

ABSTRACT

Title of dissertation: OPTIMIZING PROACTIVE MEASURES
FOR SECURITY OPERATIONS
Rock A. Stevens
Doctor of Philosophy, 2020

Dissertation directed by: Professor Michelle L. Mazurek
Department of Computer Science

Digital security threats may impact governments, businesses, and consumers through intellectual property theft, loss of physical assets, economic damages, and loss of confidence. Significant effort has been placed on technology solutions that can mitigate threat exposure. Additionally, hundreds of years of literature have focused on non-digital, human-centric strategies that proactively allow organizations to assess threats and implement mitigation plans. For both human and technology-centric solutions, little to no prior research exists on the efficacy of how humans employ digital security defenses. Security professionals are armed with commonly adopted “best practices” but are generally unaware of the particular artifacts and conditions (e.g., organizational culture, procurement processes, employee training/education) that may or may not make a particular environment well-suited for employing the best practices. In this thesis, I study proactive measures for security operations and related human factors to identify generalizable optimizations that can be applied for measurable increases in security. Through interview and survey methods, I investigate the human and organizational factors that shape the adoption

and employment of defensive strategies. Case studies with partnered organizations and comprehensive evaluations of security programs reveal security gaps that many professionals were previously unaware of — as well as opportunities for changes in security behaviors to mitigate future risk. These studies highlight that, in exemplar environments, the adoption of proactive security assessments and training programs lead to measurable improvements in organizations' security posture.

OPTIMIZING PROACTIVE MEASURES
FOR SECURITY OPERATIONS

by

Rock A. Stevens

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2020

Advisory Committee:
Professor Michelle Mazurek, Chair/Advisor
Professor Jennifer Golbeck
Professor Tudor Dumitras
Professor John Dickerson
Professor Dave Levin

© Copyright by
Rock Stevens
2020

Acknowledgments

Thank you to my family for their love and support – I promise I will make up for all of the missed vacations! Moreover, none of this would have been possible without Michelle’s support, mentorship, and super-human abilities to turn my ideas into reality. A special thanks to all of my collaborators and friends from industry that enabled these unique case studies.

Table of Contents

Acknowledgements	ii
Table of Contents	iii
List of Tables	vii
List of Figures	viii
1 Introduction	1
2 Related Work	8
2.1 Human factors in threat modeling adoption	8
2.2 Compliance programs	13
2.3 Incident Response Playbooks	17
2.3.1 Crisis preparedness in other domains	18
2.3.2 Threat exposure and readiness	19
2.3.3 Incident response playbooks	20
2.4 Human factors in security operations	24
3 Planning for Security: Human Factors in Threat Modeling Adoption	26
3.1 The Center of Gravity framework	27
3.2 Methods: Threat modeling at NYC3	31
3.2.1 Recruitment	32
3.2.2 Study protocol	33
3.2.3 Limitations	39
3.3 Results	41
3.3.1 Participants	41
3.3.2 Pre-intervention baseline	43
3.3.3 Immediate observations	44
3.3.3.1 Perceived efficacy	44
3.3.3.2 Actual efficacy	48
3.3.4 Observations after 30 days	53
3.3.4.1 Perceived efficacy	53
3.3.4.2 Actual efficacy	54
3.3.4.3 Actual adoption	55
3.3.5 Observations after 120 days	56

	3.3.5.1	Actual adoption	57
	3.3.5.2	Actual efficacy	60
3.4		Discussion	62
	3.4.1	Lessons learned	63
	3.4.2	Future work	66
4		Baselining Security: Security Implications of Policies, Laws, and Regulations	67
	4.1	Method	68
		4.1.1 Compliance-standard audit	69
		4.1.2 Expert validation process	74
		4.1.3 Limitations	77
	4.2	Results: IRS P1075	78
		4.2.1 Overview	78
		4.2.2 Findings	80
		4.2.3 Expert validation	86
	4.3	Results: PCI DSS	89
		4.3.1 Overview	89
		4.3.2 Findings	91
		4.3.3 Expert validation	96
	4.4	Results: NERC CIP 007-6	97
		4.4.1 Overview	97
		4.4.2 Findings	99
		4.4.3 Expert validation	103
	4.5	Disclosures	106
		4.5.1 Disclosure intent	107
		4.5.2 IRS P1075	108
		4.5.3 PCI DSS	109
		4.5.4 NERC CIP 007-6	109
		4.5.5 Federal-level recognition	110
	4.6	Discussion	111
5		Baselining Security: Security Implications of Policies, Laws, and Regulations in the Cloud	116
	5.1	Background	117
	5.2	FedRAMP evaluation results	119
		5.2.1 Ambiguous specifications	120
		5.2.2 Obsolete references	125
		5.2.3 Risks to data	125
	5.3	Unaccounted for threat models	128
		5.3.1 Nation-state privileged access	129
		5.3.2 Corporate aggregation and monetization	130
		5.3.3 Security of security appliances	132
		5.3.4 Ignored cyber-physical systems	133
	5.4	Comparing FedRAMP	134
	5.5	Discussion	135

6	Implementing Security: Humans factors in Incident Response Readiness	137
6.1	Setup and preliminaries	138
6.1.1	Selected frameworks	138
6.1.2	The partners	140
6.1.3	Selected scenarios	141
6.1.3.1	Brute- force login attempts	142
6.1.3.2	Valid credential misuse	143
6.2	Playbook design and evaluation	144
6.2.1	Method	144
6.2.1.1	Recruitment	145
6.2.1.2	Playbook design	145
6.2.1.3	Playbook evaluation	147
6.2.2	Participants	147
6.2.2.1	Limitations	148
6.2.3	Results	150
6.2.3.1	Participants	151
6.2.3.2	Usability metrics	151
6.2.3.3	IACD feedback	152
6.2.3.4	NIST feedback	155
6.2.3.5	Expert evaluation	158
6.2.4	Summary	162
6.3	Playbook implementation and use	163
6.3.1	Method	163
6.3.1.1	Playbook implementation	164
6.3.1.2	Playbook use during incident response	165
6.3.2	Results	168
6.3.2.1	Recruitment	168
6.3.2.2	Playbook implementation	169
6.3.2.3	Playbooks in use	171
6.3.3	Summary	177
6.4	Playbooks in other domains	177
6.5	Discussion	178
7	Implementing Security: Complementing and Repairing Baseline Security	183
7.1	Method	185
7.1.1	Survey design	185
7.1.2	Recruitment and Screening	187
7.1.3	Data analysis	188
7.1.4	Limitations	190
7.2	Results	191
7.2.1	Participants	192
7.2.2	Compliance is insufficient	194
7.2.3	Going beyond compliance	199
7.2.4	Additional measures are not a panacea	208
7.3	Discussion	215

8	Discussion	220
8.1	Better together	220
8.2	Impacts of organizational culture	223
8.3	Proposed workflows for proactive measures	228
9	Conclusion	231
A	Study Instruments	233
A.1	Survey Questions from Chapter 3	233
A.1.1	Pre-intervention survey	233
A.1.2	Post-intervention survey	236
A.1.3	Follow-up survey	240
A.1.4	NYC leadership panel questions	243
A.2	Expert Survey from Chapter 4	243
A.3	Survey instruments for Chapter 6	245
A.3.1	Design phase	245
A.3.2	Evaluation phase	247
A.3.3	Implementation phase	248
A.3.4	Utilization phase	250
A.4	Interview guide	252
A.5	Survey instruments for Chapter 7	253
B	Additional Data	256
B.1	Additional data from Chapter 3	256
B.1.1	Star Wars walkthrough	256
B.1.2	E-commerce scenario	257
B.1.3	Participant P17 example	258
B.1.4	Visualizing Center of Gravity	260
B.1.5	CoG Identification Accuracy Regression	262
B.2	Additional data from Chapter 4	263
B.2.1	Overall risk distribution from Chapter 4	263
B.2.2	Compliance audit findings from Chapter 4	263
B.3	Additional data from Chapter 6	284
B.4	Playbook examples	284
B.5	Additional data from Chapter 7	284
B.5.1	Quantitative analysis	284
B.5.2	Reported measures	286
B.5.3	List of reported compliance standards	286
B.5.4	Demographics	287
B.5.5	Codebook	287
	Bibliography	295

List of Tables

Chapter 3 study participants	42
Chapter 4 expert participants	69
Chapter 4 example inter-rater reliability matrix	72
Chapter 4 distribution of security concerns	111
Chapter 6 evaluation coverage of designed playbooks	147
Chapter 6 participant and expert demographics	149
Chapter 6 demographics summary	193
Chapter 7 regression model	208
Chapter 7 analysis of sentiment and assessment frequency	208
Chapter 3 summary of threat modeling accuracy	262
Chapter 4 security concerns identified in IRS P1075	263
Chapter 4 security concerns identified in PCI DSS	271
Chapter 4 security concerns identified in NERC CIP 007-6	274
Chapter 4 security concerns identified in FedRAMP	275
Chapter 7 factors for Cumulative Link Mixed Model	284
Chapter 7 contrasts and estimates	286
Chapter 7 codebook	287
Reported complementary measures in Chapter 7	292
Reported compliance standards in Chapter 7	293
Chapter 7 study participants	294

List of Figures

Threat modeling process using CoG	28
The six-part threat modeling study protocol and metrics.	31
Threat modeling subtask completion times	49
Threat modeling efficacy after 30 days	50
Threat modeling organizational change through task tracking	56
Example of security concern annotations	70
Composite Risk Management Framework for assessing risk	74
Distribution of security concerns identified for IRS P1075	79
Distribution of security concerns identified for PCI DSS	91
Distribution of security concerns identified for NERC CIP 007-6	100
Distribution of FedRAMP security concerns	121
Playbook comparisons of perceptions and error rates	152
Reported distribution of compliance standards	194
Visualization of complementary measure categories	195
Visualization of sentiment distribution	207
Recommended proactive measures workflow	228
Threat modeling using the CoG tabular method	260
Participant example of CoG visualization	261
Comparing security concerns in IRS, PCI, and NERC	263
IACD playbook example	284
NIST playbook example	285
Visualization explaining playbook threat detection failures	285

Chapter 1: Introduction

The digital-security threat landscape necessitates that organizations employ defensive strategies to mitigate risk exposure. Experts estimate that intellectual property and sensitive data theft due to digital espionage range from \$1 to \$2 trillion annually and can impart irreparable damage to affected businesses [13, 129, 183].

For these reasons, numerous complex security paradigms exist that offer near-term assistance for shoring up an organization’s risk mitigation efforts. A subset of solutions includes investments in “next-generation” defense technologies, hiring additional security staff, or outsourcing specialized labor to identify and mitigate threats [49]. While new technology platforms might provide an added benefit, organizations also must optimize the use of existing resources.

Proactive measures for security operations involves three factors: (1) planning to mitigate likely threats, (2) establishing a baseline of security for sustained operations, and (3) implementing security controls that can prevent adversarial access or reduce the impact of an intrusion.

Proactive planning involves assessing critical assets, building plans to protect those assets, and investing in the cognitive skills of security technicians to reduce digital-security risk [66]. Threat modeling frameworks present a methodology for

evaluating complex threats and breaking them into modular components to allow humans to reason about the threat, understand risk, and develop mitigating strategies [184].

Proactive baselining ensures digital systems are operating at a minimum level of security to defeat common threats and ensures administrators strictly adhere to practices that sustain security. In many cases, proactive baselines are mandated by various laws, policies, or regulations through compliance programs [75, 108, 159, 169].

The third and final component of proactive security involves implementing various security controls that mitigate adversarial impact. Incident response playbooks present security practitioners with a strategy for handling an imminent threat and reducing adverse organizational impacts by providing practitioners with pre-planned actions [150]. These pre-planned actions are intended to help reduce stresses that technicians may face during a security incident, allow them to gain momentum during response efforts, and ensure organizations are prepared for likely adversarial situations. Additionally, organizations can implement security measures that complement mandated baselines to either address security gaps or repair security complications caused by compliance programs.

Proactive measures for security operations are standard practices within the security industry, but little to no research portrays the efficacy of these defensive strategies in practice.

This situation suggests that opportunities exist to improve upon how organizations design and implement proactive measures. To this end, this dissertation empirically studies how security professionals and organizations proactively employ

defensive strategies in practice and identify gaps that inhibit their effectiveness. The resulting findings identify methods to measurably improve proactive security measures in real-world environments.

This thesis measures the efficacy of defensive strategies in practice to improve proactive measures for digital security operations. Specifically, **the study of proactive measures for security operations and related human factors will lead to optimizations that can be applied for measurable increases in security.**

My colleagues and I investigate this hypothesis by studying each component of proactive measures: planning, baselining, and implementing. Understanding these components individually will allow us to analyze commonalities, understand particular nuances that may apply to certain situations, and understand adaptations that were required for anecdotal “best practices” to be translated into practice for security operations.

Comprehensive evaluations reveal that, oftentimes, the efficacy of defensive security controls relies on numerous complementary organizational and human-centric factors. Chapters 3 and 6 highlight the need for educational interventions, the adoption of reinforcing training programs, and changes in communication channels for security behaviors to take hold. Chapters 3 and 4 show that even foundational security practices can result in security concerns if not correctly integrated into existing security paradigms. Chapter 7 focuses on the steps organizations take to implement security solutions that address threats not mitigated by baseline security paradigms.

This thesis evaluates the three components of proactive measures and identifies methods for optimizing through a series of research studies. Using a case study in

Chapter 3, we conduct an end-to-end analysis on the efficacy of an exemplar threat-modeling planning framework within an enterprise environment. We ask security professionals proactively to consider threats to the networks and systems, allow them to design security mechanisms that can mitigate risk, implement those controls within a production network, and measure the change in security posture over the span of 120 days. Additionally, we attempt to understand the particular aspects of threat modeling that are well-suited for the partnered organization and isolate the organizational factors that assisted in reinforcing adoption.

The results suggest that threat modeling may provide valuable benefits in an enterprise setting. After 30 days of using threat modeling, 20 employees reported that they had adopted its concepts in their daily routine and 23 employees believed threat modeling provided tangible benefits to their security jobs. These positive perceptions shaped the adoption metrics observed over the next 90 days. NYC3 employees implemented participant-designed defensive plans to mitigate threats across eight newly-identified threat categories.

As a result, these defensive strategies prevented five privileged account hijackings, mitigated 541 unique intrusion attempts, and remedied three previously unknown web-server vulnerabilities. These positive outcomes were supported and enabled by integrating threat modeling into existing processes; developing new training and mentorship programs; and improving how security professionals communicate with one another and executive-level leaders.

The study with NYC3 in Chapter 3 revealed that, prior to using threat modeling, compliance programs played a foundational role in how the organization and

its security technicians implement baseline security controls. Chapters 4 and 5 focus on potential security concerns that may exist despite perfect adherence to compliance standards. These studies feature a novel systematic methodology for evaluating the risks that organizations may inherit as a result of baseline security and compliance with various policies, laws, and regulations. In Chapter 4, a group of researchers evaluate programs that affect everyday citizens, ask industry experts to validate/reject their findings, and partner with relevant organizations to disclose their findings. This study reveals that when compliance standards are used literally as checklists, perfect compliance can result in sub-optimal security conditions. This study highlights that hundreds of issues with varying severity exist across credit card, taxpayer, and electric grid security standards. More problematic, no clearly-defined process exists for reporting security concerns associated with compliance standards. Overall, results suggest that auditing compliance standards can provide valuable benefits to the security posture of compliant organizations.

Chapter 5 extends lessons learned from Chapter 4 and identifies 46 issues that may present security threats to organizations that use FedRAMP-approved programs in cloud-based infrastructures. Additionally, thematic analysis reveals four threat models that appear to be neglected throughout the FedRAMP framework and could pose significant threats if not properly handled.

Chapter 3 revealed that while technicians were aware of security issues, their operations center did not always have a response plan codified for when the issue manifested. Chapter 6 details two case studies designed to evaluate an end-to-end implementation of incident response playbook frameworks within a security opera-

tions center. Security professionals used two exemplar incident-response playbook-design frameworks to design their own response plans for specific threats their organization may encounter. Industry experts, in turn, assessed the validity of the designed playbooks. Using the best-scoring playbooks, technicians implemented select playbooks within a partnered security operations center and conducted a series of incident response exercises to measure users' ability to perform response actions while following the prescribed response plans.

Findings suggest that playbooks, in some cases, simplify and support incident response efforts. However, designed playbooks often lacked sufficient detail for real-world use, particularly for more junior technicians. This study shows that incident response playbooks may be valuable tools for increasing preparedness against a threat, but often require extensive planning and threat modeling to determine which playbooks should be prioritized for development and customization. Additionally, baseline security mechanisms, such as the ones discussed in Chapter 4, may constrain or alter playbook design efforts to ensure compatibility with applicable compliance programs.

From Chapters 4 and 5, we understand that digital security compliance programs help organizations establish baseline security levels, but also are laden with their own security issues. As a result, security professionals are left to assess the impact of compliance on their organization and to identify ways to extend security beyond compliance mandates to fill known security gaps. Chapter 7 reports on organizations' use of *complementary measures* — policies and technical controls enacted to mend known security gaps and exceed compliance requirements. Af-

ter surveying 40 security professionals from across multiple essential-service sectors, thematic and content analysis reveal (1) numerous complementary measures that organizations use to address security gaps, (2) identification of the measures that worked particularly well (or poorly), (3) how organizations prioritize and evaluate the complementary measures they adopt.

Findings suggest that compliance programs are insufficient and that 37 of 40 organizations implement complementary measures as a result. Although the specifics of how and why organizations implement complementary measures vary, we find that organizations often adopt complementary measures in response to security incidents, to reduce costs, when recommended by external experts, or requested by (sometimes non-technical) executives. Participants found complementary measures to be beneficial, but far from perfect. Participants reported numerous instances of poorly-managed complementary processes that impact organizational security. Overall, improving compliance programs and organizational culture that supports digital security may provide valuable insight into improving security for organizations as a whole.

Chapter 2: Related Work

Prior work has focused on various methods for improving digital security postures. In this section, I review related work that focuses on threat modeling to anticipate future risk; various legal and policy frameworks that impact security; and incident response playbooks that assist with readiness for when a security breach occurs.

2.1 Human factors in threat modeling adoption

The following is a review of prior work on digital-security threat-modeling techniques and empirical studies of these models. Current research in digital security threat-modeling frameworks focuses on providing structured processes for dynamic situations and insight into possible threats to systems.

Threat-modeling frameworks are often ported into digital security from the military domain. The digital-security adaptation of the military *OODA Loop* framework focuses on improving situational awareness and serves as a decision-making model that allows defenders to rapidly observe the situation, orient themselves to the problem, decide on a course of action, and act effectively [86]. In this context, a failure to understand threats will ultimately lead to an incorrect decision for

defenders and responders — but success will allow defenders to preempt attackers and improve their likelihood of success. Relatedly, the *Cynefin framework* aims to manage cybersecurity risk and provide a process for structuring uncertainty to reveal context, develop situational awareness, and illuminate appropriate responses to complex problems [61]. The five Cynefin domains help describe problems; framing problems as terms of disorder, obvious, complicated, complex, and chaotic provide defenders with a guide for action and assist with risk management. Both of these frameworks enable defenders to reactively implement defenses against active threats while understanding their tactics, techniques, and procedures but are not structured for planning proactive defense plans. Proactive plans serve to mitigate threats before an adversary is able to enact an attack.

The *cyber kill chain* assists defenders in recognizing the essential tasks that an adversary must complete before accomplishing their goals [103, 147]. From this perspective, defenders have an advantage over adversaries; attackers must complete all stages of their kill chain for success, whereas defenders only need to thwart an adversary once at any stage. Much like OODA loop and the Cynefin frameworks, these frameworks enable defenders to reactively implement defenses against active threats while understanding their tactics, techniques, and procedures but they are not structured for planning proactive defense plans. Additionally, they heavily rely upon threat intelligence and situational awareness for response actions. However, one significant deficiency of this framework is its exclusive focus on external threats: it considers neither insider nor inadvertent threats.

The National Security Agency produced a threat modeling framework for

adversarial mitigation techniques using the output of an internal wargame exercise [154]. This framework provides an enumeration of common threat attack vectors and proposes defenses based on observed offensive tactics from the wargame. While these guidelines represent the agency’s best practices, they do not account for tailoring threats to individual organizations.

A multitude of other threat-modeling techniques have origins within industry and academia. Denning et al. developed *Security Cards*, a brainstorming framework for modeling threats in four areas: human impact, adversary’s motivations, adversary’s resources, and adversary’s methods [57]. Through the use of a deck of cards, users are prompted to consider common scenarios and contemplate how those scenarios may apply to their own situation. *Persona Non Grata*, another brainstorming framework, also focuses on attacker motivations and capabilities through the lens of an attacker’s perspective [47]. By understanding motivations and intent, PNG users can build a list of ways an adversary may affect a digital system [47]. Microsoft developed *STRIDE*, a step-by-step framework for determining how threats may affect a particular system or component [136, 137]. STRIDE automatically generates likely threats within the following areas: spoofing identity, tampering with data, repudiation, information disclosure, denial of service, and elevation of privilege. By identifying how an adversary may affect a system, defenders may plan defenses. STRIDE delegates threat priority determination to its users based on specific requirements. Both these approaches promote brainstorming, but security cards and PNG do not propose threat mitigations nor do they prioritize identified threats by criticality. As with any brainstorming approach, the output of these methodologies

is dependent on the insight, experience, and creativity of its users.

The *attack tree* methodology provides a formal structure for representing attacks on a system [179]. The root of an attack tree is typically the adversarial goal, and each leaf represents different ways to achieve it [177]. Defenders can traverse the tree to identify which leaves require action and which are adequately mitigated [184], allowing defenders to narrow their mitigation focus. This allows defenders to narrow their scope and focus on mitigating a subset threat vectors. Yet prior work has identified a key weakness of attack trees: without a complete set of root nodes, defenders will fail to account for a large number of potential attacks [184].

The Center of Gravity (CoG) framework focuses on both internal and external threats [103, 147], helping users identify critical threats in a structured way rather than leaving them to assign threat priority without guidance [136, 137]. It requires that defenders define their environment in a top-down approach, ensuring that they build a comprehensive set of adversarial “root nodes” [179, 184]. Prussian military theorist Carl von Clausewitz coined CoG in his 19th century book *On War* [222]. As an analysis of Napoleon I’s military strategy and tactics, *On War* introduced the concept *weight of effort*, a term whose definition varied but is most commonly defined as “the hub of all power and movement, on which everything depends” [219]. Today, military theorists use *Schwerpunkt* synonymously with the modern concept of CoG. The U.S. Department of Defense, which expands upon Clausewitz’s original definition, defines CoG as “the source of power that provides moral or physical strength, freedom of action, or will to act” [79]. Taking into account lessons learned from unconventional and guerrilla warfare, military theorist Joe Strange defines CoG

as “the primary sources of moral or physical strength, power, and resistance” [196]. Finally, Dale Eikmeier offers a nuanced approach, stating that the CoG is the “primary entity that possesses the inherent capability to achieve the objective” [64]. Modern military theorists have expanded the idea of *Schwerpunkt* to account for the moral and mental components of warfare [79] and unconventional and guerrilla warfare [196]. Eikmeier defines the modern, nuanced concept of the CoG as the “primary entity that possesses the inherent capability to achieve the objective [64].” Military strategists have proven the effectiveness of CoG as a proactive framework for conducting and winning real-world military operations. For example, the United States decisively defeated the Iraqi Army in Operation DESERT STORM by developing campaign plans using the CoG framework [198]. By identifying the Iraqi centers of gravity at the strategic, operational, and tactical levels, the United States optimized the use of its military strength to find and exploit Iraqi weaknesses and swiftly defeat the world’s fifth largest military [52]. CoG is applicable within any contested domain [198] and its various definitions are synonymous with the concept of centrality, which appears in network theory for social groups [115] and network theory in the digital domain [206]. CoG can be used for offensive cyberspace operations [46] and internally prioritizing digital defenses [50].

While a number of threat-modeling frameworks and approaches have been defined, few have been empirically evaluated, and none have been assessed at the enterprise level. Sindre and Opdahl compared the effectiveness of *attack trees* and *misuse cases*, an integrated approach for viewing security issues within a system’s architecture for discovering threats, using groups of students [166, 185]. They found

that attack trees were more effective overall, especially in cases where participants were asked to describe threats without an existing model. Similarly, Karpati et al. compared two misuse cases approaches: traditional *misuse cases* and *misuse case maps*, again with student participants [114]. Labunets et al. conducted an empirical study with groups of student participants [121] to measure the effectiveness and perception of *CORAS* [128], a visual framework for modeling risk-based security, and *SREP* [134], a textual framework. In this study, the researchers asked groups of students to discover threats and security requirements. *CORAS* allows users to develop brainstorming diagrams, showing relationships between assets, threats, risks, and security requirements. *SREP* allows users to describe similar information using table-based templates or structured paragraphs. They found that the visualization method aided in identifying threats, yet there was no statistically significant difference between either model for enumerating security requirements. Massacci et al. [131] used small groups of industry practitioners and students to compare the performance of four threat models [77, 85, 128, 146] against fictional scenarios in a classroom environment. They found *CORAS* to be generally the most useful, with participants citing comfort with its step-by-step methodology and easy-to-understand terminology.

2.2 Compliance programs

The following is a review of prior work on digital-security compliance programs and studies of their impact on security and organizations. Current research on com-

pliance programs focus on assisting organizations with better achieving compliance.

Digital security compliance programs within the United States date back to the Computer Security Act of 1987, which required agencies to protect sensitive systems and conduct security training [156]. Many programs implement a “carrot-and-stick” approach to compliance, in that organizations are rewarded for successful programs and levied with sanctions for compliance deviations. This section briefly reviews past studies involving digital security compliance and its impact on organizations.

Compliance audits force organizations to balance being “inspection ready” and sustaining daily operations, such as providing essential services or selling goods. Because of this careful balance, many organizations choose to perform compliance actions only before a pending audit, and then neglect further security maintenance until another audit requires them to repeat the process [170]. This behavior meets the security minimums for compliance standards, but fails to adhere to the spirit of secure practices. Moreover, evidence shows that fully-compliant organizations can still suffer data breaches. Auditors certified Target as PCI-compliant in September 2013, just before it suffered a massive data breach in November 2013 [170]. These factors show that organizations must have periodic security checks in place that exceed compliance minimums and proactively reduce risk beyond baseline technical controls and implementation processes.

Previous studies highlight cultural disconnects between developers, engineers, and compliance officials that create issues when digital security measures are “bolted on” after software development is complete [17, 45]. Clark highlights institutional changes that had to occur in his organization to incorporate NIST security guide-

lines and cybersecurity framework to ensure security was measurable, effective, and individuals at the organization took ownership of their role in security [45]. Clark recommended changes such as embedding compliance experts within development teams to encourage grass-roots-style compliance integration [45]. Assal and Chiasson discuss the importance of integrating security throughout the software development process, the sometimes necessary organizational restructuring that may need to occur to support integration, and other organizational optimizations that need to occur to support more secure behaviors [17]. Thomas et al. conducted a systematic study of application developers to understand how they reason about compliance and security to identify ways to overcome organizational behaviors and improve tools for secure software development [209].

