier Bradley Manning from nefariously accessing classified information in 2010 (Smith, 2014).

2. Migrating Workflow Processes

In his interview, the NSA CIO also mentioned that security is applied throughout, at multiple layers of the cloud stack (Smith, 2014). On-premises systems have focused on encryption for data at rest and in transit and have been marred by repeated successful attacks because of incomplete security throughout the data security life cycle. Transformational migration mandates security through the complete data life cycle: creating, storing, processing, sharing, archiving, and destruction (CSA, 2017). This is

achieved by retraining a workforce that is no longer assigned to a subset of applications (Bommadevara et al., 2018). Instead, security is each member's responsibility. Transformational migration "requires mission owners to have a cloud-literate IT workforce to design, install, and configure the applications" (Odell et al., 2015, p. 4-4). Additionally, the Institute of Defense Analysis recommends that the DoD focus its hiring on professionals who have previous development experience in cloud environments (Odell et al., 2015). Cloud applications are often poorly documented. It is therefore incumbent on the consumer to have staff that can interface with the cloud service provider with similar work function responsibilities. Reorganizing will require migrating application professionals to a new dynamic, transformational workforce with the dexterity to remediate issues at multiple cloud logical layers. Both the newly hired and the retrained workforce will develop and continuously tune applications, addressing security at multiple layers and for the complete data life cycle.

3. Retraining Workforces for Cloud Computing

The DoD must train its cloud operators in greater depth, more frequently, and over a longer period of time than mandated for other information systems. The Institute for Defense Analyses reported encouraging initiatives in this vein, such as the Defense Acquisition University's plan to create a cloud module to educate technical staff about securely implementing within cloud environments (Odell et al., 2015). Transformational migration requires "an IT workforce who is familiar with and excited about cloud computing, and who can translate the benefits to DoD mission owners" (Odell et al., 2015, p. 6-1). However, this training fails to specifically address the underdeveloped training for forensic digital incident-response within cloud environments, which earnestly requires cloud service provider cooperation. Transformational migration for retraining places a significant burden on the cloud service provider to first foster greater collaboration to help forensic investigators understand the cloud service provider operating environment.

D. RECOMMENDATIONS ADDRESSING COMPLICATIONS IN THE IMPLEMENTATION OF SECURITY CONTROLS

The DoD CIO and DISA have acknowledged:

In addition to specified controls, commercial cloud service providers would benefit from the implementation of DoD-specific security services that all cloud service providers could leverage. As of July 2015, this is only a proposed development, but its implementation would help spur cloud adoption and ease integration of cloud service providers into the DoD security framework. (Odell et al., 2015, p. 5-2)

Additionally, the cloud controls that are more difficult to implement than traditional architectures pose a further challenge. One potential solution is to use a forensics as a service (FaaS) provider to address forensic aspects that are more challenging to implement, but this increases supply chain management, with yet an additional vendor to manage.

E. RECOMMENDATIONS ADDRESSING DIGITAL FORENSIC INCIDENT-RESPONSE

This thesis agreed with three recommendations that the NIST Cloud Computing Forensic Science Working Group (2014) advanced to improve the forensic digital incident-response process. The first is that consumer boundaries within cloud service provider architectures require better definition. This recommendation aims to segregate forensic data in a multi-tenant architecture to remove the possibility of breaching other tenants' confidentiality during a forensic investigation. The second recommendation, related to the first, advocates for the integration of forensic evidence collection tools into the cloud service provider architecture (NCC FSWG, 2014). The third—implementing structural changes to consumer boundaries that guarantee consumer confidentiality during forensic investigations and integrated cloud service provider forensic acquisition tools—will remediate these two high-priority digital forensic science obstacles. This thesis also advises that cloud service providers welcome greater collaboration with academic and forensic science researchers to improve the field's understanding of the challenge-laden cloud service provider environment.