Other organizations found that threat modeling could proactively identify security gaps that may exist in compliant solutions [11, 45]. Some organizations have even overhauled their physical network topology to meet federally-mandated requirements, restructuring their teams and network architecture to limit the scope of auditable systems within their environment [101]. In practical terms, there are significant impacts on service availability, financial resources that may be required for redundancy during re-engineering, and even contractual considerations for reorganizing employees and contractors.

Numerous studies focus on how humans perceive compliance standards and modify their behaviors based on those perceptions. Julisch highlighted numerous factors that shape organizational decision-making when investing in compliance measures, often seeking new security technologies that are out-of-the-box compliance

ready [113]. This creates a market that offers solutions to administrative requirements but these solutions are not plug-and-play; extensive research is required for the integration of these security technologies so that the security controls are fully utilized and that new gaps are not created. Beautement et al. describe the “compliance budget,” the human factors behind the implementation of compliance controls; their research illuminated ways to improve security and compliance readiness through resource allocation optimization [23]. In their research, they highlight approaches that managers can use to influence employees and improve security behaviors. Much like a financial budget, Beautement et al. advocates for prioritization of efforts as to not exceed the compliance budget, or reach a point that investments begin to have a diminishing return. Building upon previous works, Puhakainen and Siponen found that training employees to better understand compliance standards can improve organizational behaviors and shift employees toward implementing more secure practices [172]. Their research finds that motivation is key for intervention training and employees need to be shown the direct correlation between their habits and system security. The study further highlights the importance of communication between organizational leaders and employees to reinforce better security practices.

Additionally, Hu et al. found that managers who “lead by example” and implement top-down management initiatives encourage employees’ compliant security behaviors [98].

Correctly understanding and implementing legal obligations in compliance program texts have been studied by many researchers. Breaux et al. focused on the difficulty of implementing compliance programs, specifically the ambiguity and com-

plexity of the legal language that is used to describe rights and obligations that compliance programs require [31–34]. Similarly, Agarwal et al. proposed a flexible and modular compliance assessment framework that would help companies understand their legal obligations [5].

2.3 Incident Response Playbooks

The following is a review of prior work on digital-security incident response playbooks and other types of incident preparedness. Current research on playbooks focuses on identifying likely vulnerabilities, likely threats, best practices for response, and ways to better prepare organizations and employees for threats.

Digital security *playbooks* — a structured action plan for incident response — are designed to help organizations prepare for security breaches and enable them to respond quickly and appropriately. High-stress situations such as an ongoing data breach may disrupt technicians’ cognitive abilities [60, 97, 125]; playbooks are designed to present documented best practices to prompt action and momentum during incident response, as well as supporting post-incident root-cause analysis. A *local instance* of a playbook is specifically tailored to a particular site and should include fine-grained controls and actions for responding to specific events.

The following works represent types of playbooks in other domains that may be applicable to digital security, previous crisis readiness proposals, and a look at two exemplar frameworks used for designing incident response playbooks.

2.3.1 Crisis preparedness in other domains

Business continuity plans (BCPs) help minimize financial losses, ensure the continuation of core functions, ensure resource availability, and prepare employees through training. Many organizations are required by insurance or regulations to have BCPs. Numerous references provide reporting templates for communicating essential information and how-to guides for audits [27, 94]. BCP training varies, but typically involves walk-through rehearsals to prepare for disruption events [94]. Other researchers have focused on aggregating lessons learned and training scenarios for a vast array of situations that may cause damage to a business: terrorist attacks, supply chain disruptions, and even managing negative media coverage [80]. BCPs typically contain fine-grained detail to assist with implementation and auditing.

Federal government agencies maintain playbooks for natural disaster continuity and health emergency preparedness, among other crises [224]; libraries of pre-made disaster response playbooks are available for reference [67, 148, 227].

In the medical field, crisis resource management combines standard medicinal practices with non-technical skills to ensure exposure to best practices for likely emergency situations [41]. Crisis resource management includes fuses standard medicinal practices with non-technical skills such as labor distribution (who should be doing what and when), communication (who needs to know what and when), and anticipation and planning (understanding the likely steps that should follow) [41]. Studies found that simulated rehearsals allowed participants to rehearse their response action “playbooks” and that participants felt confident that

the lessons learned would transfer to real-life situations [175].

Pilot training uses simulation to allow aviators to prepare for dangerous situations (and even cyber attacks) prior to ever entering a real cockpit [68, 186]. This readiness and preparedness also extends beyond the cockpit, with much emphasis placed future readiness when presented with large-scale disruptions such as natural disasters or acts of terrorism [59, 92]. Allowing international organizations and pilots alike to rehearse and refine their playbooks improves their ability to handle threats when they arise.

2.3.2 Threat exposure and readiness

Previously, I proposed an incident response readiness concept known as digital calcification [189]. The premise of this proposal was that data breaches, insider threats, and other forms of cyber attacks are pervasive enough that organizations must be prepared to respond to likely security incidents. Unprepared organizations can exacerbate the impact of these threats, leading to a loss of consumer trust and confidence. Organizations should allow trusted entities to attack systems and develop a comprehensive understanding of security gaps; repeated exposure to these attacks may lessen the stress for incident responders over time and improve the security of the environment overtime by implementing lessons learned from the events. This “calcification” concept — the development of stronger structures that can better withstand damage — fused together concepts like the Netflix “Chaos Monkey” with incident response playbooks. The Chaos Monkey concept is that no system is

off-limits during stress testing and that critical resources should be resilient enough to withstand cyber attacks, power outages, and system failures [24]. Calcification should include all personnel at an organization, from technicians to C-level executives, so that each individual can understand the role they play in crisis readiness and lessen the effect of successful attacks in the future.

Additionally, at the International Conference on Cyber Conflict, I proposed novel methods for exposing the U.S. military to many modern-day threats posed by the integration of digital networks within lethal combat platforms [190].

2.3.3 Incident response playbooks

Playbook frameworks allow security professionals to (1) design playbooks that can be tailored to their unique environment, (2) engineer solutions that fulfill security requirements, and (3) describe the actions that defenders should take during incident response events. Additionally, playbooks guide responders towards executing root-cause analysis using industry best practices and providing essential “first steps” that help prevent inaction during the early phases of an incident response effort.

Bollinger et al. highlight using playbooks to secure high-value systems, defend against various threat models, conduct investigations, generate reports based on findings, and modernizing your technologies to remain relevant against emerging threats [28]. In their book, Bollinger et al. present lessons learned from their creation of an incident response team at Cisco and share their insights into organizational

planning, adaption, training, and integrating lessons learned iteratively. Other researchers discuss playbook benefits such as preparing incident response checklists that identify “must-do” tasks, report templates, and communication templates to share with stakeholders as part of business continuity plans; they also identify playbooks playing an essential role in education and helping professionals adjust to a new program or responsibility [89, 138]. Others advocate for using playbooks to prepare an entire organization for incident response. C-level executives and network defenders alike have a critical role to play during an incident; playbooks allow each individual to proactively understand their role and prepare responses before an incident occurs [138].

Security professionals may use playbook frameworks during cybersecurity training exercises to design playbooks and prepare for future incident response events that may occur within their corporate networks [116]. These cybersecurity exercises immerse technicians within developed scenarios to measure and evaluate their ability to perform their security functions [116]. While some scenarios pit teams of professionals against one another, nearly all exercises involve notional events within practice networks that are not truly representative of the real-world incidents or environments [190]. Using cybersecurity exercises to train technicians with non-native tools or networks may coach “exercise-isms” that do not translate into real-world scenarios and may significantly inhibit the transfer of knowledge from training into reality. Executing cybersecurity exercises organically at organizations may require extensive time and financial resources, but it helps ensure lessons learned from the exercises are tailored to the networked environments within which participants nor-

mally operate.

Numerous playbook frameworks exist [15,28,116,138,165], but we focus on the following two frameworks given their support from the United States Government and they have no-cost, openly-available guides for use.

IACD. The IACD framework is the result of collaboration between the U.S. Department of Homeland Security, National Security Agency, and Johns Hopkins Applied Physics Laboratory to leverage automation within incident response [104]. The defining feature of IACD playbooks is a visual approach for capturing essential response actions that both humans and automated systems must take to handle an incident. The IACD website has numerous, publicly available guides and examples for practitioners to reference [105].

Through case studies, IACD provides a method by which defenders can (1) conduct proof-of-concept evaluations of new sensors and technologies; (2) provide insights to potential challenges that an organization may face; (3) identify gaps in technology, policies, or standards; and (4) gather requirements to facilitate standards development [104]. IACD playbooks are intended to capture security processes that reflect governance or regulatory requirements as well as industry best practices [105].

The IACD framework breaks playbook design into 10 steps. The first step is to identify the initiating condition: the event or situation that triggers use of the playbook (e.g., a database breach) and how that event is detected (e.g., an automated email alert sent to an administrator). The second step involves listing all possible actions that could occur in response to the initiating condition, typically via mind

mapping. Practitioners should reference existing best practices to identify possible actions. Next, playbook designers designate each identified action as required or optional. For example, generating a written report that details the incident from beginning to end — which may provide invaluable insight after the event but does not contribute directly to response efforts — should be labeled optional. Optional tasks are aggregated in an *action options box*: a menu of available tasks that can assist with the investigation but may not be required. Steps 4-7 involve grouping actions by function, ordering required actions sequentially, and interleaving optional actions where appropriate. The designer produces a diagram showing these ordered relationships.

Next, update the action options box. Listed items in the action options box typically involve designing automated-system prompts for humans-in-the-loop to authorize a follow-on action or designing actions that humans will manually queue for automated execution. An example would include an automated prompt asking for human confirmation to disconnect a system from a specified network. In step 9, the designer verifies that the playbook terminates either in a desired end state or in a new initiating condition that flows into another playbook. The final step ensures that the playbook satisfies applicable regulatory controls and requirements.

NIST. The NIST Computer Security Incident Handling Guide (hereafter referred to as the NIST framework) focuses on expeditious recovery after a security incident [44].

Using this framework, designers break a security incident down into three phases, creating playbook content for each phase. The preparation phase occurs

before an incident occurs and requires analysts to identify critical assets that must be protected during the incident of the type under consideration. Playbook content for the detection and analysis phase should help defenders identify the incident’s entry point, breadth of impact, potential consequences, and containment methods. Content for phase three — containment, eradication, and recovery — should guide defenders in patching or isolating the attacker’s entry point and other similar potential entry points, increasing monitoring, and safely bringing services back online. Each phase of a NIST playbook should emphasize communication and metrics tracking: ensuring essential personnel are informed, victims are notified, and the scope of impact is thoroughly documented. While playbook designers may or may not deem communications as required actions in IACD playbooks, communication is required throughout NIST playbooks.

Unlike IACD, the NIST guide does not typically result in a visualization of response actions (although it could). Instead, a NIST playbook typically provides detailed textual descriptions, intended to be drawn from institutional procedures or best practices, for each phase.

2.4 Human factors in security operations

The following is a review of prior work on human factors within security operations centers and their overall impact on security.

János and Dai found that organizational behavior and culture impacted the efficacy of security operations [110]. Kokulu et al. similarly found numerous de-

iciencies within SOCs stemming from insufficient training, poor communication, and evaluation criteria disconnected from meaningful performance metrics [117]. Research from Alomar et al. discusses breakdowns in trust, communication, and resourcing that inhibit the effectiveness of vulnerability disclosure programs [8]. Furnell et al. identified multiple usability concerns in incident response tools as well as the occasional need for internally-developed tools [72]. Sundaramurthy et al. highlight the consequences of “build-once-sell-to-everyone” security vendor models on SOCs and also found that in-house, tailored solutions may best support analysts’ needs [203].

Focusing on security analyst performance, Sundaramurthy et al. observed burnout rates within SOCs and identified possible solutions for sustaining morale and completion of security tasks [202]. Dykstra and Paul found that analysts’ fatigue and stress levels increase throughout the day, affecting their ability to perform security tasks and suggesting analysts’ tools and environment need to offset frustration where possible [60]. Other researchers focused on reducing the impacts of information overload to help incident responders improve mitigation efforts against true-positive attacks [87].

Chapter 3: Planning for Security: Human Factors in Threat Modeling Adoption

In this chapter¹ I present the first case study of threat modeling as a proactive planning framework in a large, high-risk enterprise environment. During proactive planning, threat modeling frameworks present a methodology for evaluating complex threats and breaking them into modular components. This is intended to allow humans to reason about the threat, understand risk, and develop mitigating strategies [184].

Through a partnership with New York City Cyber Command (NYC3), I introduced 25 employees to an exemplar threat-modeling approach through group training sessions. I tracked the impact of this threat modeling training on NYC3's security posture quantitatively, through analysis of 120 days of log data, and qualitatively, via pre-, post-, and 30-day-post-training surveys with participants.

The results suggest that threat modeling may provide valuable benefits in an enterprise setting. Participants' perceptions of threat modeling were very positive: after 30 days, 23 participants agreed that it was useful in their daily work and 20 reported that they have adopted its concepts in their daily routine. Collectively,

¹Published as [191, 195]

participants developed 147 unique mitigation strategies, of which 64% were new and unimplemented within NYC3. Additionally, participants identified new threats in eight distinct areas within their environment, such as physical access-control weaknesses and human configuration errors. Within one week of developing these plans, NYC3 employees started implementing participant-designed plans to mitigate these eight newly-identified threat categories. In the 120 days following the study, NYC3 implemented participant-designed defensive strategies that prevented five privileged account hijackings, mitigated 541 unique intrusion attempts, and remedied three previously unknown web-server vulnerabilities. The observations and metrics from this study provide a scaffolding for future work on threat modeling and enterprise-employee security training.

With regard to proactive measures as whole, Chapters 4, 5, 6, and 7 show that threat modeling is insufficient independently as it does not provide practitioners with action plans for when incidents are underway — although this study shows that threat models reveal likely areas of vulnerability and help develop plans to optimize current resources to lessen risk. Threat models and risk mitigation strategies must also be compatible with applicable baselines for integration into real-world security operations centers.

3.1 The Center of Gravity framework

In this study, I introduced NYC3 employees to the Center of Gravity (CoG) framework, which originated in the 19th century as a military strategy [222]. As

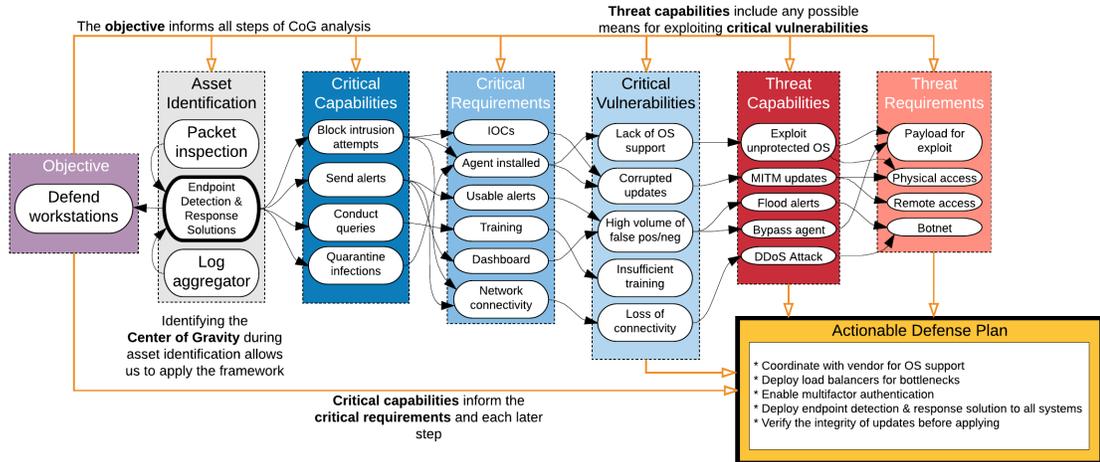


Figure 3.1: Step-by-step process for threat modeling with CoG, using participant P17’s responses as an example.

a military concept, a center of gravity is the “primary entity that possesses the inherent capability to achieve the objective [64].” As a threat modeling approach, CoG focuses on identifying and defending this central resource. This approach is applicable within any contested domain [198] and is synonymous with centrality, which appears in network theory for social groups [115] and network theory in the digital domain [206]. CoG supports planning offensive cyberspace operations [46] and prioritizing digital defenses [50].

The constraints of the partnership with NYC3 — in particular, the requirement to minimize employees’ time away from their duties — only allowed researchers to introduce and examine one threat modeling framework. I selected CoG because it incorporates many key characteristics from across more pervasive frameworks: CoG provides practitioners with a top-down approach to identifying internal points of vulnerability, similar to STRIDE [136, 137], and it assists with assessing vulnerabil-

ities from an adversarial perspective, similar to attack trees, security cards, persona non grata, and cyber kill chain [47, 57, 103, 179]. Uniquely among popular threat modeling approaches, it allows organizations to prioritize defensive efforts based on risk priority.

It is essential to understand the process of applying the CoG approach. Figure 3.1 illustrates these steps using an example provided by one participant.

To begin using CoG, analysts must start by codifying the long-term organizational objective, or “end state,” of defensive measures. An end state provides the *why* for implementing defenses and allows an individual practitioner to understand their own specific security objective with respect to the organization.

Once the practitioner understands their objective, the next step is to identify all of the assets currently in use that support accomplishing the objective. In this context, an asset can be a system, a service, a tool, or anything relevant to accomplishing the objective (not just security-specific assets). The practitioner then identifies the CoG as the pivotal asset on which all other assets depend for accomplishing the objective.

Once the practitioner identifies the CoG, they can deconstruct it into three components: critical capabilities, critical requirements, and critical vulnerabilities [64]. *Critical capabilities* (CC) are distinguished by two key features: they support the practitioner’s objectives, and the CoG would cease to operate without them [79]. For each CC, there are one or more *critical requirements* (CR), defined as supporting resources that enable the CC to function [79]. Eikmeier distinguishes between capabilities and requirements using a “does/uses” litmus test [64]: If the

CoG does something, that something is a capability, and if it uses something, that something is a requirement. *Critical vulnerabilities* (CV) are directly related to critical requirements; CVs are thresholds of diminished CRs that make the CoG inoperable [181]. Practitioners identify CVs by asking the following question for each CR: what would cause this requirement to no longer function as intended? Some CVs are binary, such as the complete loss of a CR, but others may cause a reduced functionality beyond some threshold, preventing the CoG from accomplishing the objective.

Building a thorough list of critical vulnerabilities allows the practitioner to understand how their objectives can be threatened. The practitioner should consider both malicious and accidental threats to collectively describe the worst-case situation for their organization and objectives. The CoG approach models all threats with a singular, unified motivation: exploiting critical vulnerabilities. This allows practitioners to develop a list of threats that can encompass nation-state hackers, insiders, poorly trained users, and others. The practitioner iterates over the list of critical vulnerabilities to develop a corresponding list of *threat capabilities* (TC). For each CV, they ask: what could take advantage of this vulnerable condition? From the list of TCs, they enumerate all of the threat requirements (TR) needed to support each capability.

The final step in the CoG analysis process is building an *actionable defense plan* (ADP) that can neutralize identified threat capabilities and requirements, mitigate critical vulnerabilities, and protect the identified CoG. Each component of an ADP, designed to dampen or eliminate one or more potential risks, is referred to as a

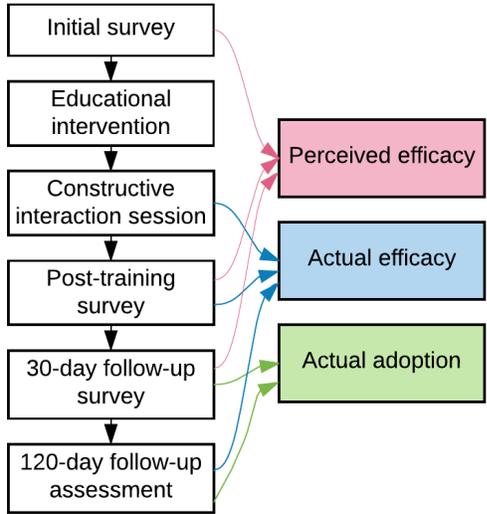


Figure 3.2: The six-part study protocol and metrics.

mitigation strategy.

3.2 Methods: Threat modeling at NYC3

To evaluate the impact of introducing threat modeling to an organization that had not previously used it, I partnered with NYC3 to introduce a specific threat-modeling framework (CoG) and observe the effects. NYC3 is responsible for protecting the most populous city in the U.S. and its government from cyber attacks. The Government of the City of New York (GoNYC) includes 143 separate departments, agencies, and offices with more than 300,000 employees that support 8.6 million residents and 60 million yearly visitors [161]. It maintains nearly 200,000 external IP addresses and has its own Internet Service Provider, with hundreds of miles of fiber-optic cable and dozens of major points of presence. Further, the city is responsible for maintaining industrial control and mainframe systems. The

participant pool consisted of civil servants and private-sector contractors who work directly with NYC3.

This study focuses on the *efficacy* of threat modeling, which in this context is defined as the ability to achieve a desired outcome. Both *effectiveness*, the ability to successfully achieve an outcome, and *efficiency*, the ability to reduce effort to achieve an outcome, comprise efficacy.

Because of introducing threat modeling in NYC3’s operational environment, a comparative experiment was not possible; instead, I designed a primarily observational study to obtain as much insight as possible — both qualitative and quantitative — into the effects of introducing threat modeling within an enterprise environment. The study includes six components (as shown in Figure 3.2), that occurred from June through November 2017, and was approved by the University of Maryland Institutional Review Board. Due to the study’s sensitive nature, some details are generalized about defenses and vulnerabilities to protect NYC. Additionally, sensitive information is redacted when quoting participants and generalized job descriptions so as to not deanonymize participants.

3.2.1 Recruitment

NYC3 leadership sent all of its employees an email that outlined the voluntary nature of the study as well as the motivation and goals. The email informed NYC3 employees that they would be introduced to new techniques that could potentially streamline their daily duties, and that the findings from the study would

be directly applied to defending NYC3 systems and networks. The study occurred during participants' regularly scheduled work hours and participants did not receive any additional monetary incentives for participating.

3.2.2 Study protocol

The multi-part study protocol is described as follows.

Protocol pilot. Prior to deploying the protocol with participants, researchers conducted three iterations of the study using non-NYC3 employees (two security practitioners and one large-organization chief information security officer) to pre-test for relevance, clarity, and validity. The study protocol was updated based on pilot feedback and overall study flow. The final protocol described below, derived from three pilot iterations.

Baseline survey. Establishing a baseline for NYC3 defensive practices allows researchers to compare the security posture before and after the training intervention. Participants were asked about their specific work role, responsibilities, and demographics; their understanding of organizational mission statements; which assets they use to accomplish their daily duties; their sentiment towards NYC3's current security posture; and their perceived self-efficacy for performing digital security tasks.

The 29-question online survey (App. [A.1](#)) used a combination of open-ended, close-ended, and Likert-scale questions. All self-efficacy questions were based on

best-practices and question-creation guides from established educational psychology studies [20]. The post-training survey and 30-day follow-up survey used an identical structure . Capturing self-efficacy before, immediately after, and 30 days after receiving the educational intervention allowed measurements of how each participant perceived the model’s efficacy. Measuring efficacy perceptions is important, as self-efficacy has been shown to be an important component of individual success at performing job duties in enterprise settings [19]; one key component of self-efficacy is belief in the efficacy of the tools you use to complete tasks.

Educational intervention. After completing the initial survey, groups of participants received in-person instruction on the history of CoG, its application as a threat modeling technique, the CoG process outlined in Section 3.1, and two examples of applying the framework. Participants completed one of three independent sessions based on what was most convenient to their work schedule.

The 60-minute educational intervention was based on on fundamentals from adult learning research and the experiential learning theory (ELT) [118]. Kolb and Kolb found that adults learn new concepts better through ELT by (1) integrating new concepts into existing ones, (2) accommodating existing concepts to account for the new concepts, and (3) “experiencing, reflecting, and acting” to reinforce the new concepts [118]. Social learning theory (SLT) further supports this process, indicating that adults learn new patterns of behavior best through direct experience [21]. Thus, the class was designed to reinforce each concept with a hands-on exercise using scenarios relevant to the audience and their domain knowledge.

During the class, the instructor introduced participants to tabular and graph-based methods performing CoG analysis [197]; both examples are provided in App. B.1.4. The tabular tool allows users to record their responses to each subtask of the CoG framework; each section supports data in follow-on sections. The graph-based method provides users with an alternative, complementary method for eliciting the same data. Previous research indicates that various learning styles benefit from multiple forms of data elicitation [118].

During the first classroom example, the instructor guided participants through a scenario drawn from the Star Wars movie franchise to determine the CoG for the Galactic Empire. The instructor provided step-by-step instructions for using the tabular and graphical tools throughout. In the second example, the participants worked together without instructor guidance to apply CoG and framework tools to a fictional e-commerce scenario. Both fictitious scenarios are described in App. B.1.

Prior to providing the intervention, the instructor observed NYC3 employees at work for four days to better understand their operating environment. The instructor developed the fictitious scenarios so that they did not reflect any specific conditions within NYC3. These scenarios served as better tools in lieu of NYC3-specific scenarios to reduce bias during training that would inadvertently coach participants towards providing “approved solutions.”

To control for variations in instruction, each group had the same instructor. The instructor is a member of the research team with extensive subject-matter knowledge and experience, including six months of formal university training on threat modeling. The instructor communicated this experience prior to each class

to establish a baseline of credibility with the group. During each class, participants could ask questions at any time, and the instructor maintained a running log of these questions. To maintain consistency across class sessions, the instructor incorporated answers to these questions at relevant points in future sessions, and emailed the answers to participants who had attended previous sessions.

Performance evaluation session. After all participants finished the educational intervention training, they each completed a 60-minute individual session where they applied CoG to their daily duties. For example, P17 used the framework in his role as a security analyst to develop plans for better defending NYC endpoint workstations (See App. B.1.3). This phase of the study provided hands-on reinforcement learning, as recommended by ELT and SLT [21, 118].

Each session was audio recorded and the interviewer provided participants with clean worksheets and whiteboards for brainstorming (App. B.1.4), and allowed participants to bring in any notes from the previous educational intervention training. Without notifying the participants, the interviewer logged task completion times for each step, in an effort to measure the efficiency of the framework without putting undue pressure on participants.

The interviewer used the constructive interaction method for communicating with the participants, asking them to openly communicate throughout each subtask in Section 3.1 [144]. During each step, the instructor re-stated participants' previous verbal comments or documented responses to assist with data elicitation but did not introduce any new concepts to prevent data bias. For consistency, the same

interviewer completed all performance evaluation sessions.