The open-source Linux Caine distribution includes a digital forensic tools suite that may address several of the challenges raised by the NCC FSWG: Autopsy (recovering lost

or obfuscated data, artifact extraction, and time line analysis), Guymager (disk imaging without damaging the cloud service provider's image), Fred (scanning folders and files for deleted or obfuscated data), and Photorec (data carving confirming complete recover of damaged or deleted files) (Decusatis, Carranza, Ngaide, Zafar, & Landaez, 2015). Integrating forensics as a service (FaaS) applications embedded in cloud service provider environments to readily manage the investigation of cloud log files would increase investigation utility without forcing cloud service providers to reveal more about their environments than they are prepared to (Raju, & Moharil, & Geethakumari, 2016). Digital forensics as a service (DFaaS) uses customized, multilateral service-level agreements, beyond the consumer and cloud service provider, but standardizes technical, legal, and boundary forensics aspects identified by the NCC FSWG (Keserwani & Samaddar, 2017).

This thesis advocates for the inclusion of specific terms within contractual service-level agreements. Service-level agreements can contain service-level objectives. Service-level objectives specifically indicate measures to determine performance or underperformance. This thesis endorses service-level objectives that include cloud service provider obligations for preservation and access to evidential sources under cloud service provider control (CSA, 2013). Many of those sources will be logs and should be specifically named by function (e.g., guest operating system logs; CSA, 2013). This reiterates that all critical contractual terms require documentation prior to a digital forensic incident.

THIS PAGE INTENTIONALLY LEFT BLANK

VI. CONCLUSION

This thesis set out to identify the detailed deviations between digital forensic incident-response in on-premises and cloud computing architectures and offer risk management recommendations for migration. This thesis answered the research question by identifying, in detail, the divergence of digital forensic incident-response procedures in cloud computing starting in Chapter III.C. The thesis also diagnosed four cloud computing–related risk management issues that the DoD must tackle: customer misconfigurations, cloud leaks, complications in implementing security controls, and digital forensic incident-response challenges (Chapter I.B.1-4). Furthermore, the thesis identified that when humans (customers) misunderstand the ramifications of service model selection and the related complications with implementing security controls, customer misconfigurations occur, which can lead to cloud leaks and, if serious enough, result in a digital forensic incident-response investigation. Chapter V’s recommendations advocate for adopting a transformational migration mindset by consolidating the series of organizational steps necessary to achieve a successful cloud migration. For more in-depth treatment of definitions, see Chapter II; for cloud-specific threats, see Chapter II.E; for cloud security controls, see Chapter II.F; for service model selection, see Chapter III (para. 1); for critical cloud structures for securing data, application, infrastructure, and configuration settings; see Chapter III.B.2; and for federal cloud computing policy for nine incident-response controls, see Chapter IV.B.5. There is obviously more research required into digital forensic incident-response.

Recommendations for follow-on or tangential research are:

- The publication of Security Technical Implementation Guides (STIGs) for all cloud services (infrastructure as a service, platform as a service, and software as a service). It could begin with the three largest cloud service providers and would result in nine STIGs to enforce strict configuration settings to reduce misconfigurations.

- The development of a series of required robust continuous cloud training modules that retrain the DoD workforce to securely interface with cloud service providers.
- Experimentation with the open-source Linux Caine distribution digital forensic tools suite to determine full capabilities and limitations in addressing forensic investigative challenges.
- The vetting of the newly proposed Deliver Uncompromised strategy to shift the DoD supply chain to approved cloud service providers.
- The development and implementation of DoD-specific security services that all cloud service providers could leverage.
- The testing of automated cloud penetration tools for each of the DoD cloud service providers akin to CloudSploit, which is dedicated to scanning AWS.
- The integrating of embedded forensics as a service applications in cloud service provider environments to address digital forensic investigative challenges.
- Experimenting with forensics as a service applications in external cloud service provider environments to determine efficacy in reducing digital forensic investigative challenges.
- The customization of multilateral digital forensic as a service service-level agreements that standardize technical, legal, and boundary forensic aspects identified by the working group.
- The endorsement of a service-level agreement, representing majority acceptance by the largest cloud service providers by market share, that includes cloud service provider obligations for preservation and access to evidential sources under cloud service provider control.