At the completion of each session, the interviewer retained a copy of the completed worksheets, photographed the whiteboards, and returned the original worksheets to the participant to help guide their responses for the second online survey. The aggregated worksheets and time logs support measurements for the actual efficacy of the CoG framework (See Section [3.3.3.2](#)).

The performance evaluation interviewer transcribed responses to the open-ended questions after each session using the audio recordings. Two researchers jointly analyzed all open-ended survey questions and each transcription using iterative open-coding [199]. In alignment with this process, researchers coded each research artifact and built upon the codebook incrementally. Researchers resolved all disagreements by establishing a mutually agreed upon definition for coded terms. From here, researchers re-coded previously coded items using the updated codebook and repeated this process until all responses were coded, all disagreements resolved, and the codebook was stable.

Post-training survey. In this 27-question online survey (App. [A.1](#)), conducted immediately after the performance evaluation session, surveys collected responses measuring the framework’s actual and perceived efficacy. The survey asked participants to re-apply CoG to their daily duties, which allowed them to account for any new details they might have considered since the previous session. Additionally, surveys asked them to re-evaluate their perception of the NYC3 baseline security posture and their ability to complete digital security tasks. Using this information,

it is possible to measure changes in how participants view the organization and their own abilities [71]. Further, the survey asked participants to evaluate their ability to complete digital security tasks solely using the CoG framework and to answer comprehension questions measuring their current understanding of the framework.

Follow-up survey. The 13-question follow-up survey (App. A.1) measured framework adoption, knowledge retention, and perceived efficacy 30 days after researchers departed. To measure the extent to which participants adopted CoG analysis without instructor stimulus, surveys asked participants to describe whether and how they used the information derived from CoG analysis or the framework itself within their daily duties. These questions allow researchers to understand participants' ability to apply output from the framework, measure their adoption rates at work, and measure their internalization of CoG concepts. The survey also continued to use self-efficacy questions supplemented with survey questions from the technology acceptance model (TAM) [55].

Long-term evaluation. After 120 days, researchers evaluated the efficacy of adopted defense plans for protecting NYC3 systems. The evaluation used a combination of NYC3 incident reports and system logs extracted solely from defensive measures that participants recommended and implemented because of their use of CoG threat modeling. NYC3 deployed these new defensive measures in “blind spots,” so each verified intrusion attempt or vulnerability clearly links an improved security posture to these new defensive measures.

3.2.3 Limitations

All field studies and qualitative research should be interpreted in the context of their limitations.

One threat-modeling framework was used in this study: although the sample represents 37% of the NYC3 workforce, 25 participants (in many cases with no overlap in work roles) would not have been sufficient to thoroughly compare multiple approaches. Testing multiple models within participants was impractical due to the strong potential for learning effects and the need to limit participants' time away from their job duties. As such, it is possible that other threat-modeling or training approaches would be equally or more effective. However, the results still provide insight as to how threat modeling in general can benefit a large enterprise.

Described in Section 3.3.3.2 below, two NYC3 leaders jointly evaluated the defense plans produced by the participants. More, and more independent, evaluators would be ideal, but was infeasible given confidentiality requirements and time constraints on NYC3 leadership.

The results may be affected by demand characteristics, in which participants are more likely to respond positively due to close interaction with researchers [96, 167, 212]. This is mitigated through (1) anonymous online surveys that facilitated open-ended, candid feedback, (2) removing researchers from the environment for 30 days before the follow-up survey, and (3) collecting actual adoption metrics. Further, there may be selection bias in which those NYC3 personnel most interested in the topic or framework were more likely to participate because the purpose of

the study was explained during recruitment,; this is mitigated by asking NYC3 leaders to reinforce that (non-)participation in the study would have no impact on performance evaluations and by recruiting a large portion of the NYC3 workforce.

NYC3's mission, its use of pervasive defensive technologies, and its adherence to common compliance standards indicate that NYC3 is similar to other large organizations [108,150,151]; however, there may be specific organizational characteristics of NYC3 that are especially well (or poorly) suited to threat modeling. Nonetheless, the results suggest many directions for future work and provide novel insights into the use of threat modeling in an enterprise setting.

TAM has been criticized (e.g., by Legris et al. [123]) for insufficient use coverage. Additionally, the positive framing of TAM questions may lead to social desirability biases [62]. To address coverage, surveys use TAM in conjunction with the Bandura self-efficacy scales for a more complete picture. Moreover, reusing validated survey items and scales in this study is a best-practice in survey design that has been shown to reduce bias and improve construct validity [69, 83]. Lastly, surveys elicited participant feedback with a negative framing explicitly after each performance evaluation session, and implicitly when assessing threat modeling adoption at the 30-day evaluation. Eliciting feedback through negatively-framed mechanisms allowed participants to provide their perceptions from both perspectives.

For each qualitative finding, a participant count indicates prevalence. However, participants who did not mention a specific concept during an open-ended question may simply have failed to state it, rather than explicitly disagreeing. Statistical hypothesis tests are not used for these questions.

3.3 Results

Below are the results of the case study evaluating threat modeling in an enterprise environment, drawing from transcripts and artifacts from performance evaluation sessions, survey answers, and logged security metrics. Participant demographics, baseline metrics, immediate post-training observations, 30-day observations, and observations after 120 days are reported.

The findings are organized within the established framework of perceived efficacy, actual efficacy, and actual adoption [121, 145, 166]. Participants' perceived efficacy and belief that they will achieve their desired outcomes directly shape their motivation for adopting threat modeling in the future [18]. Actual efficacy confirms the validity of perceptions and further shapes the likelihood of adoption. Lastly, regardless of perceived or actual efficacy, a framework must be adopted in order to demonstrate true efficacy within an environment. Through these three measurements, the study provides security practitioners with the first structured evaluation of threat modeling within a large-scale enterprise environment.

3.3.1 Participants

Qualitative research best practices recommend interviewing 12-20 participants for achieving data saturation in thematic analysis [84]. To account for employees who might need to withdraw from the study due to pressing work duties, 28 participants were recruited for the study. Of these, 25 participants completed the study (Table 3.1), above qualitative recommendations, and also reached saturation in the

ID	Duty Position	IT Experience	Training (yrs)	Education
P01	Leadership	16-20	6-10	SC
P02	Data Engr.	16-20	6-10	G
P03	Sec Analyst	11-15	0-5	SC
P04	Sec Engineer	11-15	0-5	BS
P05	Governance	16-20	6-10	SC
P06	Sec Engineer	6-10	11-15	P
P07	Sec Engineer	0-5	6-10	G
P08	Net Admin	21-25	6-10	G
P09	Sec Engineer	11-15	0-5	SC
P10	Sec Engineer	11-15	6-10	BS
P11	Net Admin	16-20	6-10	BS
P12	Sec Engineer	25+	6-10	G
P13	Sec Analyst	0-5	0-5	BS
P14	Sec Engineer	11-15	0-5	BS
P15	Sec Engineer	16-20	25+	SC
P16	Support Staff	6-10	0-5	BS
P17	Sec Analyst	16-20	16-20	G
P18	Sec Engineer	21-25	16-20	G
P19	Sec Analyst	21-25	6-10	SC
P20	Leadership	11-15	6-10	G
P21	Sec Analyst	0-5	6-10	G
P22	Leadership	11-15	6-10	G
P23	Sec Analyst	16-20	6-10	BS
P24	Leadership	0-5	0-5	BS
P25	Leadership	0-5	0-5	G

Table 3.1: Participant Demographics. The columns show: participant identifiers (coded by interview date order), duty position, years of IT experience, years of IT training, and education. The abbreviations in the education column stand for high school graduate, some college, bachelors degree, associates degree, graduate degree, and prefer not to answer.

performance evaluation sessions. For the rest of this paper, all results refer to the 25 participants who completed the study. This sample represents 37% of the NYC3 employees as of August 8, 2017.

Technicians such as network administrators and security engineers account for 18 of the participants; the remainder fulfill supporting roles within NYC3 (e.g., leadership, policy compliance, and administrative support). This composition is similar to the actual work role distribution across NYC3, with 50 of 67 employees serving

as technicians. Prior to this study, one participant had a high-level understanding of the military applications of CoG, and none of the participants had any applied experience using any threat-modeling framework.

All participants had at least some college education, with ten holding a graduate degree and eight holding a bachelor's. Additionally, 15 possessed at least one industry certification. Participants had an average of 14.7 years of information technology and security experience in large organizations, with a mean of 8.5 years of formal or on-the-job training.

3.3.2 Pre-intervention baseline

A baseline of how participants deployed defensive strategies prior to the training helps measure the impact of threat modeling within NYC3 systems. Most commonly, they prioritized defending high-impact service-based systems such as NYC.gov (n=7) and adhering to compliance frameworks (n=7), followed by applying risk management strategies (n=6) and assessing which systems are most susceptible to attack (n=3). Participants reported using the following guidelines and programs for assessing NYC's digital security posture: city-specific policies and executive orders such as the NYC remote access policy [162] (n=6), NIST Cybersecurity Framework [150] (n=4), and NYC3's one-time accreditation process for adding new technologies to their network (n=2). Of these guidelines, participants stated that none of the programs were applied frequently enough, with P5 stating that "compliance is only as good as your last assessment." With too much lapsed

time between audits, defenders cannot establish an accurate assessment of the environment's security posture over time. The remainder of respondents (n=13) said they were unsure about which programs or policies were applicable.

3.3.3 Immediate observations

In contrast to the baseline survey, performance evaluation session observations and post-training surveys indicate that threat modeling provided participants with a better understanding of their security environment, that participants felt more confident in their ability to protect NYC, and that participants could successfully apply threat modeling relatively quickly with accurate results.

3.3.3.1 Perceived efficacy

Participants' initial threat modeling perceptions are recorded in the context of new insights, framework usefulness, and changes in self-efficacy. Participants' perceptions and beliefs that they will achieve their desired outcome directly shapes their motivation for adopting threat modeling in the future [18].

New understanding. Overall, 12 of 25 participants reported that threat modeling allowed them to understand new critical capabilities, requirements, or vulnerabilities that they had never previously considered. In particular, four participants had never previously mapped threats to vulnerabilities. P16, a non-technical administrative support staffer, used threat modeling to understand the implications of wide-open security permissions on a wiki and networked share drive.

Threat modeling provided two participants with self-derived examples of why crisis continuity plans exist for large organizations. P04 stated that this new understanding would further assist him with planning for crises, allowing him to recommend to “senior management the plan of action for what should be done first.”

Of the 13 participants who did not report discovering anything new, seven stated threat modeling was simply a restructured approach to current defensive concepts like defense-in-depth [132]. Four stated threat modeling did not help them discover anything new but added additional emphasis to areas they should be concerned with.

Four participants identified an over-reliance on personal relationships (rather than codified policies) as a critical vulnerability for organizational success, which conceptually is something none of them had ever before considered. During his performance evaluation session, P24 discussed how changes in the political environment from the local to federal level can affect established trust across the GoNYC; a large turnover in personnel could halt some progress and potentially kill some initiatives. P25 stated “I had not really considered...the impact that some sort of major, non-cyber event could have on the ability to be successful,” discussing how a major terrorist event within NYC could decrease NYC3’s ability to sustain critical requirements and capabilities. Thus, both participants recommended codifying existing relationship-based agreements into legislation capable of withstanding non-digital security threats to their daily responsibilities. An example of this includes establishing a formal memorandum of understanding (MoU) with law enforcement agencies in NYC to facilitate the exchange of threat indicators.

Perceived framework usefulness. After completing the performance evaluation session, 23 participants agreed that threat modeling was useful to them in their daily work. For example, ten said the framework allowed them to prioritize their efforts. P24 developed a new litmus test for adding any defensive efforts, stating that “If the adversary doesn’t care, then it’s all just fluff [inconsequential].” P21 used threat modeling to show “what we’re lacking or what we need to concentrate [on],” such as standard cyber hygiene.

Eight participants expressed that threat modeling added much-needed structure and perspective to difficult problems. P11 feels empowered by its structure and believes it allows him to “accept the things you cannot change, change the things you can, and have the wisdom to know the difference. I feel [CoG is] along those lines; this is your world, this is what you control.” He believes threat modeling makes a positive difference with available resources, while helping to prioritize requests for future capabilities and support.

Five participants reported that threat modeling allowed them to plan defensive strategies more effectively. P05 stated that threat modeling helps him “plan effectively, document, track, monitor progress, and essentially understand the security posture.”

Threat modeling allowed four participants to comprehend how threats can affect systems within their environment; these technicians previously relied upon best security practices without fully considering threats. While applying the framework, P10 declared that “insider threats overcome the hard shell, soft core” within most

enterprise networks and that threat modeling helped him identify new ways to neutralize the impact of insiders bypassing perimeter defenses and exploiting trusted internal systems.

Four participants stated that purposefully considering their asset inventory during threat modeling allowed them to fully understand their responsibilities. Three participants stated that threat modeling provides them with a new appreciation for their position within NYC3. P14 said, “When I did my job, I didn’t think about what the purpose of the group is [within NYC3]. . . [threat modeling] aligns what we’re thinking with what I think my role is in this organization.”

Interestingly, both of the participants who did not find threat modeling useful felt that cybersecurity is too nebulous of a realm for a well-structured approach like CoG. P12, when asked to clarify his difficulties with the framework, stated that cloud environments present unique problems for defenders: we care about “the center keep of your castle, well there’s this other castle somewhere out there, we don’t know where, [and it is] part of the CoG.” However, these two participants did successfully use threat modeling to discover critical vulnerabilities within their daily work that they had not previously considered.

Changes in self-efficacy. When comparing responses from the post-training survey to baseline responses, 10 participants reported a perceived increase in their ability to monitor critical assets, 17 reported an increase in their ability to identify threats, 16 reported an increase in their ability to mitigate threats, 15 participants reported an increase in their ability to respond to incidents. Respectively, aver-

ages increased by 8.8%, 19.3%, 29.8%, and 20.0%. Using the Wilcoxon signed-rank test [226], there were significant increases in participants' perceived ability to identify threats ($W=61.0$, $p=0.031$), mitigate threats ($W=47.0$, $p=0.010$), and respond to incidents ($W=59.0$, $p=0.027$).

3.3.3.2 Actual efficacy

The study measures the actual efficacy of threat modeling using several metrics: the accuracy of participants' output, task completion times, similarities between participants' identified CoGs, and the contents of their actionable defense plans.

Output accuracy. Simply completing CoG tasks is insufficient to demonstrate success; the resulting output must also be valid and meaningful. Thus, the accuracy of participants' results is assessed via an expert evaluation from two NYC3 senior leaders. Both of these leaders received in-person training on CoG and are uniquely qualified to assess the accuracy of the provided responses given their intimate knowledge of the NYC3 environment and cybersecurity expertise. The evaluators received an anonymized set of the study results and asked them to jointly qualify the accuracy of the identified centers of gravity, critical vulnerabilities, threat capabilities/requirements, and ideal defense plans using a 6-point Likert scale ranging from zero to five with zero being "extremely unlikely (UL)" and five being "extremely likely (EL)" (See App. A.1). Additionally, the leaders were asked to indicate whether each ADP was sufficiently detailed to implement. The survey included one fictitious

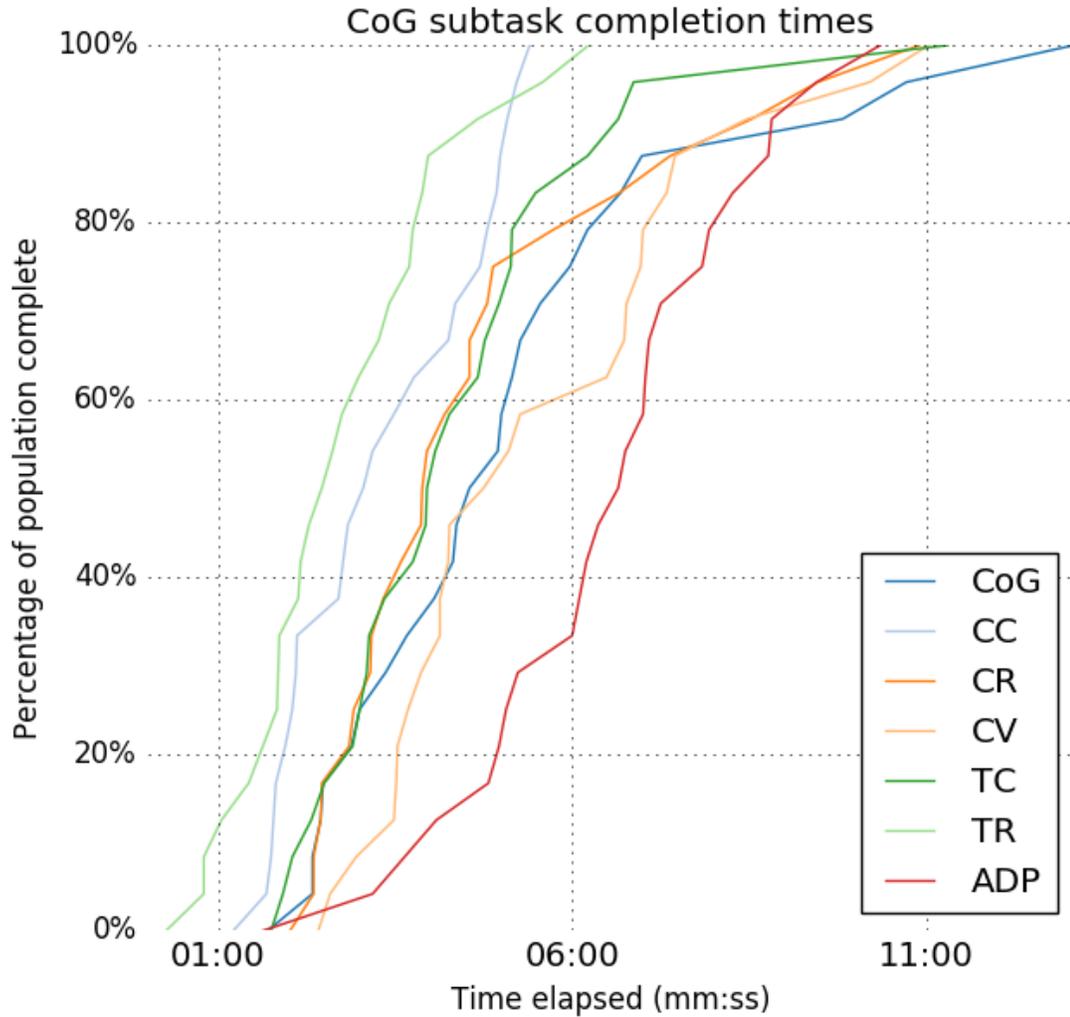


Figure 3.3: A cumulative distribution function (CDF) for participant subtask completion times.

participant entry as an attention check and validity control, which both panel members identified and rejected.

The panel concluded that: 22 of 25 identified centers of gravity were accurate with respect to a participant’s responsibilities (‘EL’=3, ‘Likely [L]’=9, ‘Somewhat likely [SL]’=10); all critical vulnerabilities were accurate for the identified centers of gravity (EL=6, L=7, SL=12); 23 of 25 threat capability and requirement profiles were accurate (EL=6, L=7, SL=10), and 24 of 25 actionable defense plans would

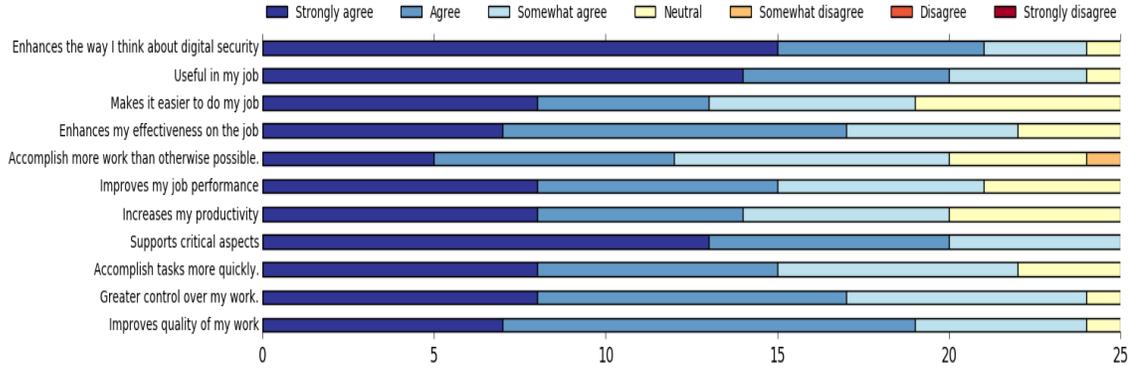


Figure 3.4: Perceived efficacy after using threat modeling for 30 days.

accurately address the identified threats (EL=5, L=11, SL=8).

A logistic regression, appropriate for ordinal Likert data, was used to estimate the effect of work roles, experience in IT, and educational background on the accuracy of the panel results. The regression included a mixed-model random effect [90] that groups results by work roles to account for correlation between individuals who fill similar positions. The initial model for the regression included each demographic category. To prevent overfitting, the regression tested all possible combinations of these inputs and selected the model with minimum Akaike Information Criterion [7]. The final selected model is given in Appendix B.1.5. The regression results show that no particular work role, amount of education, IT experience, or combination thereof enjoyed a statistically significant advantage when using threat modeling. These high success rates across the demographics support findings by Sindre and Opdahl that indicate threat modeling is a natural adaptation to standard IT practices [185].

Time requirements. Time required to apply CoG analysis helps to measure efficiency, which is a component of efficacy. On average, participants used the frame-

work and developed actionable defense plans in 36 minutes, 46 seconds ($\sigma = 9 : 01$). Figure 3.3 shows subtask completion times as a cumulative distribution function (CDF). Participants spent the greatest amount of time describing critical vulnerabilities and developing actionable defense plans, with these tasks averaging 5:27 and 6:25 respectively. Three out of five participants in a leadership role affirmed without prompting that threat modeling provided them with a tool for quickly framing difficult problems, with P24 stating “within an hour, [CoG] helped me think about some items, challenge some things, and re-surface some things, and that is very useful for me given my busy schedule.” P22 applied the framework in 22 minutes and commented during his closing performance evaluation session that he would “need much more time to fully develop” his ideas; however, he also said the session served as a catalyst for initiating a necessary dialogue for handling vulnerabilities.

CoG consistency. Analysis of the performance evaluation session results reveals that participants with similar work role classifications produced similar output. For example, 16 of 18 technicians indicated that a digital security tool was their CoG (e.g., firewalls, servers) whereas four of six participants in support roles identified a “soft” CoG (e.g., relationships, funding, and policies). Participants produced actionable defense plans averaging 5.9 mitigation strategies per plan and ranging from a minimum of three strategies to a maximum of 14.

Actionable defense plans. The contents of participants’ actionable defense plans help to further evaluate success. Participants identified real issues present within

their environment and developed means for reducing risk. Within the 25 actionable defense plans, participants cumulatively developed 147 mitigation strategies; detailed examples are described in Section 3.3.5. Participants indicated that 33% of the mitigation strategies they developed using threat modeling were new plans that would immediately improve the security posture of their environment if implemented. Additionally, participants stated that 31% of the mitigation strategies would improve upon existing NYC3 defensive measures and more adequately defend against identified threats. Participants felt that the remaining 36% of their described mitigation strategies were already sufficiently implemented across the NYC3 enterprise.

The NYC3 leadership panel indicated a majority of the actionable defense plans were sufficiently detailed for immediate implementation ('Yes'= 16). This shows that, even with limited framework exposure, many participants were able to develop sufficient action plans. To help illustrate, here is an ADP with insufficient detail using a security analyst's plan. After identifying his CoG as an Endpoint Detection and Response (EDR) system (Endpoint Detection and Response (EDR) describes a suite of tools focused on detecting and investigating suspicious activities, intrusions, and other problems on endpoint systems) and applying the framework, his ADP consisted of three mitigation strategies: "Make sure there is a fail-over setup and test it. Better change control. Better roll back procedures." While all of these address critical vulnerabilities, they provide no implementation details. In cases such as this, individuals require additional time to improve the fidelity of their responses or may benefit from expert assistance in transforming their ideas into fully

developed plans.

3.3.4 Observations after 30 days

After 30 days, participants still had a favorable opinion of threat modeling, most participants actually implemented defensive plans that they developed through the study, and that NYC3 institutionalized threat modeling within their routine practices.

3.3.4.1 Perceived efficacy

Thirty days after learning about CoG, there was a slight decrease in the perceived efficacy of the framework when compared to participant perceptions immediately after training: a 1.47% decrease for monitoring critical assets ($W=81.0$, $p=0.57$), 3.22% decrease for identifying threats ($W=131.0$, $p=0.83$), 3.58% decrease for mitigating threats ($W=94.0$, $p=0.18$), and 1.67% decrease for responding to incidents ($W=100.0$, $p=0.59$); none of these decreases were statistically significant. When comparing these 30-day metrics to the baseline, however, participants' perceived ability to monitor critical assets increased 7.4%, perceived ability to identify threats increased 16.1%, perceived ability to mitigate threats increased 26.3%, and perceived ability to respond to threats increased 18.3%. Participants' perceived ability to mitigate threats is a statistically significant increase from the baseline ($W=73.5$, $p=0.049$).

Figure 3.4 shows participants' evaluations of the efficacy of CoG analysis after

30 days. Overall, all participants agreed (“Strongly” = 13) that threat modeling supports critical aspects of their job. Additionally, 24 participants agreed (“Strongly” = 15) that threat modeling enhances the way they think about digital security. Despite the aforementioned decrease in perceived efficacy over the 30-day period, the number of participants who found the framework useful to their jobs increased from 23 to 24, as NYC3’s adoption of ADPs within their environment caused one participant to believe in the framework’s usefulness. Lastly, 245 of 275 responses to the 11 TAM questions indicated threat modeling is valuable for digital security.

3.3.4.2 Actual efficacy

Participants’ knowledge retention helps measure actual efficacy after 30 days. Measuring knowledge retention allows researchers to evaluate the longevity of organizational impacts from integrating the framework. After 30 days, participants averaged 78% accuracy on four comprehension questions. This is an increase from 69% immediately after learning the framework, suggesting threat modeling may become more memorable after additional applied experience. Each comprehension question required participants to pinpoint the best answer out of three viable responses; this allowed researchers to measure if participants understood critical relationships. In the 30-day follow-up, all participants accurately answered the critical vulnerability question, 23 correctly identified a CoG visually, 17 correctly identified a critical requirement for a capability, and 13 correctly identified a critical capability for a notional CoG.

3.3.4.3 Actual adoption

After 30 days, 21 participants reported that they implemented at least one mitigation strategy that they developed using threat modeling. In addition, 20 participants reported after 30 days that they integrated concepts from threat modeling within their daily work routines. For example, seven participants now use the framework for continually assessing risk; this is in contrast to the baseline results, where participants typically assessed risk only during audits and initial accreditation. Five participants stated that they now use threat modeling to prioritize their daily and mid-range efforts. Participants who did not adopt said they were too busy with urgent tasks (n=4) or needed more applied training (n=1).