LIST OF REFERENCES

- Amazon Web Services. (2015). How do I create and activate a new Amazon Web Services account? Retrieved from <https://aws.amazon.com/premiumsupport/knowledge-center/create-and-activate-aws-account/>.
- Apple. (n.d.). Logging. Retrieved September 18, 2018, from <https://developer.apple.com/documentation/os/logging>
- Bommadevara, N., Del Miglio, & A., Jansen, S. (2018). Cloud adoption to accelerate IT modernization. Retrieved from <https://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/cloud-adoption-to-accelerate-it-modernization>
- Central Intelligence Agency (CIA). (2014, Dec. 17). CIA creates a cloud: An interview with CIA's chief information officer, Doug Wolfe, on cloud computing at the agency. Retrieved from <https://www.cia.gov/news-information/featured-story-archive/2014-featured-story-archive/cia-creates-a-cloud.html>
- Chief Information Officer Council & Chief Acquisition Officers Council (CIO & CAO Councils). (2012). *Creating effective cloud computing contracts for the federal government: Best practices for acquiring IT as a service*. Retrieved from <https://www.cio.gov/2012/02/24/cloud-computing-update-best-practices-for-acquiring-it-as-a-service/>
- Clarke, G. (2015, Apr. 13). Self preservation is AWS security's biggest worry, says gros fromage. *The Register*. Retrieved from https://www.theregister.co.uk/2015/04/13/aws_security_sleepless_nights/
- Cloud Security Alliance (CSA). (2011). *Quick guide to the reference architecture*. Retrieved from https://cloudsecurityalliance.org/wp-content/uploads/2011/10/TCI_Whitepaper.pdf
- Cloud Security Alliance (CSA). (2013). *Mapping the forensic standard ISO/IEC 27037 to cloud computing*. Retrieved from <https://downloads.cloudsecurityalliance.org/initiatives/imf/Mapping-the-Forensic-Standard-ISO-IEC-27037-to-Cloud-Computing.pdf>
- Cloud Security Alliance (CSA). (2016). *The treacherous 12: Cloud computing top threats in 2016*. Retrieved from https://downloads.cloudsecurityalliance.org/assets/research/top-threats/Treacherous-12_Cloud-Computing_Top-Threats.pdf
- Cloud Security Alliance (CSA). (2017). *Security guidance: For critical areas of focus in cloud computing v4.0*. Retrieved from https://cloudsecurityalliance.org/guidance/#_overview

- Cloud Security Knowledge Sharing. (2018). Cloud security responsibility matrix (low differentiation) [Image]. Retrieved from <https://cloudsecurityknowledgesharing.com/dealing-with-shared-responsibility-model-in-public-cloud/>
- Committee on National Security Systems (CNSS). (2014). *Security categorization and control selection for national security systems*, CNSSI No. 1253. Retrieved from http://www.dss.mil/documents/CNSSI_No1253.pdf
- Committee on National Security Systems (CNSS). (2015). *National information assurance (IA) glossary*, CNSSI No. 4009 Retrieved from <https://www.cnss.gov/CNSS/issuances/Instructions.cfm>
- Decusatis, C., Carranza, A., Ngaide, A., Zafar, S., & Landaez, N. (2015). Methodology for an open digital forensics model based on CAINE. *2015 IEEE International Conference on Computer and Information Technology; Ubiquitous Computing and Communications; Dependable, Autonomic and Secure Computing; Pervasive Intelligence and Computing (CIT/IUCC/DASC/PICOM)* (pp. 935–940). doi:10.1109/CIT/IUCC/DASC/PICOM.2015.61
- Defense Information Systems Agency (DISA). (2017, Mar. 6). *Department of Defense Cloud Computing Security Requirements Guide*, version 1, release 3. Retrieved from https://www.complianceweek.com/sites/default/files/department_of_defense_cloud_computing_security_requirements_guide.pdf
- Department of Homeland Security (DHS). (2013). Spillage and cloud computing. Retrieved from https://csrc.nist.gov/CSRC/media/Events/ISPAB-FEBRUARY-2013-MEETING/documents/ispab_feb2013_cloud-security-challenges_rseeholzer.pdf
- European Network and Information Security Agency (ENISA). (2009). *Cloud computing: Benefits, risks and recommendations for information security*. Retrieved from <https://www.enisa.europa.eu/publications/cloud-computing-risk-assessment>
- Gordon, A. (2016). *The official (ISC)2 guide to the CCSP CBK* (2nd ed.). Indianapolis, IN: John Wiley and Sons.
- Gregg, A. (2017, Jun. 1). Booz Allen Hamilton employee left sensitive passwords unprotected online. *Washington Post*. Retrieved from https://www.washingtonpost.com/business/capitalbusiness/government-contractor-left-sensitive-passwords-unprotected-online/2017/06/01/916777c6-46f8-11e7-bcde-624ad94170ab_story.html?utm_term=.6cad14ff8b95
- Grobauer, B., Walloschek, T., & Stöcker, E. (2011). Understanding cloud computing vulnerabilities. *IEEE Security & Privacy*, 9(2), 50–57. doi:10.1109/MSP.2010.115