NYC3 started to institutionalize threat modeling after participants had discussed their results with one another and realized the important implications of their findings. One week after completing their performance evaluation sessions, six participants transformed a wall within their primary meeting room into an “urgent priorities” board (Figure 3.5) for implementing defensive actions that address critical vulnerabilities identified during this study. Their board facilitates two-week action periods and improves how the organization communicates the impact of their progress to senior leaders. NYC3 leaders have since formalized this board using project management software and other practices such as “demo days” to demonstrate the viability of their defensive efforts.

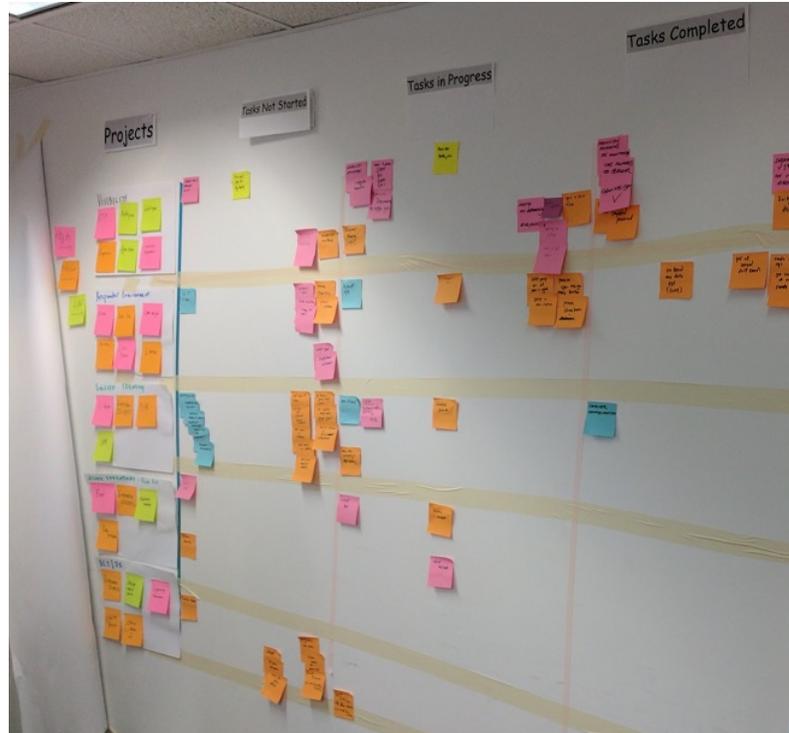


Figure 3.5: NYC3 developed an “urgent priorities” task tracker to address problems identified in this study.

3.3.5 Observations after 120 days

Observing NYC3’s environment 120 days after the study concluded allows researchers to understand the longer-term impact of threat modeling within live work environments. In total, NYC3 implemented eight new categories of controls directly based on the ADPs developed by participants in this study. Additionally, NYC3 provided researchers with access to server logs, their alert dashboard, and vulnerability reports to help measure the actual efficacy of three of these new controls.

3.3.5.1 Actual adoption

Below a sample set of ADPs that participants derived using threat modeling are detailed. NYC3 leaders monitored the implementation of these ADPs using their priorities board, and all mitigation strategies persist within the NYC environment 120 days after the study. Only provide high-level details about the ADPs are provided below to avoid placing NYC3 systems at risk.

Testing readiness. Nine participants cited resilient systems as critical requirements within their environment, and two identified untested disaster recovery plans as critical vulnerabilities. To dampen the impact of a cyber attack, natural disaster, or terrorist attack, they recommended frequently using multiple “fail-over” sites to validate functionality. Accordingly, NYC3 has begun testing fail-over servers within their local domain and plans to implement periodic, mandatory readiness tests across all NYC networks.

Securing accounts. Several participants identified user account permissions – a fundamental security control in any networked environment – as insufficiently well managed. Three participants stated that it is common for employees to migrate across the organization and retain permissions to data shares and assets they no longer need. NYC3 now directs monthly audits and re-certification of user access to narrow the impact of insider threats or stolen credentials. Seven participants recommended implementing multi-factor authentication. As a proof of concept, NYC3

implemented multi-factor authentication for 80 user accounts within a monitored subdomain.

Protecting physical network assets. Seven participants determined that if control measures restricting physical access to networking infrastructure were weak, it would create critical vulnerabilities. All expressed concern with insider threats causing damage or stealing data, but they all indicated that the most likely threat stems from accidental damage. Three participants discussed concerns with inadvertent, wide-scale power outages or power surges to networking infrastructure that could cause some issues to persist for an extended duration. These three participants recommended security escorts for all personnel, in addition to multi-factor access control near all networking infrastructure. Since the performance evaluation sessions, NYC3 has been working with federal, state, and private-sector entities on issues related to this topic.

Crowdsourcing assessments. Two participants reported that automated vulnerability assessment tools might not detect all vulnerabilities and that manual testing is needed for identifying more complex issues. Thus, P21 recommended that NYC establish a bug bounty program for public-facing services to benefit from the collective security community. Because of his recommendation, NYC3 partnered with a bug bounty service provider to conduct a 24-hour proof-of-concept assessment against one of its web services.

Sensor coverage. Ten participants acknowledged that the NYC environment is

far too vast for manual monitoring and that automated sensors play a critical role in defense. In this situation, a gap in sensor coverage can lead to unprotected systems or the successful exploitation of known vulnerabilities. Four participants recommended deploying additional EDRs on systems in specific subdomains within which NYC3 had limited visibility. Within 30 days after the threat modeling training, NYC3 technicians deployed 1331 new EDR sensors within these subdomains.

Protecting legacy systems. Three participants stated that legacy systems significantly impact their ability to secure systems; some were installed five decades ago and were never intended to be networked. Thus, they recommended segmenting non-critical legacy systems until they are replaced/upgraded. NYC3 is now working closely with partners to protect segmented systems and those that must remain online.

Protecting against data corruption. Participants P02 and P17 identified data corruption as risks to NYC3 systems. NYC3 technicians now verify the integrity of each software and indicator of compromise (IOC) update provided by third-party vendors to prevent the exploitation of update mechanisms, as seen in the 2017 NotPetya malware outbreak [182].

Reducing human error. Human error was another common theme across the threat landscape. Six participants stated that a simple typo in a configuration script, like the one that caused the 2017 Amazon S3 outage [9], could have significant impacts across multiple systems or networks. Three defenders recommended

two-person change control when updating configuration files on firewalls and EDR systems. Such controls require one person to propose a change and another to review and implement the change to reduce the likelihood of human error. NYC3 now enforces two-person change control on all modifications to access control lists.

3.3.5.2 Actual efficacy

Quantitative metrics captured in the 120 days after threat modeling training empirically support the efficacy of threat modeling. A NYC3 security analyst verified every intrusion, incident, and vulnerability within these data records. To protect the operational security of NYC3, specific threats that would enable a malicious actor to re-target their systems are not mentioned.

Securing accounts. User account logs allow researchers to analyze account hijacking attempts based on the geographic origin of attempts, time frequency between attempts, and why the attempt failed (e.g., wrong password or invalid token). Over 120 days, NYC3 recorded 3749 failed login attempts; based on frequency and subsequent successful logins, 3731 of these attempts with employees forgetting their password. Among the remaining failed logins, NYC3 successfully blocked hijacking attempts that originated from a foreign nation against seven *privileged* user accounts. Of these seven accounts, the attacker failed at the multi-factor login step for five accounts and failed due to password lockout on the other two accounts. Prior to this study, this subdomain did not have multi-factor verification enabled; these five privileged accounts were protected by mechanisms implemented solely because

of the introduction of threat modeling.

Crowdsourcing assessments. The 24-hour bug-bounty trial program yielded immediate results. Overall, 17 security researchers participated in the trial program and disclosed three previously unknown vulnerabilities in a public webserver protected by NYC3, verified through proof-of-concept examples. NYC3 validated these vulnerabilities and patched the production systems in accordance with policy and service-level objectives. After the success of this trial, NYC3 has authorized an enduring public program that will focus on improving the security posture of web applications under NYC3's purview. Such a program is a first for the City of New York and NYC3, created as a direct result of introducing threat modeling.

Sensor coverage. EDR reports allow researchers to uniquely identify which IOCs appeared in which systems, their severity level, and frequency of attempts. NYC3 deployed 1331 new sensors to endpoints that were previously unmonitored and were able to verify and respond to 541 unique intrusion attempts identified by these new sensors. Of these 541 intrusion attempts, 59 were labeled critical and 135 were labeled high severity; NYC3's partnered vendor security service manually validated each of these intrusions and verified their severity levels as true positives. One important aspect to note: if any systems had been infected prior to sensor deployment, the study would have captured both new intrusion attempts and any re-infection attempts that occurred after NYC3 deployed the sensors for the first time. According to the lead NYC3 EDR engineer, all 541 of these events could have led to successful

attacks or loss of system availability if technicians had not deployed the sensors to areas identified during threat modeling.

3.4 Discussion

This study provides the first structured evaluation of introducing threat modeling to a large-scale enterprise environment. Overall, the findings suggest that threat modeling, in this case the CoG framework, was an effective and efficient mechanism for developing actionable defense plans for the NYC3 enterprise. Defense plans created using CoG led to measurable, positive results. These results suggest that even a relatively small amount of focused threat modeling performed by IT personnel with no previous threat-modeling experience can quickly produce useful improvements.

Immediately after completing the performance evaluation sessions, 23 participants reported that they found the framework useful; after 30 days of use, 24 participants reported finding the framework useful and 20 participants reported regularly using concepts from threat modeling in their daily processes. In less than 37 minutes on average, the 25 participants developed 147 unique mitigation strategies for threats to their organization. NYC3 adopted many of these recommendations, improving their security posture in eight key areas. After 120 days, participant-designed ADPs blocked account hijackings of five privileged user accounts, blocked 541 unique intrusion attempts, and discovered (and remedied) three vulnerabilities in public-facing web servers, all of which support that introducing threat modeling

made NYC3 more secure.

Many of the ADPs that NYC3 employees developed and implemented (Section 3.3.5) contain straightforward recommendations, such as applying multi-factor authentication. This in itself constitutes an important finding: despite adhering to applicable federal, state, and local compliance standards and “best practices,” these measures were not already in use. Threat modeling offered the participants the agility to identify and implement defensive measures not (yet) prescribed in these standards. In this case, threat modeling helped the organization gain new perspective on their security gaps and proactively mitigate issues.

Many organizations are currently making significant investments in digital-security tools and capabilities [49]. The case study of threat modeling, in contrast, shows promising results that can be achieved by leveraging existing resources, without the need for new technologies or personnel. Further, the approach included only two hours of employee training, which is expected to be palatable for many organizations.

3.4.1 Lessons learned

Based on the case study, there are several observations available for the process of adopting threat modeling in a large organization.

Hands-on learning. Participants indicated that the hands-on approach to teaching threat modeling worked well. After the performance evaluation sessions, without prompting, 24 of 25 participants said that the personalized, hands-on application

allowed them to understand the framework better than the educational intervention classes alone. The logistic regression analysis on participants' CoG accuracy revealed a relatively level understanding of the framework across educational backgrounds, experience levels, and work roles. This suggests that many different practitioners can potentially benefit from this hands-on approach, supporting findings from Kolb & Kolb [118] and Bandura [21].

Mentoring and peer partnering. Multiple participants mentioned a desire for social and organizational support to facilitate the adoption of threat modeling. In their 30-day follow-up surveys, P18 and P24 stated that NYC3 would need organizational programs in place to aid wide-scale adoption of threat modeling, such as pairing junior personnel with mentors and facilitating peer-to-peer partnerships. During their performance evaluation sessions, P09 and P19 both mentioned that threat modeling would also be useful for integrating new personnel into NYC3. It is hypothesized that pairing experienced employees with junior personnel could permit mentors to orient their mentee to the environment and provide context to ongoing defensive initiatives, all while reinforcing their own understanding of threat modeling.

Further, the NYC3 leadership panel results indicated that 9 of 25 actionable defense plans were insufficiently detailed for immediate implementation. Peering would allow small teams to challenge one another and elicit details until results are adequately robust. This accords with prior studies of threat-modeling techniques, as well as peer partnering examples from other domains, that demonstrate the benefits of peer collaboration [47, 57, 61, 77, 85, 86, 103, 128, 131, 134, 136, 146, 154, 177].

Communication with leadership. After threat-modeling training, participants reported that they were better able to communicate the importance of various threats to NYC3 leadership. This was reflected in the immediate deployment of mitigation strategies, as discussed in Section 3.3.5. It is hypothesized that use of a single threat modeling framework — in this case CoG — across administrative boundaries may help to facilitate a shared language within the organization for communicating about threats. It would be particularly interesting to explicitly evaluate whether training executive-level leadership along with on-the-ground practitioners might yield useful communication benefits.

Shortcomings. Knowledge retention results show that participants struggled with framework-specific terminology; only 17 of 25 participants correctly identified critical requirements after 30 days. When institutionalizing threat modeling, it may be helpful to provide learners with quick-reference guides containing relatable examples to help clarify essential terminology. Additional emphasis on critical requirements within educational intervention training may help because critical vulnerabilities, threat capabilities, threat requirements, and actionable defense plans are all derived using critical requirements. Hands-on learning exercises using with the Eikmeier “does/uses” litmus test [64] may improve this shortcoming and participants’ overall CoG analysis understanding.

3.4.2 Future work

Further work is needed to complement and validate study findings. In this work, I took advantage of a unique cooperative opportunity to evaluate the introduction of an exemplar threat-modeling approach into an enterprise environment. In future work, comparative evaluation — ideally also in real-world environments — is necessary to understand the relative effectiveness of different threat-modeling approaches and may also help to clarify in what situations and environments different threat-modeling approaches are likely to be most effective.

To this end, threat modeling should be tested in multiple environments, to understand when and why these frameworks should be applied. Future evaluations may be able to consider how organization size, experience level and typical workload of staff members, organizational culture, and existing threat-modeling and/or security-analysis processes affect the efficacy of threat modeling. Future work should also explore less tangible organizational characteristics, such as employees' understanding of organizational objectives, hierarchical structure, lines of communication within and across groups, and the empowerment given to mid-level leaders.

In summary, study results indicate that introducing threat modeling — in this case, CoG — was useful for helping a large enterprise organization utilize existing resources more effectively to mitigate security threats. These findings underscore the importance of future evaluations exploring when and why this result generalizes to other real-world environments.

Chapter 4: Baseline Security: Security Implications of Policies, Laws, and Regulations

In this chapter¹ I present research that focuses on the implications of baseline compliance security programs on organizations. Proactive baselining is intended to ensure digital systems are operating at a minimum level of security to defeat common threats and ensures administrators strictly adhere to practices that sustain security. In many cases, proactive baselines are mandated by various laws, policies, or regulations through compliance programs [75, 108, 159, 169]. These compliance programs are typically one-size-fits-all in applicability and scope, meaning organizations are left to independently assess which threats they are still vulnerable to and how to prepare for likely security incidents.

Generally, digital security compliance programs and policies serve as powerful tools for protecting organizations' intellectual property, sensitive resources, customers, and employees through mandated security controls. Organizations place a significant emphasis on compliance and often conflate high compliance audit scores with strong security; however, no compliance standard has been systemically evaluated for security concerns that may exist even within fully-compliant organizations.

¹Published as [192, 193]

Here, I describe the novel approach for auditing exemplar compliance standards that affect nearly every person within the United States: standards for federal tax information, credit card transactions, the electric grid, and cloud-based services for the federal government. Partnered organizations help validate these findings within enterprise environments and provide first-hand narratives describing impact.

This study reveals that when compliance standards are used literally as checklists — a common occurrence, as confirmed by compliance experts — their technical controls and processes are not always sufficient. Security concerns can exist even with perfect compliance. Hundreds of issues with varying severity across these standards exist. Additionally, no clearly-defined process exists for reporting security concerns associated with compliance standards. Overall, results suggest that auditing compliance standards can provide valuable benefits to the security posture of compliant organizations.

4.1 Method

In the first step of this study, the researchers comprehensively audited three compliance standards to identify potential security concerns. To validate these concerns, the study leverages four experts to provide their assessment of the findings. Through quantitative and qualitative analysis on expert responses, the results identify discrepancies and also derive additional context for applicability within enterprise environments.

This study occurred from October 2017 through September 2018 and was ruled

not human subjects research by the ethics-compliance office. Due to the sensitive nature of unmitigated data vulnerabilities within real environments, many findings are generalized to protect networks and systems.

ID	Employment	Role	Org Size	IT Exp (yrs)	Edu.	Docs
R1	A, G	M, R	500	18	MS	I,P,N
R2	G	M, R	10k+	16	PhD	I,P
R3	A, G*, I	M, R	100	20	BS	I,N
R4	I	M, R	35	15	JD	I,P
R5	A, G*, I	M, D	100	8	BS	I,N
R6	G	M, D	100	5	BS	I,N
E1	G, I	M	150	10	BS	I
E2	G	M	150	15	MS	I
E3	G*, I	M, D	1k	18	MS	P
E4	A, G*, I	R	5k	20	MS	N

Table 4.1: Researcher and Expert Demographics. The columns show: participant identifiers (coded by interview date order), place of employment, work role, employer’s organization size, years of IT experience, education, and which documents they evaluated. The abbreviations A/G/I/* in the employment column stand for academia, government, industry, and previous experience in the indicated sector respectively. The abbreviations M/R/D in the work role column stand for management, research, and development. The abbreviations in the education column stand for bachelors degree, masters degree, and doctorate (PhD/JD). The abbreviations I/P/N in the documents column stand for IRS P1075, PCI DSS, and NERC CIP.

4.1.1 Compliance-standard audit

The team of six researchers designed the audit to systematically evaluate three unrelated compliance standards in a repeatable manner. Each researcher audited a subset of the standards, with at least three researchers per standard (as shown in Table 4.1). The objective was to identify issues that might negatively affect digital security, including policies that expose sensitive information and processes that create issues due to ambiguous implementation guidance.

- ⓘ 9.3.1.9 Session Lock (AC-11) insufficient_process irs Probability_Likely Severity_Moderate

- ⓘ 9.3.1.2 Account Management (AC-2) - shared accounts data_vulnerability irs Probability_Frequent Severity_Moderate

- ⓘ 12.1 General data_vulnerability irs Probability_Unlikely Severity_Moderate

- ⓘ 9.4.1 Cloud Computing Environments irs Probability_Occasional Severity_Negligible verbiage_issue

- ⓘ 9.3.16.5 Boundary Protection irs Probability_Likely Severity_Negligible unenforceable

- ⓘ 9.4.18 Wireless Networks data_vulnerability irs Probability_Likely Severity_Critical

- ⓘ 9.3.17.7 Information Input Validation (SI-10) irs Probability_Likely Severity_Negligible verbiage_issue

Figure 4.1: Example of recorded security concerns.

To support asynchronous collaboration, each researcher independently logged their findings (Figure 4.1). This repository maintained a running list of codified definitions and instructions for researchers to reference throughout the study.

All six researchers conducted a complete audit of IRS Publication 1075, following a content-analysis process drawn from social-science research. Each researcher independently examined each line of the standard. At each of several predetermined milestones within the document (e.g., the end of a section), the researcher would log their findings, including the section title where the issue was found, the exact phrase deemed problematic, a short description of the perceived issue, and references to related, publicly-known issues. If a researcher found multiple issues within one phrase or section, they logged each separately and each issue was given a unique identification number (this assisted greatly in performing post-collection analysis). For every logged issue, all other researchers would indicate (1) if they

found the same issue independently and (2) whether they concurred with the finding. If there was not unanimous consensus on an issue, the issue was discarded it but maintained a record of the disagreement (used to calculate inter-rater reliability). Issues without unanimous agreement were discarded instead of resolving disagreement due to time constraints; future studies using this methodology may choose to mediate disagreements within the group or use an external expert for conducting tie-breakers. Mediation may provide additional real-world discoveries that would have been discarded otherwise.

After each researcher logged all of their independently-discovered security concerns, researchers then calculated the inter-coder reliability — a measure of consistency among independent auditors — for IRS P1075. Researchers used Krippendorff’s Alpha (α) to account for chance agreements [88]. To do this, researchers normalized all data for ingest into ReCal3, an online inter-rater reliability calculator. Each individual compliance control is considered an independent item that researchers could agree (or disagree) upon. For example, each individual compliance control fell into one of three different categories: (1) all researchers identified and agreed that the control contains security concerns, (2) all researchers agreed that the control did not contain security concerns, or (3) there is a disagreement whether the control contains security concerns.

The first step of normalizing the data for inter-rated reliability was converting IRS P1075 into a binary matrix, listing each technical control in the document as a row. Columns in this matrix indicated agreement levels from each of the six researchers for corresponding control listed in the row (‘1’ indicates research believes

	R1	R2	R3	R4	R5	R6
Control1	1	1	1	1	1	1
Control2	0	1	0	0	0	1
..						
ControlN	0	0	0	0	0	0

Table 4.2: Example matrix used for calculating inter-rater reliability in compliance controls. The columns show: researcher identifiers and whether the researcher independently assessed a security concern in the security control.

a security concern is present, ‘0’ indicates a researcher believes a security concern is not present). Using the example shown in Table 4.2, there would be unanimous agreement that ‘Control1’ has a security concern but a disagreement for ‘Control2.’ Results yielded a reliability of $\alpha = 0.815$ for P1075; an α value above 0.8 indicates high reliability [120, 126]. Having developed a reliable auditing process, subgroups helped parallelize the remaining effort. Four researchers audited NERC CIP 007-6 and three researchers audited PCI DSS. One researcher (R1) audited all three guidelines to provide continuity. The subgroups attained $\alpha = 0.801$ and 0.797 respectively.

Researchers further analyzed the identified issues using iterative open coding, a process for creating and applying category labels (known as a *codebook*) to data [199]. In particular, the researchers who audited each standard coded each identified issue in that standard for perceived root cause, probability of occurrence, and severity. Researchers resolved all disagreements among coders and developed a stable codebook by establishing a unanimously agreed-upon definition for coded terms, adapting many terms from the Composite Risk Management (CRM) framework [216] and the Information System Risk-based Assessment framework [65]. After

any revisions to these definitions, researchers re-coded previously coded items, repeating this process until researchers coded all responses, resolved all disagreements, and the codebook was stable.

The final codebook described four root causes for security concerns. A *data vulnerability* is an issue that will result in a data breach or compromise of sensitive information. An *unenforceable security control* cannot be enforced as written; these controls should be reworded or removed from the compliance standard. An *under-defined process* is an issue explicitly missing instructions or details that are required for a secure implementation, resulting in security gaps. An *ambiguous specification*, in contrast, is vague or ambiguous about some implementation details, such that different readers could interpret it differently. Some interpretations could potentially result in either an inappropriate action or inaction. Throughout Sections 4.2.2, 4.3.2, and 4.4.2, audit findings are detailed using these root causes.

The following terms and definitions are used for probability: *frequent* occurs often and is continuously experienced; *likely* occurs several times; *occasional* occurs sporadically; *seldom* is unlikely, but could occur at some time; and *unlikely* is assumed it will not occur. The following terms for are used for severity: *catastrophic* results in complete system loss, full data breach, or the corruption of all data; *critical* results in major system damage, significant data breach, or corruption of sensitive data; *moderate* results in minor system damage or partial data breach; and *negligible* results in minor system impairment. Using a risk assessment matrix adopted from the CRM framework (Figure 4.2), one can then calculate each issue's risk level — a function of probability and severity — as extremely high, high, moderate, or

		Probability				
		Unlikely	Seldom	Occasional	Likely	Frequent
Severity	Catastrophic	M	H	H	E	E
	Critical	L	M	H	H	E
	Moderate	L	L	M	M	H
	Negligible	L	L	L	L	M
		E - Extremely High	H - High	M - Moderate	L - Low	

Figure 4.2: Security concern risk levels. Levels were assigned based on a Composite Risk Management risk-assessment matrix that includes both probability of occurrence and impact severity.

low [216].

Best practices suggest that empirical research should be conducted by personnel with extensive domain knowledge [171]. Accordingly, the auditing researchers possess an average of 14.3 years of digital security experience within academia, the federal government, and industry. Each researcher grounded their audit findings in their past digital security experiences. Additional information about the data set is in Appendix B.2.2.

4.1.2 Expert validation process

To obtain external validation of the findings, I established partnerships with real-world organizations and compliance subject-matter experts to confirm or reject the findings. Experts were asked to classify the identified issues in one of three categories: confirmed, plausible, or rejected. A confirmed issue indicates that the expert has previously observed security concerns associated with the issue or that observable consequences from the issue actively exist within an enterprise environment. A

plausible issue occurs when the expert has not personally observed security concerns related to the issue but agrees such security concerns could manifest within other organizations. A rejected finding indicates that there is no observable evidence of security concerns related to the issue within a live environment, or that there are related security factors not considered.

Each expert was asked a series of closed- and open-ended survey questions to elicit information (detailed in Appendix A.2). In addition to directly validating or rejecting each issue, the experts were asked to provide additional context from their personal experience. The experts received the issues in a randomized order, and received only the referenced section title, exact text from the section, and a short narrative describing the perceived issue.

After collecting data from each expert and removing rejected findings, the Wilcoxon signed-rank test is used to compare researchers' assessment of probability and severity with the experts' responses for PCI DSS and NERC CIP 007-6; the Friedman test (omnibus) with planned pairwise Wilcoxon signed-rank tests helps compare IRS P1075 responses, which had two expert validators [51, 226]. Open-ended discussions with the experts were used to discuss similarities and differences in assessments.

Partner criteria. The following criteria was used for partnering with organizations:

(1) the organization must regularly be subjected to audits, must regularly audit other organizations, or must contribute to the content of the relevant compliance standard, (2) the provided validators must have at least two years of experience with

the relevant standard, and (3) the organization must be able to mediate responsible disclosure of the findings.

After months of negotiation, researchers established memorandums of understanding with three organizations that met the criteria. Leaders within each organization nominated several compliance experts; each candidate received an email outlining the voluntary nature of the study as well as the motivation and goals. Table 4.1 shows the qualifications of the four volunteer experts. Experts completed their surveys during regularly scheduled work hours and did not receive any additional monetary incentives for participating.

Issue selection. An essential tenet for partnering with experts is minimizing disruption to their daily responsibilities. Research suggests that the quality of survey responses decreases over time, and excessive time away from work may result in an expert terminating their participation in the study [100]. To this end, surveys were designed for experts to complete within 60-90 minutes of focused effort; actual completion time averaged 84.8 minutes. Given the limited pool of experts, this required the selection of only a subset of the findings to validate; as described in detail below, the issues were selected to validate semi-randomly, while prioritizing the extremely-high-risk and high-risk issues.

Pilot. Prior to deploying the protocol with partnering organizations, surveys were piloted to pre-test relevance and clarity with security practitioners familiar with auditing and compliance standards. The study protocol was updated based on

pilot feedback. The finalized questionnaire in Appendix was created after two iterations [A.2](#). The two pilot experts currently conduct digital-security penetration testing against organizations, providing technical remediation recommendations for discovered security concerns.

4.1.3 Limitations

The recruitment letter and consent waiver explained the purpose of the study. Thus, there may be self-selection bias in which personnel most interested in the study were more likely to anonymously participate. However, this may also suggest that the experts were prepared to think more critically about reported issues.

All of the experts work directly in compliance and their intimate working knowledge with compliance standards reduces the possibility of demand characteristics — a condition in which participants unconsciously change their behavior to perform well within a study [167]. The study questions the validity of the compliance standards that serve as the basis for the experts' employment. This suggests that the experts would be in many cases predisposed to underestimate problems within these standards. Additionally, the validation method does not elicit expert feedback for false negatives – issues that the original analysis may not have detected. As such, the expert responses provide a lower bound for validity.

The partnered organizations have similar structures, missions, and technologies to other organizations that adhere to the selected compliance standards; however, there may exist specific organizational characteristics that affect their specific

implementations or inhibit generalizability. As such, validating the presence of the discovered security concerns within partnered organizations' environments does not mean that all organizations adhering to similar compliance standards have security concerns, and the rejection of one of the findings does not imply that another organization elsewhere does not have security concerns. Nonetheless, the results can indicate systemic issues that organizations need to account for when assessing their levels of digital security risk and provide novel insights into the impact of compliance standards on digital security in enterprise environments.

Lastly, each compliance was audited standard without considering other security controls in complementary documents. For this study, it is assumed that organizations implement compliance standards perfectly and limit the scope to finding security concerns in the documents as written.

4.2 Results: IRS P1075

4.2.1 Overview

IRS Publication 1075 provides mechanisms for protecting and controlling access to federal tax information. IRS P1075 applies to all U.S. federal, state, and local agencies that receive Federal Tax Information (FTI) from the IRS or from secondary sources like the Social Security Administration [108]. Of the three standards assessed, IRS P1075 is the longest standing, dating back to 1996 [107]. This study audited the 2016 revision, which was the most current version available at the time.

FTI security potentially affects every federal taxpayer. Organizations such

Risk Estimates for IRS Compliance Standard

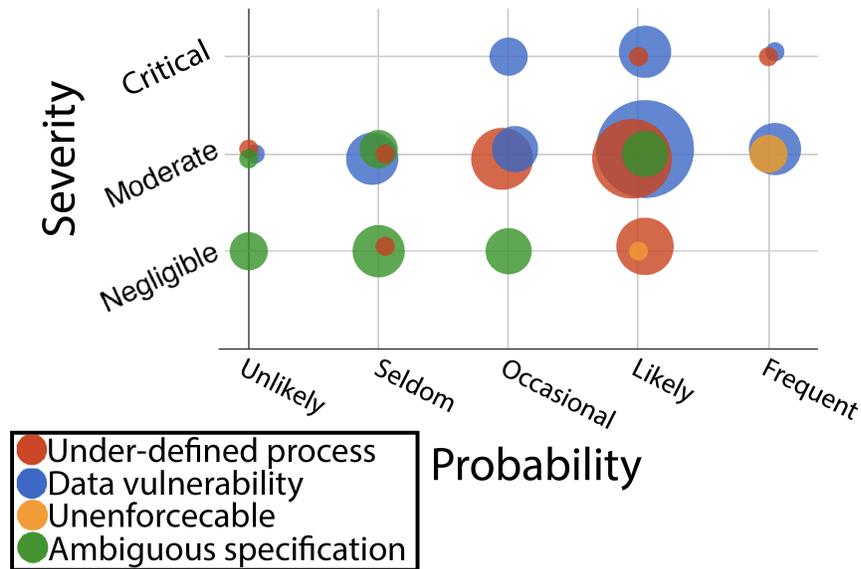


Figure 4.3: Distribution of security concerns identified for IRS P1075. Color indicates the type of security concern; each dot indicates by size how many security concerns were identified with a given type, severity, and probability. Data vulnerabilities were most common (n=37).

as the Office of Child Support Enforcement from the U.S. Department of Health and Human Services rely upon IRS P1075 for securing the networked infrastructure of child support financial records [215]. Companies such as Amazon offer cloud infrastructure services that are fully compliant with P1075, marketing their virtual private server services to customers who need a “turn-key” solution for systems that transmit or receive FTI [10].

P1075 is written for information technology security professionals responsible for securing FTI. Key provisions include definitions for terms, parties authorized to access FTI, record-keeping requirements, physical controls for securing data, technical controls for secure storage/transmission, inspection protocols, and sanctions for non-compliance. The IRS Office of Safeguards coordinates and conducts compli-

ance audits of entities possessing FTI. Of the three standards assessed, P1075 has the weakest sanctions. There are no provisions for the issuance of fines for insecure practices, and the strictest sanction available to inspectors is data revocation after failure to adhere to a prescribed corrective action plan. However, non-compliant organizations can apply for data revocation waivers that extend their access to FTI for six months; according to the standard as written, this process can continue indefinitely despite continued non-compliance. This process has the potential to minimize the impact of sanctions while allowing insecure practices to persist. Overall, IRS P1075 was qualitatively and quantitatively the weakest of three documents assessed during this study.

4.2.2 Findings

The audit of P1075 identified a total of 81 independent issues across 309 individual security controls (Figure 4.3). Of these, two issues presented an “Extremely High” risk, whereas 13 were “High,” 32 were “Moderate” and 34 were “Low” risk according to the Risk Assessment Matrix (Figure 4.2). In addition, 15 initially identified issues were discarded, including 11 discarded when researchers found implementation details that were clarified in later sections of the standard and four resulting from researcher disagreements. All four issue disagreements related to nuanced interpretations of ambiguous portions of the standard.

Security concern trends. There are five issues involving portable devices (e.g., mobile phones and laptops) and seven involving cloud-based data storage solutions.

The prevalence of these issues are associated with shifts toward bring-your-own-device regimes and an increased reliance on cloud-storage solutions over on-premises servers [76]. These emerging solutions require specialized security measures and create inconsistencies with the best security practices that professionals have developed over the past few decades [201].

Of the 81 issues identified within P1075, Section 9 had 40 technical controls with security concerns. Of note, Section 9 has several obsolete controls such as password expiration period requirements (which is shown to encourage insecure practices such as writing newly rotated passwords near user workstations) [81,205]. In this particular instance, the *IRS mandated organizations to make a worse security decision than the decision they might have made in the absence of P1075*. Below are detailed examples of findings based on their associated root cause.

Data vulnerability. There are 37 issues that would establish conditions for a data breach if controls and processes are implemented as described in the publication. One example in Section 9.3.6.8 outlines processes for restoring backups once a compromise in a live system has occurred. As written, P1075 does not require technicians to verify the integrity of backups before restoration, meaning that technicians could revert to a state that an attacker has already infected (giving them persistent access) or revert to a vulnerable state that an attacker could re-exploit, reestablishing access to sensitive data [173]. Real-world trends stemming from ransomware support the urgency of backup integrity checks [180]. This high-risk issue has a likely probability and moderate severity.

Section 9.3.5.11 includes provisions for user-installed applications. Environments that store or transmit FTI should be highly secure and should only be used for FTI — other functions and services should occur outside the FTI environment. As such, there should be little to no need for user-installed software, especially given that users are one of the primary vectors for introducing malware into environments [3]. Section 9.3 should instead mandate application whitelisting for installation attempts, limiting the subset of authorized applications that anyone can install on the system. A more stringent recommendation would include revoking user-installation privileges altogether, forcing trusted system administrators to establish a safe baseline of applications allowed to interact with FTI. This high-risk issue has a likely probability and critical severity that can place FTI at risk.

There is an extremely-high-risk issue within Section 1.4.4 “Information Received From Taxpayers or Third Parties,” which limits the responsibility for securing FTI. According to this section, the IRS is only responsible for securing data originating from the IRS as FTI, excluding data received from customers like federal tax returns. Additionally, this section includes provisions for removing FTI protections on data if an entity replaces IRS-sourced FTI with the exact same information sourced from another party. This is analogous to eliminating protections for top-secret government data simply because the same information can be bought on the black market. This mandated behavior allows organizations to bypass security measures and remove protections on the data P1075 is meant to safeguard. P1075 should enforce protection for all FTI, regardless of source.

Section 1.4.3 defines certain data as personally identifiable information (PII)

but does not protect the names of individuals associated with the person filing the return – such as a spouse or dependent. This high-risk issue may allow an attacker, for example, to develop a targeted spearphishing campaign against an individual. The definition of PII should expand to include sensitive information about all persons listed.

Unenforceable controls. There are three controls that are unenforceable. For example, Section 4.7 provides several measures for secure telework access to FTI. P1075 provides many requirements for physical data protections, such as badge-based control and on-premises guards; these are infeasible in the case of telework, as most personnel with FTI access at their private homes cannot abide by these types of controls. Additionally, IRS inspections of private residences for physical security compliance seems fraught with complications. Either the IRS should ban residential-based telework programs until it can verify that all locations with FTI access are compliant with physical security requirements, or that the standard acknowledge that these physical controls are not actually required. This high-risk issue has a frequent probability and moderate severity.

Under-defined process. There are 27 issues that reflect processes that are not sufficiently detailed for a secure implementation. One such issue within Section 8.3 states that “every third piece of physical electronic media must be checked to ensure appropriate destruction of FTI.” Given the disparate possible sizes of electronic media, this particular section should recommend accounting for logical storage size

of the media instead of its physical instantiation. This would ensure that media with larger storage volumes are highly prioritized for destruction validation. This issue has a moderate-severity, moderate-risk issue with a likely probability.

One low-risk issue occurs in Section 1.4.7, which limits human access to FTI based on “need to know” but does not consider machines or systems with a “need to access” data. Administrators must limit system access to FTI to prevent unauthorized access or manipulation of data, especially for systems performing aggregate analysis that may inadvertently disclose sensitive information.

Section 9.3.13.3 covers background checks for personnel with access to FTI. The researchers assessed that this section could create information gaps at the federal, state, and local levels. For example, information about an individual who mishandled sensitive data at a previous job may never have entered federal databases. These extremely-high-risk information gaps increase likelihood for insider threats and risks to data, and highlight the need for aggregating multiple sources of data for thorough background checks.

Section 9.3.5.8 has an issue, which outlines a procedure for establishing an Information System component inventory (i.e., a listing of authorized devices that operate within an organization). As written, this procedure does not require the inventory process to be tied to a “ground truth,” meaning there is no comparison of which devices should be operating within an organization with which devices actually are. This is dangerous, as it could permit a rogue system to persist on a network or even be inventoried as a legitimate system. Providing a rogue system with legitimate access within a sensitive environment obviates the need for an attacker

to deploy malware within the environment and reduces the likelihood that any defensive sensors would ever detect anomalous activity from the attacker. This moderate-risk issue has an occasional probability and moderate severity. Industry recommendations integrate asset inventory with supply acquisition, ensuring that only company-purchased, legitimate systems are on the network [25].

Ambiguous specification. There are 14 issues involving insufficient details that create ambiguity or uncertainty throughout P1075. P1075 uses vague terms such as “significant change” throughout, without ever defining thresholds that auditors deem to be significant. As an example, Section 9.3 outlines “Access Control Policy and Procedures” that must be reviewed (by whom?) every three years or whenever a significant change occurs. This subjectivity allows reviewers to deem any or all changes insignificant to circumvent a change review. Additionally, the document’s use of passive voice clouds the responsibility for conducting the review — ambiguous controls can create security gaps through inaction. Each mandate should use active voice and assign a responsible individual (e.g., an office manager or system administrator) for each requirement. As presently written, an individual who works in an organization’s talent recruiting department with no security training would be a sufficient reviewer for access-control policy. These moderate-risk issues have a likely probability and moderate severity.

4.2.3 Expert validation

For assessing the validity of the IRS P1075 audit, the findings were exported to New York City Cyber Command (NYC3) for evaluation. NYC3 is a city government organization that oversees the cybersecurity of 143 separate departments and agencies as well as more than 300,000 people. In addition to defending NYC against cybersecurity threats, NYC3 is responsible for ensuring compliance with government-mandated policies. In particular, the NYC3 team includes five full-time employees and three consultants who focus solely on security compliance. Each of the 143 departments within the city government also has an internal, full-time compliance teams.

IRS P1075 applies to the vast majority of these 143 entities. NYC3 advises other NYC entities on P1075 compliance and is also subjected to IRS audits. Two NYC3 governance and compliance officials assessed the validity of the findings with respect to a particular subdomain under NYC3's purview that must comply with P1075 standards. This subdomain consists of a controlled internal network that contains FTI and supports approximately 150 users. NYC3's last formal P1075 audit was in January and February 2018, where three on-site auditors used the standard as a line-by-line checklist to assess NYC3's compliance. Of note, preparation for this inspection consumed the compliance team as well as several technicians for approximately four months prior to their inspection date.

Because of their limited time availability, these two NYC3 compliance officials (hereafter referenced as Experts E1 and E2) assessed 20 issues (25% of the total 81

issues). In order to cover issues at all risk levels but prioritize significant concerns, the evaluation included both extremely-high-risk issues, and then randomly sampled 10 of the 13 high-risk issues, four of the 32 moderate issues, and four of the 34 low-risk issues.

When validating P1075 issues, E1 and E2 were able to directly examine their network for the presence of security concerns caused by issues identified by the researchers, in a kind of white-box penetration test [74]. This was possible because, unlike E3 and E4, E1 and E2 are officials with administrator privileges within the audited subdomain. The two experts analyzed the findings independently and did not discuss their findings with one another during the study. Overall, these experts confirmed 17 of the findings, rejected two issues, and indicated one issue could be plausible within their own or another environment.

When comparing the risk estimates to those of E1 and E2, there was found no statistical difference between severity estimates (omnibus $p = 0.54$), but the researchers assessed issues to be statistically more likely with medium effect ($p = 0.0001, 0.034, < 0.0001$; $r = 0.485, 0.336, 0.533$ for omnibus and then pairwise researchers vs. E1, E2 respectively). E1 indicated that his knowledge of current and on-going initiatives most likely biased his responses, making it hard for him to follow instructions to consider each issue only “if standard is followed as written and nothing else” (as written in Appendix A.2). This response supports the notion in Section 4.1.2 that participant-validated responses represent a lower-bound for this study.

The issue that E1 and E2 classified as plausible rather than confirmed comes

from Section 1.4.7 “Need to Know.” E2 indicated that NYC3 data scientists incorporate the principle of least privilege for systems, service accounts, and user accounts, which would prevent unauthorized access and manipulation of FTI. E1 added that NYC3’s PKI infrastructure assists with controlling access to “need to access” data. Both participants indicated they were unsure if this security concern was ever present within NYC3, but that it could be present within other organizations.

E1 and E2 rejected the finding for Section 1.4.3 PII, indicating NYC3 always encrypts entire tax records while in transit and rest, and that this is standard practice for organizations with access to FTI. Thus, associated individuals’ PII are always protected, invalidating the finding. However, because this is not codified within P1075, there is no certainty that other organizations adhere to this “standard practice.”

NYC3 also rejected the finding associated with Section 9.3 background checks. E1 indicated that it is standard practice to aggregate personnel records from the locations an individual has lived or worked to determine if the individual should have access to sensitive information, thus rejecting the finding. Because P1075 does not mandate data sources or how far back in history to consider, there is no certainty that other organizations conduct this practice.

Additional defenses. E1 and E2 identified several controls pervasive throughout NYC3 that help reduce or eliminate the impact of many of the researcher findings. Of note, NYC3 requires a Change Control Board (CCB) that E2 believes “is an essential risk-mitigating factor” for addressing many of the confirmed P1075 security

concerns, such as Section 9.3.5.11 “User-Installed Software.” The CCB evaluates all user requests for system modifications and holistically considers the change’s impact to security. If the CCB approves the change, it authorizes a trusted administrator to conduct the software installation and adds the change to the system’s secure baseline. Additionally, NYC3 incorporates a real-time, automatic asset manager which alerts their Security Operations Center any time a new device is added to their networks. This defensive strategy eliminates the security concern identified in Section 9.3.5.8 “Information System Component Inventory.”

It is important to note that these defenses employed at NYC3 exceed the baseline security standards required by P1075 and mitigate issues that P1075 either fails to account for even causes. One cannot assume that all organizations will recognize the need for these additional mitigations and be willing to invest in them.

4.3 Results: PCI DSS

4.3.1 Overview

The Payment Card Industry Data Security Standard (PCI DSS) applies to all entities that store, process, and transmit credit-card-holder data for major branded credit cards [169]. Guidance in this standard includes building and maintaining a secure network and systems, protecting cardholder data, and monitoring/testing networks. PCI DSS v1.0 dates back to 2004 as a program led by Visa; the PCI Security Standards Council (SSC) was formed in 2006 by American Express, Discover, JCB International, MasterCard and Visa to enhance PCI DSS [169]. Researchers

audited the 2016 v3.2; v4.0 was in development during this study.

PCI DSS affects every person within the United States who makes credit card purchases and every organization that accepts credit card payments. A U.S. Federal Reserve study showed that consumers spent \$5.98 trillion with credit cards in 2016, highlighting the importance of securing the systems that support those financial transactions [214]. PCI DSS authors designed the document to be accessible to assessors and the technicians charged with implementing the technical controls.

Qualified Security Assessors perform PCI DSS audits after attaining the appropriate inspection certifications. According to one such assessor (not an author or an expert validator), audit frequency varies for merchants and service providers depending on their number of supported annual transactions [149]. On-site audit teams vary from one to three personnel per inspection; these personnel focus full-time on auditing the PCI DSS compliance of other organizations. The assessor indicates that penalties for non-compliance are common, but vary in size based on the severity of infraction and size of customer base. Monthly fines that can range from \$5,000 to \$100,000 and continue until compliance issues are resolved. If a data breach occurs as a result of non-compliance, companies may be responsible for consumer services (e.g., credit monitoring) or may have payment-processing privileges revoked.

Risk Estimates for PCI Compliance Standard

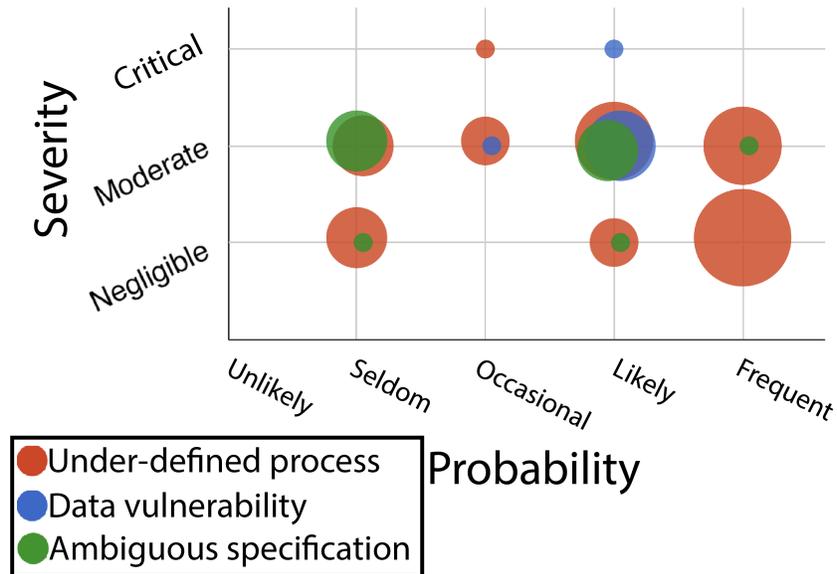


Figure 4.4: Distribution of security concerns identified for PCI DSS. Color indicates the type of security concern; each dot indicates by size how many security concerns were identified with a given type, severity, and probability. Under-defined processes were most common (n=29).

4.3.2 Findings

Within the 851 independent controls specified by PCI DSS, 46 security concerns were identified: eight high-risk, 22 moderate-risk, and 16 low-risk (as shown in Figure 4.4). Six other potential issues were discarded, all of which were under-defined processes that did not result in any insecure practices or conditions.

Security concern trends. There are four issues related to improperly identifying sensitive information. PCI DSS focuses heavily on protecting primary account numbers (PANs) that are tied to credit cards but fails to protect other information that could lead to PAN access, such as passwords or password-recovery information.

Additionally, there are 10 issues involving technical controls that lack timelines for required action. For each required action, the standard should specify either a fixed interval for repetition or for a triggering event with an ensuing deadline. Below are discovered PCI DSS issues, organized according to perceived root cause.

Data vulnerability. There are seven security concerns that could establish conditions for a data breach. One example of a high-risk vulnerability stems from Section 1.4, which includes mandates for securing personal computing systems within the cardholder data environment (CDE). Security professionals should disallow personal electronics within the CDE network segment; more broadly, all services and systems should be limited by “need to access” cardholder data. Personal devices and activities increase the likelihood of malware or other unauthorized access and generally are not necessary within CDE network segments [3]. This security concern has a likely probability and critical severity.

A tangentially-related moderate-risk security concern stems from the “Network Segmentation” section of PCI DSS, which scopes the standard’s safeguards to only the network segment that contains cardholder data. Effectively, this provision would allow an organization with no security controls outside of the CDE to pass an audit as long as the CDE itself is protected in accordance with PCI DSS specifications. Allowing vulnerable servers and systems within the same network as the CDE could provide attackers with a landing point into internal portions of the network and establish conditions for lateral movement into the CDE from adjacent network segments (through well-known attacks such as VLAN hopping). Due to the series

of security holes that must be present for such an attack to occur, the exploitation of this vulnerability should be seldom but critically severe for affected systems.

Another data vulnerability is present within the “PCI DSS Applicability Information” section, where PCI DSS defines sensitive authentication data. PCI DSS does not consider passwords to be sensitive authentication data and does not protect information an attacker could use to reset service passwords (e.g., email addresses, Social Security Numbers, and dates of birth). The social engineering attack against Naoki Hiroshima’s @N Twitter account leveraged similar pieces of information to access protected accounts [95]. Given that publicly-available articles detail how unprotected information can lead to unauthorized access, this security concern has a moderate severity and likely probability.

Under-defined process. There are 29 issues with process specifications that are insufficient for a secure implementation. Section 3.2.1 calls for assessors to select a sample of system components and examine data sources to detect cardholder data that is improperly stored. Sampling is an insufficient process, considering the simplicity of searching for cardholder data that adheres to a well-known format. Assessors should be mandated to use automated tools on all CDE systems to detect improperly stored cardholder data. Based on the moderate severity of exposed cardholder data and the frequent likelihood insufficient checks occurring, this issue has a high-level risk.

PCI DSS features two high-risk under-defined processes in “Requirement 5: Protect all systems against malware and regularly update anti-virus software or

programs” and Section 5.1. These sections rely solely on antivirus to prevent malware infections. Numerous data breaches have shown that antivirus alone cannot protect against all malware [40]. These limited-scope requirements leave organizations exposed to multiple attack vectors that will most likely occur frequently and have moderate severity. These sections should mandate additional controls such as application whitelisting, blocking access to areas that permit persistence (e.g., Windows Registry Keys), and enforcing least-privilege access.

Section 1.3.7 focuses on limiting the disclosure of private, internal IP addresses from firewalls and routers, but fails to discuss any other services that could leak the same information, such as a domain name server or internal files (e.g., Word documents) improperly exposed to search engines. Attackers have leveraged common techniques such as “Google Hacking” to discover internal network configurations and sensitive systems like a domain controller [127]; expanding the scope of this moderate-risk issue to limit external enumeration would improve its security.

Sections 11.1.c and 11.1.d actually incentivize less-secure practices. Each subsection defines additional audit checks that an assessor must conduct only if a particular security control is in place (wireless scanning and automated monitoring, respectively). Under this policy, financial sanctions associated with non-compliance could lead a security professional not to implement a security control at all rather than risk having it assessed as non-compliant — if it is not present, the organization is automatically compliant. These two particular controls would have a negligible overall impact if they were not in place; therefore, this is a low-risk issue. If the PCI SSC believes these security controls are important, they should be mandatory

rather than optional; otherwise, these sections should be eliminated entirely.

Ambiguous specification. There are 10 issues within PCI DSS in which insufficient details create ambiguity or uncertainty. An example of a high-risk security concern with a frequent probability and moderate severity stems from Section A1.1 and limits the usage of Common Gateway Interface (CGI) scripts to control privileged access to cardholder data. This control is sound but is overly narrow; in modern systems, there are a variety of applications that could access or manipulate cardholder data in ways similar to CGI scripts. Simply replacing “CGI scripts” with “applications” would improve the clarity of this control.

Section 11.3.3 discusses corrective action plans for vulnerabilities discovered during penetration tests. The section does not specify how soon after a penetration test vulnerabilities must be addressed, nor the party responsible for fixing the vulnerabilities. Based on the researchers’ past experiences with organizations delaying remediation, this security concern has a high risk level with a frequent probability of occurring and a moderate severity. Moreover, the non-validator assessor confirmed that in his experience, organizations often delay remediation, and typically dedicate one to two full-time employees for 30-40 days prior to an inspection to ensure remediation is complete just in time [149]. This section should specify a time limit (based on vulnerability severity) for addressing issues discovered during a penetration test and clarify the party responsible for fixing the vulnerable condition.

4.3.3 Expert validation

To assess the PCI DSS findings, researchers partnered with an organization that is a PCI SSC member. Expert E3 represented this organization, possessing 18 years of experience advising the security practices of large financial organizations, assessing organizations' adherence to PCI DSS security controls, and conducting digital security assessments against networked environments. E3 confirmed past utilization of PCI DSS as a line-by-line checklist as they audited organizations in the past.

E3 was asked to assess all eight high-risk issues and a randomly-sampled subset of seven moderate issues and five low-risk issues; this accounts for 43% (20 of 46) of the issues from the audit. E3 confirmed 18 of the issues and categorized the remaining two as plausible, although he had not experienced them.

There was no statistical difference between probability estimates between E3 and the auditors ($p = 0.77$), but E3 assessed issues to be statistically more severe with medium effect ($p = 0.003$, $r = 0.469$). During the post-survey discussion with E3, he stated that the financial impacts of digital security breaches involving cardholder data caused him to increase his assessed impact of each issue — had these issues been present within another business sector, E3 would not have assessed them as severely.

The first issue assessed as plausible rather than confirmed involves Section 1.3.7 and information disclosure. E3 indicated that internal data exposure is “inconsequential if boundary configuration is correct,” meaning an administrator is

successfully limiting which inbound connections from external entities are allowed to communicate with private IP addresses. However, E3 acknowledged that the security concern would exist if these additional controls are not in place.

The second issue E3 flagged as plausible rather than confirmed involves Section 5.1's reliance on anti-virus software. According to E3, organizations have lessened reliance on anti-virus for protection; he argued that Section 5.1 would have minimal impact on organizations with defensive strategies for protecting network segments, user accounts, and key resources.

Additional defenses. E3 recommended account-protection solutions such as multi-factor authentication to mitigate concerns such as VLAN attacks or insufficient protection of passwords.