- ISACA. (2009). *Cloud computing: business benefits with security, governance and assurance perspectives*. Retrieved from <http://www.isaca.org/Knowledge-Center/Research/ResearchDeliverables/Pages/Cloud-Computing-Business-Benefits-With-Security-Governance-and-Assurance-Perspective.aspx>
- Joint Chiefs of Staff. (2012). *Cyber incident handling program*, CJCSM 6510.01B. Retrieved from <http://www.jcs.mil/Portals/36/Documents/Library/Manuals/m651001.pdf?ver=2016-02-05-175710-897>
- Keserwani, P., & Samaddar, S. (2017). Customization of Service level agreement for digital forensics as a service. *Proceedings of the 7th International Conference on Computer and Communication Technology* (pp. 139–150). doi:10.1145/3154979.3154993
- Kostoska, M., Gusev, M., & Ristov, S. (2016). An overview of cloud interoperability. *Proceedings of the Federated Conference on Computer Science and Information Systems*, 8, 873–876. doi:10.15439/2016F463
- Larson, S. (2017, Nov. 17). Pentagon exposed some of its data on Amazon server. Retrieved from <https://money.cnn.com/2017/11/17/technology/centcom-data-exposed/>
- LINFO. (2006, Apr. 29). Re: Vendor lock-in definition. Retrieved from www.lininfo.org/vendor_lockin.html.
- Liu, F., Tong, J., Mao, J., Bohn, R., Messina, J., Badger, L., & Leaf, D. (2011). *NIST cloud computing reference architecture*, Special Publication 500-292. Retrieved from <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication500-292.pdf>
- Mell, P., & Grance, T. (2011). *The NIST definition of cloud computing*, Special Publication 800-145. Retrieved from <https://csrc.nist.gov/publications/detail/sp/800-145/final>
- Metheny, M. (2013). *Federal cloud computing*. Retrieved from <https://www.safaribooksonline.com>
- Nakashima, E., & Sonne, P. (2018a, Jun. 8). China hacked a Navy contractor and secured a trove of highly sensitive data on submarine warfare. *Washington Post*. Retrieved from https://www.washingtonpost.com/world/national-security/china-hacked-a-navy-contractor-and-secured-a-trove-of-highly-sensitive-data-on-submarine-warfare/2018/06/08/6cc396fa-68e6-11e8-bea7-c8eb28bc52b1_story.html?utm_term=.9d17b563bcb5