As discussed for P1075 above, both the issues E3 assessed as only plausible and his recommended additional defenses hinge on additional security controls beyond the PCI DSS standard; one cannot necessarily assume non-mandated controls will be applied.

4.4 Results: NERC CIP 007-6

4.4.1 Overview

The North American Electric Reliability Corporation (NERC) Reliability Standards define the requirements for planning and operating North American bulk power systems (BPSs), defined as large interconnected electrical systems consisting

of generation and transmission facilities and their industrial control systems [159]. All BPSs within the continental United States, Canada, and the northern portion of Baja California, Mexico must comply with NERC Reliability Standards, meaning that these security controls affect most people living within these areas. The NERC Critical Infrastructure Protection (CIP) Committee, which oversees the set of standards, comprises representatives from 30 companies and municipalities across North America [160]. Although NERC is an international not-for-profit, its regulatory authority stems from section 215(e) of the Federal Power Act and Title 18 Code of Federal Regulations §39.7. The set of standards that make up CIP date back to 2009; in this study, auditors used CIP 007-6, which is the 2014 revision of the Systems Security Management standard. CIP 007-6 includes key sections for securing ports and services, patch management, malicious code prevention, event monitoring, and access control.

NERC Regional Entities are the organizations responsible for conducting audits and monitoring adherence to the compliance standards within their assigned geographic region. On-site audits typically last one week and occur every three years. According to the expert validator E4, a NERC Regional Entity employs four to seven auditors per assessment, drawn from a pool of full-time employees. Auditors typically conduct 7-30 audits per year. E4 also noted that organizations allocate a large portion of their operating budgets toward compliance and often spend one year preparing for their audit.

Of the three standards assessed, NERC has the strongest sanctions (which can actually create security concerns, as discussed in Section 4.4.3). The maximum fine

for a compliance violation is \$1M (U.S.) per day; NERC or the applicable Regional Entity determines the monetary fine [158]. According to the expert participant, fines for NERC non-compliance are common. Recently, NERC levied a \$10M fine against Duke Energy for 127 security infractions between 2015 and 2018 [91].

Qualitatively and quantitatively, CIP 007-6 had the strongest security controls of the three documents assessed (shown in Figure 4.5), but numerous issues exist that create security gaps within compliant organizations.

4.4.2 Findings

NERC CIP has 79 individual controls. The internal audit identified 21 total issues: one extremely-high-risk, four high-risk, six moderate-risk, and 10 low-risk. One additional issue was discarded as a duplicate entry.

Security concern trends. Seven of the 21 issues identified deal with overly vague terms such as “when feasible” or “unnecessary” without defining feasibility or necessity. For example, Section 5.7 calls for limiting authentication attempts or generating alerts when feasible. The subjectivity of these statements can lead to misinterpretations of the standard and potentially permit insecure actions. Mandatory compliance standards should be mandatory; either administrators must limit authentication attempts or it is merely a suggestion. Additionally, none of the 21 issues identified specify which entity is responsible for specific actions, which can lead to inaction. Notably, NERC identified “confusion regarding expectations and ownership of tasks” as a key problem contributing to Duke Energy’s non-compliance

Risk Estimates for NERC Compliance Standard

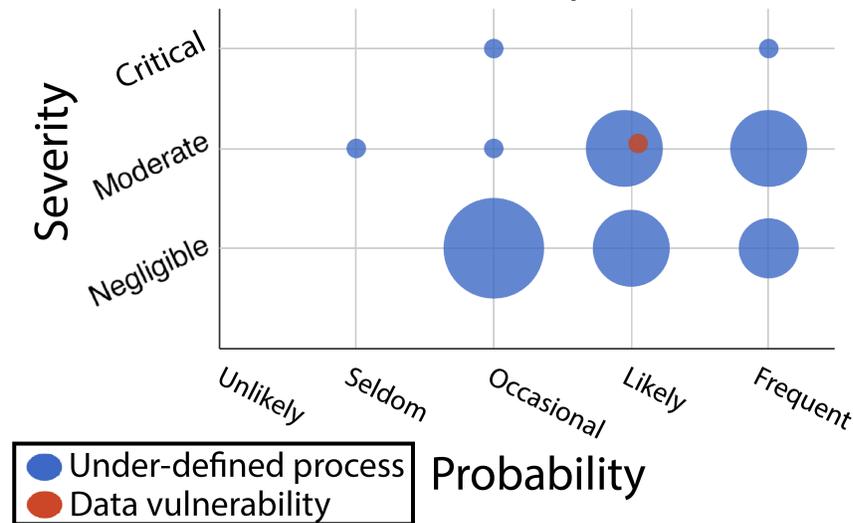


Figure 4.5: Distribution of security concerns identified for NERC CIP 007-6. Color indicates the type of security concern; each dot indicates by size how many security concerns were identified with a given type, severity, and probability. Under-defined processes were most common (n=20).

and eventual fine [91]. Below details examples of findings, organized by their perceived root cause.

Data vulnerability. Based on the assessment, CIP 007-6 only has one moderate-risk issue pertaining to a data vulnerability. Section 5.1 states that administrators should “[h]ave a method(s) to enforce authentication of interactive user access, where technically feasible.” This caveat allows legacy equipment with no provision for authenticating authorized users to endure within a secure environment. It is well-documented that legacy systems often have no password, transmit unencrypted passwords, or never change passwords from their default settings [48]. This permits attackers and insider threats to easily gain control of legacy systems, which could range from sensitive databases to the logical system “off switch.” Secure, authenti-

cated access should be a hardware and software requirement for all systems in this critical environment, reducing the likelihood of such an attack.

Under-defined process. The remaining 20 issues involve processes that are not sufficiently detailed for a secure implementation. Section 2.1 has an extremely high risk, as written, due to the critical severity and frequent probability of a security concern occurring within critical environments. The issue involves the implementation of a patch management program for improving the security of systems. Throughout all of the NERC CIP documents, no mandate exists that organizations must maintain a representative test environment for patch evaluation. Applying patches directly to live systems that provide power — including to critical infrastructure such as hospitals — could result in outages and corresponding loss of life; one such incident occurred in March 2008 and caused a nuclear power plant to shutdown [119]. Testing patches prior to live deployment allows administrators to observe potential effects within their environment and reduce the likelihood that unforeseen outages will occur as the result of the patch [211].

There is a potential loophole in Sections 2.1 and 2.2, which rely upon validated sources for patches against known vulnerabilities. If the entity responsible for patching systems does not provide sources, then there is no requirement for patching. Additionally, CIP 007-6 does not account for patches from external sources beyond the list of valid providers. Do administrators have a requirement to apply a patch for a known vulnerability if it is from an outside source? According to Cardenas, there are instances where applying a patch may violate the certifica-

tion of certain control systems [39]. This loophole presents a high risk due to the critical severity associated with unpatched systems in these environments and the occasional probability of their presence.

Section 5.3 requires administrators to “identify individuals who have authorized access to shared accounts.” Shared accounts have a moderate-risk threat, as administrators are unable to deploy granular controls on a by-need basis. Shared accounts also inhibit auditing, as the compromise of a privileged shared account could lead to the spread of malware or outages that administrators cannot positively attribute to one individual. Researchers from Sandia National Lab identified this security concern in 2003 [188].

Section 5.4 outlines provisions that allow systems to retain their default usernames and passwords if documentation supports that the “vendor passwords were generated pseudo-randomly and are thereby unique to the device.” The auditors believe that vendor-generated pseudorandom credentials can present a threat to BPSs if the pseudorandom algorithm is predictable (e.g., basing its seed on a unique identifier such as a serial number). This type of exploit requires in-depth knowledge about the vendor’s algorithm and might seldom occur despite posing a moderate risk to the environment. Compliance authors should eliminate this provision entirely and mandate that administrators change all system credentials before allowing a system to communicate with a BPS.

There is a high-risk issue in Section 4.3 concerning event log retention. CIP 007-6 requires facilities to retain 90 days of consecutive logs and demonstrate proof of such practice over a three year period. This relatively short-term rolling requirement

can interfere with incident investigations, given that advanced persistent threats can operate within networks for years before being detected [2, 220]. Organizations should be mandated to ship logs to a data warehouse for long-term storage and investigation support if needed.

4.4.3 Expert validation

A government organization that focuses on national security issues assisted in validating the CIP 007-6 findings. Expert E4, as the organization's representative, has 20 years of experience conducting digital security assessments against BPSs. E4 confirmed first-hand utilization of NERC CIP standards as a checklist for past audits. E4 has served on numerous executive councils and federal-level panels addressing cybersecurity concerns within industrial control systems. Most notably, E4 was a contributing author to many of the NERC CIP standards.

Due to the complexity of NERC CIP, the 60- to 90-minute survey could include only nine audit findings (43%). The extremely-high risk issue and all four high-risk issues were included, and included two randomly-sampled moderate-risk and two low-risk issues. Of these, E4 confirmed one issue and one broader trend, rejected one issue, and categorized the remaining seven issues as plausible.

When comparing the auditors' risk estimates to those of E4, there was no statistical difference between severity estimates ($p = 0.18$), but the auditors assessed the issues to be statistically more likely with a large effect ($p = 0.01$, $\eta^2 = 0.603$). E4, addressing these comparison differences, indicated that CIP 007-6 relies heavily on

the broader framework of CIP standards and that security controls in other CIP documents help harden the overall environment. Like E1, E4 commented that he was unable to assess CIP 007-6 only “if standard is followed as written and nothing else,” as directed (Appendix A.2). As such, E4 indicated that he rated each issue as less likely given his broader understanding of the compliance framework.

The issue E4 rejected involves the loophole identified in Sections 2.1 and 2.2 for patch management. E4 stated that “each item in the [system] baseline needs a source identified or evidence that a source no longer exists.” In his experience, he never encountered an external source that could provide a trusted, proprietary patch. However, E4 acknowledged that if a component is no longer supported or a source no longer exists, it is highly likely that the component will remain unpatched against all future publicly-disclosed vulnerabilities.

E4 confirmed the log-retention issue identified in Section 4.3, attributing the known gap between log retention and investigation windows to two factors. Primarily, the specification is written to account for the limited log retention capacity on most devices within a BPS environment. Second, most administrators and BPS owners are unwilling to connect to and aggregate event logs on an external platform. Placing an additional device within the environment (for logging) increases the number of devices an attacker can exploit and is one more device potentially subject to financial sanctions.

E4 also confirmed the risks of not specifying a responsible party for tasks, a trend the researchers identified, and referenced the aforementioned Duke Energy fine as an example.

Additional defenses. E4 noted that the best additional defense for mitigating the issues identified was to upgrade system components to more modern devices that can implement up-to-date best practices (e.g., multi-factor authentication, strong passwords, limiting login attempts. As with P1075 and PCI DSS, organizations that only meet the minimum required by the standard will not be able to take advantage of these defenses. E4 confirmed that while some facilities exceed this “minimum baseline” and systematically replace obsolete devices, he has also audited facilities that only follow the standard exactly as written.

Other recommendations. E4 described additional security concerns that the auditors did not identify. Subsets of NERC CIP security controls apply to BPSs based on how much power the system produces, creating three tiers of compliance: the highest tier of power producers are subject to all security controls, while the lower tiers of power producers must comply with decreasing subsets. E4 believes this perversely allows attackers to use publicly-available information to locate facilities that must adhere to fewer security controls and then systematically target the controls that may not be present. E4 therefore argues that NERC must standardize controls across all facilities to mitigate the targeting of smaller stations.

Additionally, E4 stated that the zero-defect culture and high fines associated with NERC’s sanctions program can incentivize minimum-effort security. Organizations that undertake additional security precautions beyond NERC CIP mandates may discover vulnerabilities that would not otherwise be identified. NERC levies

finer for non-compliance even when organizations self-report such vulnerabilities, potentially punishing organizations for transparency. E4 believes this behavior inhibits sharing of information across the power sector and collectively lowers security for all facilities. He argues that NERC could reverse this trend by eliminating fines associated with self-reporting and providing “credits” to organizations that contribute to the overall health of the power sector.

When discussing concerns with log retention, E4 recommended that all facilities should contribute toward a common log aggregation center, where security professionals could conduct in-depth security-breach investigations spanning all NERC-compliant facilities.

4.5 Disclosures

Researchers made an effort to disclose the findings responsibly. Compliance standards typically have a request-for-comment (RFC) period that allows for the submission of comments, concerns, and recommendations during a fixed window. During this study, none of the standards assessed had an open RFC, and no clearly defined channel existed for reporting security concerns, either directly to affected organizations or at the federal level. Using the partners as mediators, all of the findings were sent to the IRS; the PCI Security Standards Council; a contributing author of the NERC CIP standards; the United States Computer Emergency Readiness Team (US-CERT); the MITRE Corporation’s Common Vulnerabilities and Exposures (CVE) team; and the Department of Homeland Security. Even though this

study was completed between October 2017 to September 2018, as of June 2020, researchers are still actively working with the U.S. Government to help organizations understand the impact of the findings. Overall, there have been varying levels of success with disclosure attempts, as described below.

4.5.1 Disclosure intent

Before disclosing any of the study findings, researchers envisioned that the findings could help the U.S. federal government establish a centralized repository of best-practices and lessons learned associated with compliance controls. This information could (1) help authors of compliance programs adopt language that has been proven to be effective, (2) help organizations understand potential risks they could inherit, and (3) allow compliance programs to incrementally evolve at speed with technology to remain secure and relevant.

To achieve this intent, disclosure occurred through contacts at high levels in the federal government, at organizations responsible for creating compliance programs, and directly at the organizations that use the affected compliance standards. This involved extracting contact information from the audited standards, extracting information from publicly-available official sites, contact-chaining through personal contacts, and searching through social media for appropriate contact information. Below are the attempts and shortcomings in trying to achieve the disclosure intent.

4.5.2 IRS P1075

The IRS, NIST National Vulnerability Database (NVD), US-CERT, and the MITRE Corporation were contacted to disclose the P1075 findings. US-CERT was the first organization to respond to the disclosure attempt. After exchanging several emails, their technicians concluded that “CVEs are assigned for specific vulnerability in implementations. Each issue that requires a ‘separate patch’ can get a CVE [213].” In a series of email and phone exchanges, it was argued that each of the recommendations provided are “patches” for the vulnerable portions of the compliance standards, but US-CERT stated that the “patches” identified must be tied to a specific piece of software. Future research that correlates security concerns to compliant software may be eligible for CVE identification numbers using US-CERT’s definition.

Both NIST NVD and the MITRE Corporation indicated that compliance documents are outside their scope of responsibility, with MITRE stating “that a reasonable person can conclude that IRS Publication 1075 was never intended to have a level of abstraction that was sufficient to direct secure coding [143].” Contradicting this argument, the partners confirmed that auditors are indeed using compliance standards such as P1075 as a line-by-line checklist to confirm controls at levels as granular as access control lists on firewalls.

The IRS was contacted nine times via personal contacts, emails, and phone calls over the span of three months. To date, the IRS has not provided any form of acknowledgment other than the automated responses from `SafeguardReports@`

irs.gov, the only point of contact listed in IRS P1075.

4.5.3 PCI DSS

Unlike P1075, there was success in responsibly disclosing the findings to members of the PCI Security Standards Council. Researchers established a memorandum of understanding with a PCI SSC member organization; in turn, this organization provided the findings to the PCI DSS Version 4 Working Group.

The recommendation for improving the “Network Segmentation” section of PCI DSS has already been implemented within Version 4, prior to the opening of their RFC submission window. This change will apply PCI DSS guidelines to the entire networked environment and not only an isolated subnet with cardholder data – this change could help reduce the likelihood that an attacker could gain access via unprotected portions of the network. Additionally, the v4 Working Group is considering incorporating all feedback associated with the ambiguous specification findings.

4.5.4 NERC CIP 007-6

Expert E4, after providing feedback, noted that the recommendations would be included at future working groups for CIP revisions. However, it could be years before the next CIP update (potentially taking the recommendations into account) is released. Additionally, the partnered organization for CIP disclosure is incorporating the feedback into a comprehensive evaluation of electric grid security. More

than any other expert, E4 provided years' worth of lessons learned from CIP audits and helped explain why the standard was written the way it is. Given that the group of researchers had little experience with industrial control systems or the electric grid prior to this study, Expert E4's insight was truly invaluable for assessing the validity of the findings.

4.5.5 Federal-level recognition

To approach problems with federal-level compliance standards in a top-down manner, researchers met with representatives from the NIST National Cybersecurity Center of Excellence (NCCoE) to discuss the findings [152]. Researchers highlighted that IRS P1075 Section 9 (which contains 49% of the P1075 security concerns discovered) is copied from older versions of NIST SP 800-53 (NIST has since updated SP 800-53 twice). NCCoE offered to incorporate the findings into future document revisions. In ongoing revisions that began before the meeting, NIST acknowledged in draft SP 800-53v5 that organizations may inherit risk when implementing mandated security controls; that is, standards may create security problems [164]. Specifically, NIST describes deliberate efforts to remove ambiguity, improve understanding of responsibility, and keep controls up to date, corroborating many findings from the study.

Next, researchers contacted the Department of Homeland Security (DHS) National Protection and Programs Directorate. Several personnel within the Federal Network Resilience Division expressed interest in assisting with the findings; how-

Document	Controls	Total Issues	Extr. High	High	Mod	Low
IRS	309	81	2	13	32	34
PCI	851	46	0	8	22	16
NERC	79	21	1	4	6	10

Table 4.3: Security concerns, by document and assessed risk

ever, the DHS Office of External Affairs for Cybersecurity and Communications directed the contacts to cease communication and did not provide any alternative mechanisms for disclosure. This decision continues to provide friction between the agent contacts at DHS and the organization – the agents are motivated to help remedy the discovered issues. Through open publication, these agents are now able to use the findings and shape future compliance development on their own.

4.6 Discussion

This research provides the first structured evaluation of security issues within digital-security compliance standards. This study finds that when compliance standards are used as checklists, with “by-the-letter” implementation of security controls, security concerns can be created. The systematic approach identified security issues spanning multiple root causes and varying levels of risk (shown in Table 4.3). This section highlights common issues across the audited compliance standards, potential mitigations, recommendations for reconsidering compliance programs, and opportunities for future work.

Common issues. When considering the findings, some common issues become apparent. All standards assessed exhibit under-defined processes and vague writing.

While issues of vague writing may not seem as immediately dangerous as, for example, failing to identify passwords as sensitive data requiring protection, they have important implications when standards are treated like point-by-point checklists.

Many issues stem from passive voice, creating ambiguity concerning who is responsible for exactly what actions. Using the active voice to construct compliance controls is a best practice that helps eliminate uncertainty and ensure there is a responsible party for requisite actions [109]. If it is not feasible to eliminate passive voice (perhaps because it would prescribe organizational structure too strongly), standards authors could perhaps include supplemental best-practice recommendations for identifying responsible personnel. In addition, the standard might require each implementing organization to create a written plan identifying who is responsible for each requirement.

Further, numerous compliance controls did not have clear deadlines for action. Compliance standards should define expected periodicity (e.g., every 30 days) or thresholds for action (e.g., within 12 hours of an event). These issues with deadlines seem especially concerning in light of observations by several auditors that many problems are only mitigated during an immediate run-up to a compliance audit, as part of preparations to pass.

Terms such as “when feasible” and optional guidelines create confusion about what is actually required and may provide an illusion of more security than what is actually provided. In some cases, this wording reflects practical limitations: for example, updating legacy power systems to include modern security controls (NERC CIP) could require multi-million-dollar equipment investments and degrade near-

term power availability. Nonetheless, categorizing clearly insecure systems as “compliant” simply because there is no feasible alternative is counterproductive. Instead, compliance standards could adopt a third category that does not punish the affected organization but still indicates clearly to administrators and auditors that the situation is suboptimal and further precautions are needed. For clarity, authors should move optional guidelines into supplemental documents separate from mandatory compliance.

Each compliance standard has weak controls for user-access review and revocation procedures. To mitigate insider threats, compliance standards could mandate frequent review of active user accounts, as well as access termination before formally notifying an employee who is terminated.

Lastly, and perhaps most concerning, none of the compliance standards assessed have mechanisms for reporting security concerns. Without a direct line of communication with a governing body, it is likely many discovered security concerns will remain unaddressed. The lack of a centralized CVE database-like construct for reporting problems with compliance standards affects both governing bodies and compliant organizations. Governing bodies do not have a reference for common mistakes when developing compliance standards, meaning issues are likely to repeat across multiple standards. Additionally, this lack of transparency prevents industry-wide alert notifications for issues within a compliance standard; if a researcher discovers a valid security concern, all affected parties should be notified. Further, no standard could be expected to perfectly capture all needed security controls; as several of the experts noted, strong security practices often require going

beyond the minimum established by a standard. A centralized repository would also present an opportunity to recommend additional best practices to build upon compliance and mitigate any reported gaps.

Recommendations. The work highlights difficulties that can arise when compliance standards are used as checklists, regardless of their original intent. This approach seems inevitable when a standard is associated with potentially large penalties for non-compliance, but little or no incentive for going beyond the minimum requirements. This state of affairs suggests a need for rethinking the compliance paradigm more broadly.

First, authors of compliance standards should take into consideration that their standards might be used as an audit checklist. Whenever possible, guidelines should be broadly applicable across a particular domain but concrete enough that line-by-line compliance will provide meaningful security. Of course, writing guidelines that achieve this ideal is difficult and may sometimes be impossible; standards authors should explicitly consider tradeoffs between generalizability and secure implementation when making choices. Providing supplemental documents describing potential such issues could help standards implementers manage resulting risks.

Secondly, authors should identify opportunities to craft compliance standards that improve audits beyond checklist assessments and consider an organization's overall security culture. Provisions for a rewards program could incentivize organizations to bolster security. As examples, organizations that take proactive measures beyond minimum requirements or organizations that publish digital security lessons

learned could receive some limited safe harbor against future sanctions. As discussed during the audit of NERC CIP standards, an organization that responsibly discloses and remedies a vulnerable condition is still liable for financial sanctions. Allowing organizations to self-report issues with less fear of sanctions could incentivize better behavior and increase transparency, with potential benefits for the entire associated sector [16].

Another consideration for standards authors is that rapidly changing technology necessitates rapidly updated security mechanisms. An effective standards update mechanism should allow easy reporting of issues and enable fast revision of the standard itself, while avoiding imposing costs on organizations that cannot immediately meet the new requirement. Newly updated standards could provide suggestions for transitioning and require organizations to provide a plan for becoming compliant with the updated requirement within some specified time period.

Chapter 5: Baseline Security: Security Implications of Policies, Laws, and Regulations in the Cloud

In this chapter¹ I present research that builds upon Chapter 4 and focuses on the security implications of the FedRAMP compliance program on U.S. federal organizations that leverage cloud-based services.

Researchers identified 46 issues that may present security threats to organizations that use FedRAMP-approved programs. Additionally, researchers identified four threat models that appear to be neglected throughout FedRAMP and could pose significant threats if not properly handled.

Alongside Chapter 4, this chapter supports the notion that baseline security programs are insufficient to provide adequate security against current threats due to security concerns that may be introduced by the following compliance programs. Complementing baseline security with proactive planning and implementation strategies present opportunities to provide better coverage against security threats.

¹Published as [194]

5.1 Background

In response to the coronavirus pandemic, millions of people moved to online videoconferencing for work and school. Zoom, and other products like it, skyrocketed in popularity and usage. Hackers and security researchers quickly discovered security vulnerabilities in some of these services, including the ability to hijack meetings. CitizenLab also found that Zoom’s cryptographic strength was less than advertised [130]. The Department of Defense, NASA, Google, and others banned use of Zoom in response. While Zoom is continually evolving plans to bolster security [26], this national-level attention on Zoom calls attention to a larger problem for the United States government.

Since April 2019, Zoom for Government has been compliant with the Federal Risk and Authorization Management Program (FedRAMP) [75]. Achieving this authorization required adherence with a set of government-defined security controls for cloud-based services, but despite this certification, security issues remained, highlighting the danger in assuming that compliance implies security. FedRAMP is more comprehensive and flexible than many other compliance programs, yet dangerous gaps remain. For example, despite FedRAMP existing to protect government systems and information, no security control in FedRAMP prohibits cryptographic keys used by FedRAMP-compliant programs from being generated by a foreign nation. This could be exploited to allow a hostile nation to read sensitive information belonging to federal organizations and its employees.

As of May 2020, FedRAMP has authorized 188 programs for use, with 49

additional programs currently in evaluation. FedRAMP's security (or lack thereof) impacts more than 5 million systems and devices across more than 150 government agencies. Given this scope, it was necessary to dissect the controls within FedRAMP to understand security gaps. Specifically, the focus was **to identify the security controls that could lead to sub-optimal security conditions within an organization despite being compliant with FedRAMP.**

FedRAMP is currently based on the NIST 800-53 revision 4 standard, which was originally published in April 2013 and updated in January 2015. A host of new threats to information security have emerged since this time: organizations have migrated toward bring-your-own-device strategies that let employees attach their personal devices to private networks and organizations have shifted from on-premises servers to the cloud, to name just two examples.

FedRAMP incorporates controls for a diverse set of security considerations, including several categories of protection. The FedRAMP control naming convention uses the two-letter category abbreviations below combined with numbers to indicate the groupings of controls.

- Access control (AC): security mechanisms that govern how systems and data are accessed
- Audit and accountability (AU): requirements for assessing the implementation of controls
- Identification and authentication (IA): mechanisms for verifying users
- Incident response (IR): programs and plans for handling security incidents

- Media protection (MP): mechanisms for protecting various storage devices
- Physical and environmental protection (PE): requirements for safety and health
- Risk assessment (RA): requirements for understanding risk and risk mitigation
- Security assessment and authorization (CA): controls for conducting penetration tests
- System and communication protection (SC): controls for ensuring privacy and availability
- System and information integrity (SI): controls for data resiliency
- System and services acquisition (SA): controls and restrictions for the procurement of digital systems and devices

FedRAMP will likely adopt the NIST 800-53 revision 5 standard once it is finalized, but the standard is still under revision at the time of writing, with final draft comments due in May 2020. This slow pace of updating offers stability to cloud computing companies, which can design security and compliance programs to a known standard, but also potentially leaves new and emerging threats unmitigated.

5.2 FedRAMP evaluation results

A line-by-line audit of FedRAMP controls revealed a number of security concerns. Throughout FedRAMP, the focus was on identifying the security controls

or policies that can lead to sub-optimal security conditions when implemented as written.

Each security concern identified carries an associated risk. Frameworks such as the Composite Risk Management framework (shown in Figure 4.2) calculate risk as a function of the probability of an event occurring and the severity associated with that event. Using this model, researchers can assess that a likely event with a negligible severity carries a low risk to an organization, while a likely event with a catastrophic severity (loss of life or significant financial loss) carries an extremely high risk to an organization.

In total, the audit of FedRAMP identified a total of 46 independent issues across 325 security controls. Of these, one issue presented an “Extremely High” risk, with four rated as “High,” 13 as “Moderate” and 28 as “Low” risk (depicted in Figure 5.1). These 46 issues fall into three categories: an ambiguous specification, an obsolete reference, or a risk to data. Below are detailed examples from these categories.

5.2.1 Ambiguous specifications

Ambiguous specifications occur when two organizations can implement drastically different security controls and both are compliant; these types of issues represent the bulk of the findings. An example would be if ACME Company implements the control AC-06(09) for ensuring an “information system audits the execution of privileged functions” by auditing after each individual occurrence, but the Wid-

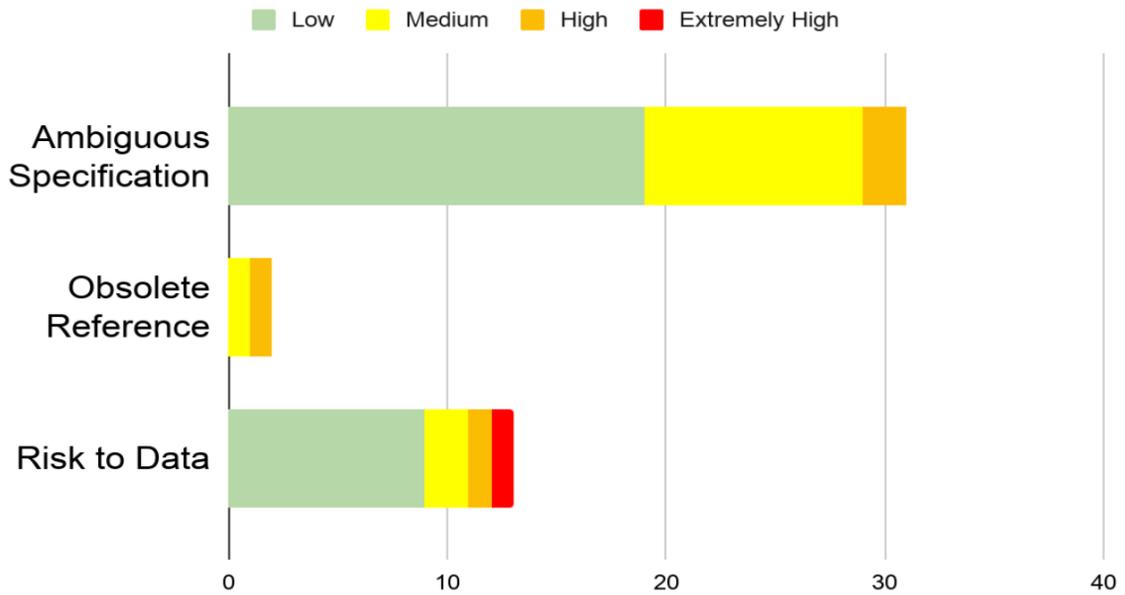


Figure 5.1: Distribution of the 46 security issues identified within FedRAMP, by category.

gets'R'Us Company only conducts such audits every 60 days. Both companies are compliant — they both perform the mandatory audit — but ACME should be considered more secure. In the context of this example, an organization's interpretation and implementation of this control could be the difference between detecting a malicious threat during the initial stages of a compromise or only after attackers have already stolen sensitive victim data from the network.

The FedRAMP program uses certified third party assessment organizations, or 3PAOs, as auditors to review and assess the FedRAMP compliance of any organization applying for authorization. The 3PAO, therefore, will be the primary arbiter of any ambiguous specifications or questionable implementations. While the FedRAMP program aims for all 3PAOs to operate with equal rigor, differences can arise. The 3PAO for ACME may drill into the implementation details of the audit

check, as well as all functionality within ACME’s system, to ensure that all privileged functions are run through the same system audits. The 3PAO for Widgets’R’Us, meanwhile, may fail to dig deeply enough to catch that not all functionality is audited, or may fail to recognize that auditing privileged functions every 60 days is not compliant with the spirit of the control.

In total, researchers identified 31 unique issues involving ambiguous specifications: two high-risk, 10 medium-risk, and 19 low-risk. Inconsistencies in organizational and 3PAO interpretations of these controls may result in an increased threat to organizations using FedRAMP-compliant programs.

There are 11 ambiguous specifications, with varying levels of risk, that fail to incorporate a time-based factor: how often should a task be performed or how soon after an event should a task be performed? A high-risk example includes AU-06(01) which requires automated analysis of data artifacts to support investigations into suspicious activities. Examples of artifacts include records of every website a user visits, every time someone attempts to log into a user account, or every time an antivirus program generates a suspicious activity alert. These artifacts, depending on the size of the network, could amount to billions of individual records and require 4-8 petabytes of storage a day. This control leaves up to the organization (and its 3PAO auditor) essential decisions like artifact retention periods, correlation frequencies, and report availability. Given that advanced persistent threats can operate within networks for years before being detected, an organization must adopt a log retention policy that would allow them to investigate compromises that may have occurred 6-12 months in the past [220]. Correlation

frequencies and report availability are intertwined; how often an organization aggregates and correlates data from network streams, end points, and service platforms directly shapes how soon they can process this data and provide a meaningful report that can support an investigation. Real-time processing can be a substantial monetary investment, but correlating petabytes of enterprise data from the past 6 months from a cold start may take too long. Obviously there is a middle ground here, and researchers recommend the inclusion of a best-practice timeline for how often organizations should conduct data correlation for threat analysis.

There are 10 ambiguous specifications involving authentication mechanisms that may under some interpretations allow an attacker to gain access to resources. Two of these issues involve weak passwords. AC-18(01) protects wireless access to systems using authentication and encryption, but as-is, would allow an organization to use encryption algorithms with known cryptographic weaknesses or trivially weak WiFi authentication passwords, such as the letter ‘a.’ IA-05(04) requires password strength checks using arbitrarily-defined requirements (such as complexity and length) but does not compare user-generated passwords against common passwords or passwords found in data breaches. (For example, ‘P@\$\$Word123’ passes most complexity checks but should not be used.) OWASP and Have I Been Pwned maintain repositories of commonly used passwords and account breach data and provide interfaces for services to check passwords against; however, use of services such as these is not required under any FedRAMP control. Understanding and mitigating the use of commonly used passwords and credential reuse across multiple accounts would improve the security of user accounts.

There are five issues involving multi-factor authentication (MFA) and authenticators. These five controls permit SMS and email authenticator codes. Depending on implementation, these additional layers of authentication may simply present additional hurdles that a capable attacker can likely bypass. SMS-jacking is a known attack vector that allows an adversary to port their victims' phone numbers to phones that the attacker controls, allowing them to receive victim SMS authenticator codes [4]. Attackers can also intercept unencrypted emails through man-in-the-middle attacks or traffic sniffing, allowing them to gain access to email-based authenticator codes. Hardware-based authenticators (such as Yubikeys or Titan keys) and software authenticators (like those from Google Authenticator or Duo Security) have other complications but typically offer a much more secure approach to MFA.

Five controls allow organizations to provide their own definition of secure. AC-01 and AC-03 allow an organization to develop its own access control procedures and policies. These serve as the basis for most other controls within FedRAMP. Using an exaggerated example, imagine that ACME Company empowers an employee with 15 years of experience in access control to create its access control program, while Widgets'R'Us subcontracts the task to the bagel vendor in the front lobby. As long as both companies produce the requisite documents, both are compliant, but ACME is more likely to have a robust program.

Controls such as MP-07, IR-01, and AC-04 allow organizations to define what information is considered sensitive, how data can be exchanged between interconnected systems, and how the organization should conduct incident response investi-

gations respectively. None of these controls should be arbitrarily defined, but rather rooted in best practices and iteratively updated after each internal evaluation, security exercise, or real-world data breach.

5.2.2 Obsolete references

Obsolete references occur when the document mandates the use of an outdated policy or references a document that has since been superseded – there are two instances of this in FedRAMP.

Throughout the document, FedRAMP references “FIPS Publications 140-2,” which was replaced by 140-3 in September 2019. Additionally, IA-05(01) requires organizations to enforce password expiration policies that NIST SP 800-63 has since rescinded – this high-risk issue is shown to encourage insecure practices such as writing newly rotated passwords near user workstations [205]. These two issues highlight a greater concern: FedRAMP has not been updated since August 2018. As technologies and best practices evolve over time, FedRAMP authors must reconcile the need to remain secure with the requirement to remain adaptive. Compliance programs such as FedRAMP should reference other security documents for best practices, but only the most recently updated versions.

5.2.3 Risks to data

There are *risks to data*, defined as security controls that expose sensitive information to an attacker, as the category posing the greatest potential risk to or-

ganizations. In total, there are 13 risks to data: one extremely high, one high, two medium, and nine low-risk issues.

Organizations using FedRAMP-compliant solutions must consider who has access to protection mechanisms and how they can be accessed.

AC-17(02) specifies the protection mechanisms for remote access to systems. IA-05(02) details requirements for PKI-based authentication. SC-12 allows organizations to define requirements for key generation, distribution, storage, access, and destruction; SC-12(02) and SC-12(03) specify requirements for symmetric and asymmetric keys respectively. None of these controls mandate protection requirements for cryptographic keys. Given that FedRAMP relies on cryptographic keys for many security controls, an attacker can exploit this policy weakness to target and gain access to unprotected cryptographic keys. Control AC-17(02), which governs remote access to systems, would present an extremely high-risk situation for organizations if keys are not adequately protected.

MP-05(04) outlines protection mechanisms for media, but it does not include protection mechanisms for keys or passwords used to encrypt the stored data. Sending passwords in cleartext emails or SMS would drastically reduce the efficacy of password-protected devices that have been intercepted by an adversary.

FedRAMP lacks oversight over the systems that provide security in an environment. Who is responsible for securing the security systems? SC-08(01) mandates that systems enforce data integrity checks during transmission, but does not consider tamper controls against those checks. Consider two scenarios that could occur if an adversary gained control over an integrity checker. They have

the ability to enact a denial of service attack by flagging all inbound and outbound traffic as corrupted, causing endless re-transmissions of data. Additionally, they have the ability to modify content in transmission and verify its integrity – this could be exceptionally damaging to an organization if the attacker modified business records to annotate significant financial losses or fired an entire company via a cryptographically-signed email.

SA-10(01) enforces integrity checks against software updates and patches but does not consider the compromise of an update server. In some situations, it may be appropriate to confirm the validity of updates from external vendors for critical networking devices, endpoint protection software, and workstations. These update servers are a juicy target for attack, as they would give an adversary the ability to exploit an entire customer base from one system.

Similarly, RA-05(01) mandates the use of vulnerability scanners and SI-03 mandates the use of malicious code scanners. Both solutions should be FedRAMP compliant and would require privileged access to data and systems to perform their intended functions. Neither solution is accounted for within FedRAMP as a potential threat vector. Attackers could manipulate these scanners to provide false negatives for alerts, allowing them to bypass defenses and gain access to vulnerable systems. These systems could also become an internal attack platform for adversaries taking advantage of their privileged, trusted access within the internal network. FedRAMP controls calling for anti-virus software to be run on all systems similarly require running software that executes with privileged access on all systems, including some, such as Linux or Unix servers, where the anti-virus software itself may create more

of a risk than the actual chance of viruses attacking these platforms. Organizations should consider having tamper-resistant controls on all platforms that maintain elevated access within their networks and closely monitor them for deviation from normal behaviors.

Organizations must consider insider threats. AC-04 controls information flow between interconnected systems but provides for local-network transmission of unencrypted controlled information. An insider threat or an adversary who has bypassed perimeter defenses could intercept these transmissions, placing controlled information at risk. As a best practice, sensitive information should always be encrypted at rest and in transit and protected by appropriate restricted access controls. The use of role-based access controls or other restrictions that prevent viewing and manipulating data when not required for a user's current job should be in place. This kind of control can be implemented in a variety of ways and will be subject to interpretation by 3PAO auditors.

5.3 Unaccounted for threat models

In analyzing the technical threats and trends across the results, there are four threat models, or meta-level profiles of threat actors and their possible methodologies, that appear to be absent from FedRAMP risk management considerations. These four threat models generically encompass the technical issues identified, and are helpful to frame the way that weaknesses in compliance standards can be taken advantage of by malicious actors. Abstracting the use of specific vulnerabilities into

a threat model is a way that empowers defenders to identify and prioritize defenses against many adversaries and attacks. Below threats are described in terms of scope and motivation that undercut FedRAMP. While these focus on risks to FedRAMP certified cloud computing platforms and web hosted software, improvement for other compliance programs can also be informed by these threat models.

5.3.1 Nation-state privileged access

Security issues and gaps, as exemplified throughout FedRAMP, present opportunities for foreign nations to access the private or sensitive data of compliant organizations. This may become more prevalent within the services provided by multinational corporations that provide encrypted solutions to a global customer base.

Private keys and passwords used for encrypting data must be protected from foreign government access. Compliance loopholes that permit direct access to encryption keys and passwords could allow nation-states to bypass privacy controls. For example, a foreign government could mandate that companies generate encryption keys and store them in databases accessible to the company or the government on demand. On-demand access would obviate encryption for data at rest or in transit, allowing foreign governments to decrypt information, at will, to conduct various forms of espionage to include intellectual property theft, personnel tracking, and communication eavesdropping.

5.3.2 Corporate aggregation and monetization

The second threat model considers businesses that desire to aggregate customer information for monetization. Knowing how customers use a service, where they choose to use the service from, when customers are most likely to use a service, as well as knowing issues encountered when using a service can shape essential business decisions. Businesses can craft and deliver targeted ads, forecast inventory requirements, build security patches, or make future business investment decisions based on this data. Most customers understand this is the status quo. However, one must consider a threat model that exceeds the status quo and breaches customers' expectations of privacy by monetizing information that should be unreadable by the service provider.

In circumstances where companies identify loopholes in privacy laws or outright disregard privacy considerations, a company may be motivated to access the encrypted communications of their customers to further enrich known information. Issues such as the ones discovered in FedRAMP could allow service providers to gain compliance and still bypass encryption. The ability to access private keys, remotely access account information, and clone MFA authenticators could allow an organization to impersonate users and farm information that can assist with further monetization.

Further, unless information is encrypted end-to-end and the software provider does not hold the encryption key – a relatively rare situation – user activity may be encrypted in transit and at rest, yet still be processed by the service for targeted ads

or personalized features. None of FedRAMP's encryption controls prohibit business access to private data and may present risks if sensitive information is disclosed. While this concern is typically addressed during service contract negotiations, it is important for compliance programs to explicitly address the issue.

Even though highlighted issues relate to corporations having too much user data, one must also consider the implications of denying certain information. Some companies use the user activity sent to them in order to make software safer and more desirable for future purchases. One such example is Microsoft's use of customer stack traces, generated after crashes caused by users, to locate and fix security issues [93]. Organizations that are highly concerned about the confidentiality of their data may not allow these stack traces and other automated error reports to be automatically shared. (In the experience, many FedRAMP-compliant programs forego sharing stack traces for this reason.) This excludes some of the most likely targets of sophisticated attacks from these automated vulnerability detection programs, which may mean that exploits are detected only once they are used against other targets, potentially limiting defensive effectiveness across the user base.

The examples in this section demonstrate that compliance programs must reconcile the balance between service providers having too much user information and not enough. As is, both sides of the argument present risks to organizations that most reasonable users would not be willing to accept.

5.3.3 Security of security appliances

The third threat model considers attackers who are motivated to exploit security systems to bypass defenses and gain access to vulnerable systems. It is infeasible and inefficient for most companies to develop their own in-house security solutions or encryption mechanisms. Organizations, for the most part, rely on commercial security appliances or third party service providers for their security. Inherently, the security of these applications have wide-reaching implications.

Security controls, such as those in FedRAMP, inherently trust security applications and do not provide mechanisms for checks and balances. These applications require privileged access to data and systems without explicit oversight. Web proxies, for example, exist to reduce network bandwidth usage and provide security stand-off from the Internet – but these proxies may also have insight into all users’ web traffic. Antivirus applications prevent the execution of malicious code on workstations – but these applications may have the highest level of access to sensitive files and the core of the operating system.

The recent increase in remote work has made virtual private network (VPN) attacks even more attractive to malicious parties [168]. While these attacks are not new [153], the scale of users of these systems has greatly increased over the last few months, creating an enticing pool of targets. In the rush to move to remote work, many enterprises have neglected to ensure that their VPN servers are fully up to date, and these unpatched VPN servers provide a gateway to access corporate data that is otherwise unreachable. The VPN client software that runs on users’

workstations and laptops provides another attack surface, particularly if the client software is out of date and has unpatched vulnerabilities. Both VPN servers and clients are security mechanisms often mandated by compliance programs that provide encrypted access to corporate information; however, the lack of explicit controls on protecting these protection mechanisms from compromise place organizations at risk.

There are a few methods that may provide requisite security checks. First, monitoring and whitelisting the permitted behavior of these protection mechanisms may prevent hijacking – for example, the service-level account of the VPN server should not be allowed to access internal file shares after business hours. Site reliability engineering provides many zero-trust recommendations that may help solve similar problems [25].

5.3.4 Ignored cyber-physical systems

Digitally-connected physical safety and security controls must also be included in an organization’s security plan. As physical systems become more intertwined with information systems, technicians must actively mitigate the risks to organizations that could be posed by electric-power systems, closed-circuit television cameras, biometric scanners, or fire suppressors. Controls such as PE-13(03) from FedRAMP mandate that organizations use an automatic fire suppression capability for information systems when a facility is not continually monitored. Consider the devastation that could occur if an attacker gained privileged access to a sprinkler

system over a holiday and flooded the facility. An attacker could turn off power to essential services or expand its access within a network by using unprotected cyber-physical systems.

5.4 Comparing FedRAMP

Compared with other compliance programs for protecting taxpayer information, credit card data, and the electric grid, there are fewer issues in FedRAMP. This is attributed to the vast number of listed collaborators, frequent integration of lessons learned, and use of public requests-for-comments which allow interested parties to assess draft documents and provide recommended improvements. Additionally, FedRAMP mandates security controls based on three different levels of impact – low, moderate, and high. This acknowledges that systems and networks have different value and associated risk and should be protected accordingly. But as discussed, such flexibility sometimes comes at the expense of under-defined or ambiguous specifications.

The maintainers of the FedRAMP program appear to have recognized some of the problems highlighted here, such as varying interpretations of ambiguous specifications. FedRAMP partners with the American Association for Laboratory Accreditation to oversee its 3PAO program, and the two organizations have undertaken a process to review and update the 3PAO training requirements and to evaluate the technical competence of 3PAOs. Additionally, these organizations have released standards that include a requirement that 3PAOs who operate internationally show

how they minimize the risk of foreign parties interfering with the FedRAMP certification process. A 3PAO who is operating under the influence of a foreign nation may, for instance, be more likely to be lenient in reviewing controls in certain areas that might make it easier for the foreign partner to attack systems or exfiltrate sensitive data.

5.5 Discussion

Compliance programs like FedRAMP often contain security issues or gaps that can allow risks to persist even in compliant organizations. This study provides tips and recommendations to help businesses, security researchers, and government organizations mitigate many of these risks. First and foremost, organizations and security professionals must remember that compliance establishes only a minimum level of protection. Compliance may be helpful (even required) to achieving the end goal of protecting an individual or organization's goals and assets, but it is not sufficient.

Organizations should understand the impact that compliance programs may have on overall security. All organizations should audit the compliance standards they follow to identify gaps that are relevant to their specific requirements and develop mitigating strategies accordingly.

As a whole, the United States government must adopt a mechanism that permits always-open request-for-comment periods for compliance programs that allow security researchers to identify weaknesses and recommend fixes. Additionally, there

is a need for faster revision of compliance documents to maintain relevance with emerging technologies and threats. These revisions should permit grace periods, to allow organizations to migrate towards new controls without facing sanctions for non-compliance.

The coronavirus pandemic and the rapid adoption of work-from-home solutions such as Zoom only highlights the need for strengthening compliance programs. Hopefully this work can help make organizations, businesses, and citizens aware of potential security issues until the federal government implements more secure and flexible options for compliance.

Chapter 6: Implementing Security: Humans factors in Incident Response Readiness

In this chapter, I present research that focuses on the end-to-end usability of incident response playbooks within an enterprise environment. Proactive security measures require professionals to implement security controls that will help mitigate adversarial impact. Incident response playbooks present security practitioners with a structured action plan for incident response efforts after a threat is underway. These pre-planned actions are intended to help reduce stresses that technicians may face during a security incident, allow them to gain momentum during response efforts, and ensure organizations are prepared for likely adversarial situations.

Although playbooks are a common practice in the security industry, they have not been systematically evaluated for effectiveness. This chapter takes a first step toward measuring playbooks, using two case studies conducted in an enterprise environment. In the first study, twelve security professionals created two playbooks each, using two standard playbook design frameworks; the resulting playbooks were evaluated by experts for completeness and correctness. In the second, five personnel use the created playbooks in no-notice threat exercises within a live security-operations center.

Playbooks, in some cases, do simplify and support incident response efforts. However, playbooks designed using the frameworks examined often lack sufficient detail for real-world use, particularly for more junior technicians. This study shows that incident response playbooks may be valuable tools for increasing preparedness against a threat, but often require extensive planning and threat modeling to determine which playbooks should be prioritized for development and customization. Additionally, baseline security mechanisms may constrain or alter playbook design efforts to ensure compatibility with applicable compliance programs.

6.1 Setup and preliminaries

This section details the playbook frameworks evaluated, the two partner organizations, and the two incident-response scenarios selected as targets for playbook design.

6.1.1 Selected frameworks

This study uses two frameworks that have U.S. government support and offer freely available guides and examples.

IACD. The IACD framework was created by the U.S. Department of Homeland Security, National Security Agency, and Johns Hopkins Applied Physics Laboratory to leverage automation within incident response [104]. The defining feature of IACD playbooks is a visual flowchart capturing essential response actions for both humans and automated systems to take (Figure B.4 in the Appendix).

The IACD framework breaks playbook design into 10 steps. (Section 6.1.3.1 provides a running example in greater detail.) The first step is to identify the initiating condition: the event or situation triggering playbook use (e.g., a database breach) and how that event is detected (e.g., an automated email alert sent to an administrator). The second step involves listing all possible actions that could occur in response to the initiating condition, typically via mind mapping. Practitioners should reference existing best practices to identify possible actions. Next, playbook designers designate each identified action as required or optional. For example, generating a written report that details the incident from beginning to end — which may provide invaluable insight after the event but does not contribute directly to response efforts — should be labeled optional. Steps 4-8 involve grouping actions by function, ordering required actions sequentially, and interleaving optional actions where appropriate. The designer produces a diagram showing these ordered relationships. In step 9, the designer verifies that the playbook terminates either in a desired end state or in a new initiating condition that flows into another playbook. The final step ensures the playbook satisfies applicable regulatory controls and requirements.

NIST. The NIST Computer Security Incident Handling Guide (hereafter: the NIST framework) focuses on quick recovery after a security incident [44]. Using this framework, designers break a security incident down into three phases and create playbook content for each. The preparation phase occurs before an incident and requires analysts to identify critical assets that must be protected from a particular threat. Play-

book content for the detection and analysis phase should help defenders identify the incident's entry point, breadth of impact, potential consequences, and containment methods. Phase three content — containment, eradication, and recovery — should guide defenders in patching or isolating the attacker's entry point and other similar potential entry points, increasing monitoring, and safely bringing services back on-line. Each phase of a NIST playbook should emphasize communication and metrics tracking: ensuring essential personnel are informed, victims are notified, and the scope of impact is thoroughly documented. While playbook designers may or may not deem communications as required actions in IACD playbooks, communication is required throughout NIST playbooks.

Unlike IACD, NIST does not typically result in a visualization of response actions (although it could). Instead, NIST playbooks typically provide detailed textual descriptions, intended to be drawn from institutional procedures or best practices. Section [6.1.3.2](#) provides a detailed running example.

6.1.2 The partners

To evaluate playbook frameworks' usability within organizations that had not previously used them, researchers partnered with two organizations specializing in digital security. For anonymity, they are referred to as the network defense center (NDC) and the security development team (SDT).

NDC manages networks spanning multiple countries and 600 user accounts, with a service-level agreement to maintain availability levels at or above 98.9% while

securing highly-sensitive customer intellectual property. NDC had 12 employees during the first case study and 13 during the second study.

SDT develops secure applications for nearly 1500 worldwide customers, often building custom solutions for niche requirements. SDT employs 28 developers.

Both NDC and SDT have mandates to secure their development and production environments from malicious attacks, insider threats, and natural disasters. Both organizations have personnel with a range of security experience: a few entry-level and the majority with more than 10 years' experience.

Prior to this study, one co-author of this paper served as a supervisor within both NDC and SDT, enabling deep understanding of both organizations' missions, cultures, customers, and risks. This co-author observed NDC's response to three security incidents within one year. The technicians' responses were ad-hoc, rather than drawing on predetermined plans, policies, or procedures. By the start of the first case study, this co-author was no longer affiliated with either partner.

6.1.3 Selected scenarios

As playbooks are designed to address specific incident scenario, two scenarios are used in the case study. Leaders from NDC and SDT collaborated to select two scenarios from the MITRE ATT&CK database, using the following three criteria: (1) both organizations could realistically encounter them; (2) each organization should be able to quickly and consistently respond to them at any time; and (3) neither organization had a standard policy or procedure in place to handle them.

Brute-force login attempts [140] and valid credential compromises [142] were the two selected scenarios.

6.1.3.1 Brute- force login attempts

In a brute-force login attack, an adversary attempts to gain unauthorized access by guessing commonly used or randomly generated passwords. This study focuses on protecting user-level domain accounts from locally-originating attacks. Below details some of the essential tasks associated with detecting, responding to, and eliminating brute force attempts from within the network using the IACD framework.

The initiating condition is the detection of multiple password-guessing attempts against one or multiple systems. A centralized log repository must continuously audit and correlate login failures from across the network. If a brute-force pattern is detected, the system should generate an alert (e.g., a dashboard push notification or email to a technician).

Required actions, in sequential order, might include: identify the system(s) being attacked; identify the potential attack source; isolate source and/or victim nodes; install new sensors for traffic monitoring; identify compromised accounts; conduct root-cause analysis; perform root-cause mitigation; and restore accounts/services. Two optional action groups might be prioritizing assets (determining which resources are most important to isolate first) and producing reports (helping responders understand the situation and make better decisions).

In IACD Step 9, the playbook terminates in a desired end state: the root cause has been patched and affected services and accounts have been restored. The final IACD step is to validate that the playbook satisfies regulatory controls and requirements, such as log retention policies.

6.1.3.2 Valid credential misuse

Valid credential compromise can occur, for example, when a database breach reveals credentials from one account that can be reused at another site. This study focuses on protecting user-level domain accounts from local abuse.

Next is a sample NIST playbook that uses *honeypots* — usernames and passwords for valid but fictional accounts — to detect credential misuse [112].

The preparation phase includes the creation of honeyword accounts, ensuring security team members understand the significance of the honeywords, and deploying an automated log-event parser to scan for login attempts associated with the honeyword account and generate an alert if found.