- Nakashima, E., Sonne, P. (2018b, Aug. 13). Pentagon is rethinking its multibillion-dollar relationship with U.S. defense contractors to boost supply chain security. *Washington Post*. Retrieved from https://www.washingtonpost.com/world/national-security/the-pentagon-is-rethinking-its-multibillion-dollar-relationship-with-us-defense-contractors-to-stress-supply-chain-security/2018/08/12/31d63a06-9a79-11e8-b60b-1c897f17e185_story.html?utm_term=.60664aebdfb8
- National Institute of Standards and Technology (NIST). (2004). *Standards for security categorization of federal information and information systems*, FIPS 199. Retrieved from <https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.199.pdf>
- National Institute of Standards and Technology (NIST). (2006). *Minimum security requirements for federal information and information systems*, FIPS 200. Retrieved from <https://csrc.nist.gov/publications/detail/fips/200/final>
- National Institute of Standards and Technology (NIST). (2010). *Guide for applying the risk management framework to federal information systems*, NIST Special Publication 800-37, Revision 1. Retrieved from <https://csrc.nist.gov/publications/detail/sp/800-37/rev-1/final>
- National Institute of Standards and Technology (NIST). (2012). *Information security*, NIST Special Publication 800-30, Revision 1. Gaithersburg, MD: Author.
- National Institute of Standards and Technology (NIST). (n.d.). Using risk management to improve privacy in information systems. Retrieved September 18, 2018, from [https://csrc.nist.gov/CSRC/media/Presentations/Using-Risk-Management-to-Improve-Privacy-in-In-\(2\)/images-media/day3_research_1035-1125.pdf](https://csrc.nist.gov/CSRC/media/Presentations/Using-Risk-Management-to-Improve-Privacy-in-In-(2)/images-media/day3_research_1035-1125.pdf)
- Nissen, C., Gronager, J., Metzger, R., O'Donnell, R., & Rishikof, H. (2018). *Deliver uncompromised: a strategy for supply chain security and resilience in response to the changing character of war*. Retrieved from <https://www.mitre.org/publications/technical-papers/deliver-uncompromised-a-strategy-for-supply-chain-security>
- NIST Cloud Computing Forensic Science Working Group (NCC FSWG). (2014). *NIST cloud computing forensic science challenges*, Draft NISTIR 8006. Retrieved from <https://csrc.nist.gov/publications/detail/nistir/8006/draft>
- Odell, L., Wagner, R., & Weir, T. (2015). *Department of Defense use of commercial cloud computing capabilities and services*. Retrieved from <http://www.dtic.mil/dtic/tr/fulltext/u2/1002758.pdf>
- Raju, B., Moharil, B., & Geethakumari, G. (2016). FaaSaaS: Enabling forensics-as-a-service for cloud computing systems. *Proceedings of the 9th International Conference on Utility and Cloud Computing* (pp. 220–227). doi:10.1145/2996890.3009904

- Saltzer J. H., & Kaashoek, M. F. (2009) *Principles of computer system design: An introduction*. Retrieved from <http://ebookcentral.proquest.com/lib/ebook-nps/detail.action?docID=534968>
- SANS Institute. (2016). *Implementing the critical security controls in the cloud*. Retrieved from <https://www.sans.org/reading-room/whitepapers/critical/implementing-critical-security-controls-cloud-36725>
- Scarfone, K., Souppaya, M., & Hoffman, P. (2011). *Guide to security for full virtualization technologies*, Special Publication 800-125. Retrieved from <https://www.nist.gov/publications/guide-security-full-virtualization-technologies>.
- Secretary of the Air Force. (2017). *Cyber incident handling*, Air Force Instruction 17-203. Retrieved from http://static.e-publishing.af.mil/production/1/af_a3/publication/afi17-203/afi17-203.pdf
- Shanahan, P. (2017, Sep. 13). *Accelerating enterprise cloud adoption* [Memorandum]. Retrieved from <https://fcw.com/articles/2017/09/20/dod-cloud-adoption-risks.aspx>.
- Singh, S., & Singh, D. (2017). Cloud computing: Security issues and challenges. *International Journal of Advances in Engineering & Technology*, 10(3), 338–343. Retrieved from <http://search.proquest.com/docview/1930762746/>
- Smith, D. (2014, Sep. 29). Exclusive: Inside the NSA's private cloud. Retrieved from <https://www.networkworld.com/article/2687084/security0/exclusive-inside-the-nsa-s-private-cloud.html>
- VanRoekel, S. (2011, Dec. 8). *Security authorization of information systems in cloud computing environment* [Memorandum]. Retrieved from <https://www.fedramp.gov/resources/documents-2016/>
- Woods, J. (2011). Five options for migrating applications to the cloud: Rehost, refactor, revise, rebuild or replace. Retrieved from https://gartnerinfo.com/futureofit2011/MEX38L_A2%20mex38l_a2.pdf.

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