The detection and analysis phase starts when a human analyst receives an alert (e.g., a dashboard push notification or an email). The analyst should then investigate breadth of impact; for example, if the honeyword was created on a domain controller, then the analyst may assume there has been a compromise of all accounts on the domain controller.

The final containment, eradication, and recovery phase involves root-cause analysis to determine how the domain controller was initially compromised and

generate a “fingerprint” to check for similar compromises on other systems. Next, all affected accounts must be denied access until they have changed their password. Affected users and compliance entities (as applicable) must be notified of the breach. Finally, the incident must be fully documented, and there may also be regulatory requirements for follow-on security assessments.

6.2 Playbook design and evaluation

This section details the first of two case studies exploring the usability of playbook frameworks. Participants from NDC and SDT were familiarized with the IACD and NIST frameworks and asked to each design two playbooks using the two frameworks and the two selected scenarios. This permitted the measurement of participants’ perceptions of the process, and external experts evaluated the designed playbooks for thoroughness and accuracy. Although most participants considered the frameworks reasonably easy to use, about half of the designed playbooks were rated as insufficiently detailed for real-world use.

6.2.1 Method

This case study assesses participants’ perceptions of the usability of the frameworks as well as whether the resulting playbooks would be usable in a real-world setting. In this context, usability is defined in terms of learnability (ease of first-time use), efficiency (timely task completion), errors, and satisfaction [157]. A mixed-methods study, as shown in Figure 3.2, addresses these questions.

This case study occurred from September through December 2019 and was approved by an ethics review board. To protect the participants and partner organizations, sensitive information, including job descriptions and identified vulnerabilities, is redacted and generalized.

6.2.1.1 Recruitment

Researchers partnered with NDC and SDT to recruit employees performing daily security functions. Leadership from both organizations announced the study during group meetings, describing the motivation and goals while emphasizing that participation was voluntary. Employees were told that participants would be introduced to new techniques that could be useful in their work, and that playbooks from the study would be adopted into daily practice. Employees and contractors were permitted to participate during regular work hours but were not otherwise compensated. NDC/SDT leaders emphasized that participation in the study would have no impact on performance evaluations.

6.2.1.2 Playbook design

Participants received group-based, in-person instruction on using IACD and NIST frameworks, using an exemplar scenario not included in the main study: responding to spearphishing links [141]. These 30-minute introductory sessions were based on fundamentals from adult learning research, including learning through examples and hands-on implementation [21,118]. Each group had the same instructor,

an author who possesses five years of experience designing incident-response scenarios for organizations. The instructor communicated this experience to each class to establish credibility.

Next, each participant designed two incident response playbooks, one for each threat scenario (Section 6.1.3), using publicly-available references and relevant entries in the MITRE ATT&CK database [139]. The assignment of frameworks to scenarios was randomized, as well as the order of tasks, to mitigate ordering effects and other biases. Each participant used each framework and each scenario once.

After designing each playbook, participants completed an online survey based on the System Usability Scale (SUS) [36]. Next were open-ended follow-up interviews, averaging half an hour, with each participant. Questionnaires and interview guides for all segments of both case studies are given in Appendices A.3 and A.4 respectively.

For all surveys and interviews in both case studies, two co-authors jointly analyzed all open-ended questions using iterative open-coding [199], building the codebook incrementally. All disagreements were resolved by mutually refining codebook definitions and then re-coding responses accordingly. This process continued until all responses were coded, resolved all disagreements, and the codebook was stable.

6.2.1.3 Playbook evaluation

Participant perceptions are valuable, but to fully understand usability, it is also important to measure error rates. Three expert evaluators — each with extensive experience with playbooks in enterprise environments — were asked to assess whether participants’ playbooks were valid and sufficiently detailed for use during incident response.

For each playbook (anonymized before review), evaluators completed an online survey with closed- and open-ended questions about whether the playbook accomplishes its goals, contains enough detail to be implemented in a real environment, and contains any likely sources of error.

6.2.2 Participants

ID	NBF	IBF	NCM	ICM
P1	E1, E2	–	–	E1, E2
P2	E3	–	–	E3
P3	E1, E3	–	–	E1, E3
P4	–	E1, E2	E1, E2, E3	–
P5	–	E3	E3	–
P6	–	E1, E3	E1, E3	–
P7	–	E2	E2	–
P8	–	E1	E2	–
P9	E2, E3	–	–	E2, E3
P10	E3	–	–	E3
P11	–	E3	E3	–
P12	–	E2	E3	–

Table 6.1: Evaluation coverage of designed playbooks. The columns show combinations of framework to scenario, denoted by N: NIST, I: IACD, BF: Brute Force, and CM: Credential Misuse.

Because of limited time availability, each evaluator examined a subset of play-

books. To ensure consistency, 10 of the 24 total playbooks were assigned to two different evaluators. In the event of disagreement on any key attributes, the third evaluator was asked to review the playbook, and their response was used to break the tie; one playbook required a third evaluation. The remaining 14 playbooks received a single evaluation each. Table 6.1 shows the distribution of evaluations.

6.2.2.1 Limitations

All field studies and qualitative research should be interpreted in the context of their limitations.

The recruitment materials explained the purpose of the study. This may have resulted in self-selection bias: participants most interested in the study topic opting to participate.

The results may also exhibit demand characteristics, in which participants are more likely to respond positively due to close interaction with researchers [96, 167, 212]. Online surveys help mitigate this and promote candid feedback; additionally, both positively- and negatively-framed questions ensure participants could provide both perspectives.

NDC and SDT use organizational structures and technological resources common to many security-conscious organizations of similar size; however, specific organizational characteristics may inhibit generalizability. This is an inherent limitation of an in-depth field study. Nonetheless, the results may illuminate systemic issues that organizations need to account for when adopting playbook frameworks.

ID	Org	Role	Exp (yrs)	Study Phase	Order
P1	NDC	Manager	11+	D	NBF:ICM
P2	NDC	Technician	0-4	D, IR1/2/3	ICM:NBF
P3	NDC	Manager	11+	D, I, IR1/2	NBF:ICM
P4	NDC	Manager	11+	D, IR1	IBF:NCM
P5	NDC	Manager	11+	D	NCM:IBF
P6	SDT	Manager	11+	D	NCM:IBF
P7	SDT	Technician	11+	D	IBF:NCM
P8	NDC	Technician	5-10	D, IR2	NCM:IBF
P9	SDT	Technician	11+	D	ICM:NBF
P10	SDT	Technician	11+	D	NBF:ICM
P11	SDT	Technician	5-10	D	IBF:NCM
P12	SDT	Manager	11+	D	ICM:NBF
P13	NDC	Technician	0-4	IR3	–
E1	–	Senior Mgr	11+	E	–
E2	–	Senior Mgr	11+	E	–
E3	–	Senior Mgr	11+	E	–

Table 6.2: Participant and expert demographics. The columns show participant identifier, employer organization, work role, years of experience, participation phase of the study, and order of playbook creation. The abbreviations in the fifth column represent design, evaluate, implement, and incident response exercise. Abbreviations in the sixth column represent NIST, IACD, brute force, and credential misuse.

As disclosed in Section 6.2.1.1, one co-author had previously served as a supervisor within NDC and SDT. This case study started five months after the co-author had departed these organizations; eight participants were new hires after the co-author had departed. Additionally, all study execution decisions were made with close NDC and SDT leadership oversight.

This study is not a direct comparison of two frameworks, but rather an observational case study attempting to identify benefits and shortcomings for each and for playbook frameworks in general. The sample (n=13) is small, but it represents 100% of NDC’s full-time workforce (58% total workforce during this case study) and 21% of SDT’s employees.

For each qualitative finding, participant count provides context. However,

participants who did not mention a specific concept when responding to survey or interview questions may simply have failed to state it, statistical hypothesis tests are not appropriate for these questions.

To limit biases, the partners did not have existing policies or procedures to handle the selected incident response scenarios. This forced participants to build plans around technologies not yet in place, which may have contributed to a lack of detail in many playbooks (see results). However, this limitation is also realistic: Evaluator E2 said his organization often faces similar situations, and IACD cites the identification technology gaps as a key function of playbook design [104].

6.2.3 Results

Below are the results of the first case study, including participant demographics, participant feedback on the playbook design process, and expert evaluations of the accuracy and completeness of the designed playbooks. Overall, participants reported a somewhat favorable perception of playbook design frameworks and their ability to assist with incident response efforts. In general, they appreciated thinking proactively and identifying solutions to realistic threats they might face in the future. However, expert evaluation suggested about half of designed playbooks were insufficient.

6.2.3.1 Participants

In total, 15 people participated in this case study, including 12 NDC/SDT employees who designed playbooks and three expert evaluators (Table 6.2). Qualitative research best practices recommend 12-20 participants for data saturation in thematic analysis [84]. The sample represented 58% of NDC's workforce and 21% of SDT's workforce at the time of the study. Prior to the study, all design participants (P1-12) said they knew that playbooks were an industry best practice, but none had used a playbook to respond to an incident. All design participants had completed at least one year of entry-level, on-the-job training for their job role; overall, they averaged 10.6 years of digital security experience.

Three expert evaluators (E1-3) were recruited via email, based on participant contact lists aggregated during previous research. Each has extensive experience with designing, implementing, and using playbooks: E1 is the director of security operations center with more than 300 employees; E2 is the Deputy CISO of one of the largest cities in the U.S.; and E3 is the CISO of a major U.S. financial institution. Collectively, they averaged 16.7 years of digital security experience.

6.2.3.2 Usability metrics

Score distributions from the survey questions show user satisfaction; this survey extended the SUS with four additional questions designed to elicit perceptions of usability for others, following guidance from Brooke [35, 36]. Participants rated the IACD and NIST frameworks 63.5 and 67.7 out of 100 respectively ($\sigma = 16.3$,

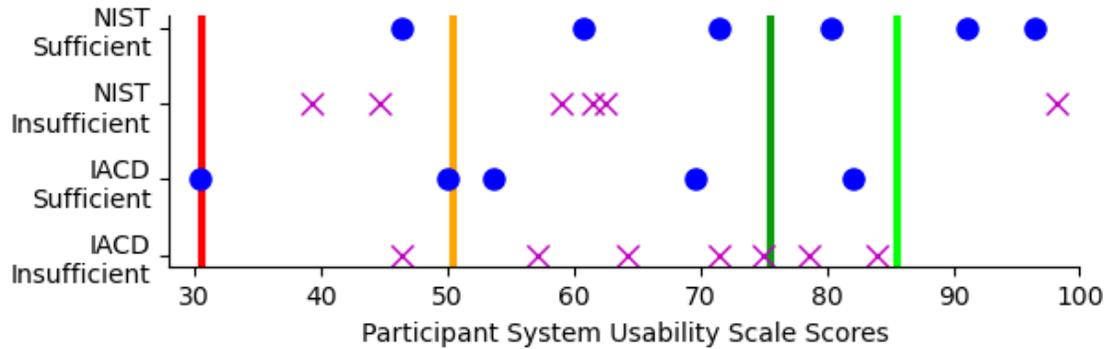


Figure 6.1: A comparison of participant usability perceptions and error evaluations from experts. The vertical lines (left to right) indicate poor, okay, good, and excellent usability.

20.1; see Figure 6.1). This corresponds to a rating of “okay” on a standard scale from poor to excellent [35].

All participants completed both assigned tasks. Average completion times were 32.8 minutes ($\sigma = 6.1$) for IACD and 42.1 minutes ($\sigma = 7.4$) for NIST, are considered acceptable for learnability for a complex task of this type. There were no observed, noticeable differences in task completion time based on order.

6.2.3.3 IACD feedback

Participants identified a variety of positive and negative features of the IACD framework.

Visualization is a key benefit. Most participants (n=10) identified the graphical depiction of required tasks as IACD’s most positive attribute. P1 and P4 both indicated that visually distinguishing between human and automated tasks helped them focus on their roles during incident response, better understand how to leverage

automated systems, and ensure they are compliant with mandatory controls.

Playbooks can help make up for lack of experience. P10 noted that IACD “helped me organize my thoughts and guided me through problem-solving”; even though he had never responded to the given scenario in a real event, he believed the framework was helpful in eliciting the necessary steps to handle the situation. P12 said “it allows organizations with a preponderance of junior defenders to execute something without the guidance of a senior defender,” especially when a speedy response is critical.

Even experienced professionals had difficulty with some terminology and instructions. Several participants (n=7) had difficulty grouping similar activities and functions, a core step that many following steps build on. IACD does not provide a list of common groups to choose from, requiring users to determine their own groupings. P8 indicated that it took him approximately one hour to develop a playbook, and a majority of that time was spent attempting to identify appropriate groupings to use. P9 had to reference the guide frequently when using the framework, and three other participants said they were never confident they provided enough information.

Resulting playbooks did not have enough detail or account for enough contingencies. Two participants felt the resulting products were too abstract for technicians to follow during incident response events. P7 said “the diagram is nice and easy to follow, but probably also needs a document to go with it explaining

in more detail what each action entails.” These comments foreshadow many of the difficulties junior technicians during playbooks utilization (Section 6.3.2.3).

P4 wanted to see more emphasis placed on loops (and their exit conditions), parallel activities that can be conducted simultaneously by both humans and systems, and accounting for multiple possible end states based on conditional transitions. P12 felt similarly: “The point of a playbook is to recapture the initiative from the attacker by having a several iterations of the OODA loop unrolled,” but he felt the IACD did not account for multiple paths. P12 suggested more modeling akin to the cyber kill-chain framework [228].

Identifying the initiating condition is most important. Ten participants agreed the “identify the initiating condition” step was the most important. All 10 described this first step as setting conditions for all follow-on steps, and noted that failure to recognize the initiating condition would significantly delay or even prevent incident response; this again foreshadows complications observed in the second case study.

P11 mentioned that the initiating condition will be in the playbook “table of contents,” which technicians will reference when identifying a playbook for responding to an event. The technician will therefore “need to be able to correlate what they believe to be occurring with how you laid out the [initiating condition] entry into the playbook.” All participants used five or fewer words to describe their initiating conditions; playbook designers must use concise yet descriptive terms to cue defender actions.

Identifying regulatory requirements is least important. Eight participants indicated that “identify regulatory controls and requirements” was the least important section of the playbook, noting that regulatory compliance was not relevant to their job role or was someone else’s responsibility. P9 said: “Compliance is less of an issue than actually solving problems.” This sentiment aligns with prior work suggesting technicians view compliance as inhibiting security [17, 45].

6.2.3.4 NIST feedback

As with IACD, participants identified benefits and drawbacks to the NIST framework.

The framework was easy to understand. The most prevalent positive feedback was that NIST playbooks were easy to understand (n=5). P6 stated NIST was “[v]ery clear on what steps I needed to follow and what outputs are expected after each step,” and P4 noted that the “[r]esulting text could be passed on to anyone to help them perform initial triage.”

The framework prompted for detail. Participants liked that the NIST framework prompted them to include as much detail as desired for response actions. They felt that fine-grained details would reduce uncertainty during response actions taken by junior defenders, rather than requiring novices to figure out on the fly how to implement abstract instructions. Two self-identified managers said the detail-oriented

design of the framework would help security engineers to understand and implement controls or systems required by the playbook. For example, the NIST framework prompted these two participants to describe in detail the expected content for an alert email, providing guidance to security engineers who would be tasked with building or configuring the alert system.

The framework supports proactive planning. Two participants appreciated that the NIST framework allowed them to think about problems before a full-blown crisis occurred. P6 noted that NIST offered him an option to “identify possibly solutions, identify gaps in technology, and have at least an initial plan in place for handling the situation.”

NIST playbooks may be less useful for novices. Participants (n=5) were concerned that it might be difficult for a novice to quickly orient themselves to a NIST playbook given the lack of visual aids, reflecting the importance of accommodating various learning styles [118]. “It’s all just a bunch of words. During a crisis, you need something concise and clean to follow,” P8 stated, after using NIST but prior to using IACD. “I liken it to if IKEA’s instructions were text only. They wouldn’t be as valuable.” Participants suggested adding a headline-style title, executive summary, and visual cues to NIST playbooks.

Even more detail may be needed. Five participants wanted the NIST framework to require even more fine-grained detail. They felt the NIST framework was too open-ended, and would have liked the framework to prompt for exact commands

that an analyst should execute, rather than requiring them to reference another guide or have the commands memorized. Two other participants felt the framework did not adequately prompt for decisions and branching plans to account for incident variability. Two participants believed NIST did not account for partially-complete tasks. P11 felt that if there is not a check on task completion, it could result in unnecessary actions just because it is in a playbook or missed opportunities to do something in parallel. These comments were similar to comments about the IACD framework, suggesting that the participants were looking for detailed, pre-planned responses that account for branching investigation paths.

Examples and instructions were again a challenge. Two participants wished for multiple NIST playbook examples as a reference during the design process. P4 noted that the NIST framework was too abstract in places, making it hard to understand what is required for each step.

Detection and analysis is most important, but no strong consensus. A plurality of participants rated “detection and analysis” the most critical response step (n=5). All five indicated that knowing a security event is underway and they need to take action, even if it is a false positive, is invaluable for a defender. P9 stated: “if you miss it, your plan for responding is useless.” This finding parallels the importance participants placed on the initiating condition in IACD playbooks and suggests the importance of user interfaces that deliver critical information without overwhelming the analyst [30].

For similar reasons, participants narrowly scored containment, eradication, and recovery as the least important phase in the NIST framework (scoring 4% lower than the “preparation” phase). Three participants said this step is only important if you have detected a problem and successfully investigated to confirm the incident is a true positive.

6.2.3.5 Expert evaluation

Overall, the playbooks participants created lacked sufficient detail and suggest that amount of experience did not have a significant impact on accuracy. The expert evaluators assessed six of 12 IACD playbooks and five of 12 NIST playbooks as insufficiently detailed for use during incident response; when asked if the playbook would be likely to adequately respond to the associated scenario, IACD playbooks averaged 2.71 ($\sigma = 1.40$) out of 5 while NIST averaged 3.0 ($\sigma = 1.57$).

Comparing participants’ SUS scores to evaluator judgments of sufficiency helps demonstrate how perceived usability mapped to effective outcomes (Figure 6.1). For NIST playbooks, sufficient playbooks were generally associated with higher SUS scores than insufficient playbooks, suggesting that participants understood whether or not they were succeeding. For IACD playbooks, however, there is no clear relationship. This may indicate an important mismatch between perceived success with the framework and outcomes.

Next are some common themes observed by the evaluators across playbooks from both frameworks.

Missing “implied” tasks. The experts noted that many playbooks were missing what one evaluator referred to as “implied” tasks: necessary for the defensive strategy to succeed, but not codified within the playbook (IACD=4, NIST=3).

As one example, E1 noted that P3’s IACD playbook was missing investigative steps necessary to confirm whether an alert is a true or false positive. To do this, the analyst must determine (in the case of credential misuse) where the login attempt originated from and why it occurred. In particular, when receiving an alert related to a honeyword, the analyst should check whether an administrator is performing a standard periodic login to ensure the account does not expire, before assuming a valid attack; this step was not included in the playbook. Further, E1 suggested that if the login attempt does appear to be a true positive, the analyst should then take steps such as searching Internet forums for the honeyword to investigate how and when the credentials were leaked.

Imprecise language may cause delay. Evaluators identified three IACD playbooks and four NIST playbooks with imprecise language or instructions that could delay response efforts. For example, some playbooks rely on client applications running on all workstations. If a technician remotely pushes commands to the clients without first ensuring all clients are running, the technician may have to re-run the commands once they identify abnormalities in the results (wasting minutes or even hours).

Missing essential communications. The experts agreed that most playbooks

missed at least some essential communications; the six lowest-scoring IACD playbooks and six lowest-scoring NIST playbooks all lacked this information. A sufficiently detailed playbook should include specific information (name, position, email address, and/or phone number) about who to contact in different circumstances. Expert E2 compared this to a bomb threat checklist, which allows users to collect essential information and communicate it to the right people (e.g., calling 911) [58]. Business continuity and disaster preparedness experts recommend that incident responders have essential contact information, as well as fill-in-the-blank forms for communicating must-know information, readily available in case of crisis [224].

Missing humans in the loop. A few playbooks (IACD=2, NIST=2) relied too heavily on automation rather than including humans in the decision process. For example, several playbooks included automatic account disabling during a brute-force attack; while superficially sensible, this could result in locking out administrators, hampering the response. E2 emphasized that when decisions affect critical services or accounts, a human decision maker must be involved.

Too linear. Evaluators noted that, especially in the case of IACD, playbooks did not account for parallel actions that humans and automated systems could accomplish concurrently, increasing efficiency and reducing overall investigation time. For example, while a technician responding to a brute-force attack searches network traffic for any attempts occurring in real time, automated systems could query log data and locate attempts from the prior hours or days. Steps 3 and 4 of the IACD frame-

work ask the playbook designer to order tasks sequentially and group by functions, which may inhibit designers from planning parallel tasks.

Some include details and best practices. The experts did recognize multiple playbooks (IACD=2, NIST=4) with high-quality, fine-grained detail. P10 included references to best practices such as data loss prevention that would prevent users with insufficient privileges from accessing sensitive data in the event of credential misuse. Several other playbooks detail steps for determining what information, if any, was stolen from the network in the event of a successful brute force attack or credential misuse.

One playbook was not just incomplete but incorrect. The evaluators identified one playbook (P7, NIST, credential misuse) as potentially impeding a technician's ability to respond to the incident. E2 believed this playbook's response events were ordered incorrectly, which could lead an incident responder to miss valuable information in one step that would be required later. Both E1 and E2 noted that this playbook lacked root-cause analysis and therefore could not lead to a successful resolution. Without root-cause analysis, it is possible for an attacker to regain access or spread throughout a network undetected while responders focus on inconsequential details.

After reviewing these expert findings, researchers conducted a follow-up interview with P7, a software developer with more than 10 years of experience in reverse engineering and secure code development but little exposure to network de-

fense or incident response. P7 said they understood the scenario and gave the NIST framework a slightly-below-average SUS score of 62.5. However, the playbook framework guidelines and reference material about each tested scenario from the MITRE database could not make up for P7's lack of relevant experience. While it might be unsurprising that a secure software developer struggled to develop an operations playbook, none of the playbook literature specified prerequisites or qualifications for designing playbooks [28, 104, 150, 200].

6.2.4 Summary

This case study suggests that the two frameworks have only moderate usability. Although all participants completed each playbook design task in under 45 minutes, only about half were considered by experts sufficiently complete and correct for real-world use. Participants found the idea of playbooks, and some individual features of the two frameworks, useful, but also identified key weaknesses. In particular, participants placed high importance on identifying the triggering action (in both frameworks) and appreciated the visual process overview associated with IACD, but also wanted detailed checklists. Participants appreciated that the NIST framework came closer to prompting for this amount of detail, but wanted the framework to go even further, such as requiring exact syntax for system queries and commands.

6.3 Playbook implementation and use

The second case study is detailed next, which investigates playbook performance in practice. Examined are two facets of using a playbook: implementing processes to support the playbook’s incident detection and response plan, and then executing the response plan during an incident.

6.3.1 Method

This case study, which took place from December 2019 to March 2020, uses two playbooks that received high scores from the expert evaluators and spanned both frameworks and scenarios: P4’s IACD playbook for brute-force login attempts and P11’s NIST playbook for credential misuse. One participant engineered security solutions based on the playbooks, and then evaluated their usability during three controlled insider-threat events. Selecting high-scoring playbooks help approximate a best-case scenario for playbook use. This study assesses efficiency, errors, and satisfaction to measure usability.

The study did not obtain legal approval to modify network monitoring solutions at SDT, so no SDT employees actively participated in this second case study. Despite this, it is considered acceptable to use P11’s playbook (which was designed for SDT) within NDC, because (1) it scored well in the expert evaluation and was noted for containing fine-grained detail, and (2) the playbook could easily be implemented as-is within NDC, because neither organization had any pre-existing solution in place, so implementation could start from a blank slate. As in the first study

(Section 6.2), NDC leaders informed potential participants about the study, while emphasizing its voluntary nature. As before, participants were allowed to participate during work hours but not otherwise compensated; as before, participants were assured that participation (or not) would have no effect on performance evaluations.

6.3.1.1 Playbook implementation

In this phase, one participant implemented new technical controls, based on the requirements from the selected playbooks, to detect and respond to brute-force attacks and the misuse of valid credentials within their live network.

It was not feasible for more than one participant to perform implementation while interacting with the live network; however, this phase was necessary to enable evaluation during controlled events (Section 6.3.1.2). NDC leaders nominated one participant for this phase, and that technician subsequently volunteered. As such, findings are not generalized from this process, but observation comments are provided from the (previously unexplored in the literature) implementation process.

After the participant implemented the controls specified in the two playbooks, they completed a survey about the experience (Appendix A.3) and conducted an in-depth follow-up interview (Appendix A.4). This survey and interview questions were grounded in the technology acceptance model, but used both positive and negative framing to mitigate social desirability bias [55, 62, 123].

6.3.1.2 Playbook use during incident response

The main goal of this case study was to evaluate usability of the selected playbooks during actual incident response. Given the unpredictability of actual attacks, we worked with NDC leadership to conduct no-notice incident response exercises that would trigger playbook use. Similar approaches have been described in compliance programs, but to researchers' knowledge have never been used to investigate playbooks [75].

After the completion of the first case study (Section 6.2), NDC leadership informed all of their employees they would begin using the selected playbooks as part of their daily duties. Copies of each playbook were provided to each participant, and also placed in a binder in NDC for easy access. Each technician received a 30-minute orientation by NDC leaders on how and when to use the playbook. Technicians were also asked to review the differences between the playbook they had designed in the first case study and the ones that were selected for use. Additionally, technicians were required to verify their use of the playbook in a logbook at the beginning and end of each shift.

In order to maximize ecological validity, technicians were not informed that the study would include incident response exercises testing the playbooks. This deception is described in more detail below.

NDC leaders identified one employee as a trusted agent to simulate the insider threat. Researchers then coordinated directly with the trusted insider to schedule the simulated attacks; neither NDC leadership nor technicians received any advance

notice of when they would occur. Researchers triggered three no-notice incident response exercises on December 2, 2019 (IR1, brute-force); January 13, 2020 (IR2, credential misuse); and March 2, 2020 (IR3, credential misuse) to evaluate NDC’s ability to use the playbooks over time.

To initiate the brute-force attack, the trusted agent used a script to rapidly attempt logins against actual user-level domain accounts throughout the enterprise, using randomly generated passwords. During the brute force attack, the trusted insider executed 50 total login attempts against two domain accounts. To initiate the credential misuse attack, the trusted agent successfully logged into a designated user-level domain account configured as a honeyword account.

After each exercise, we asked each participant to complete a survey about the experience (Appendix A.3). Researchers conducted in-depth follow-up interviews with each participant and with the trusted insider (Appendix A.4). Researchers also reviewed NDC network and system logs related to the exercises.

For this study, incident response efforts taking less than 140 minutes are considered to be a success. According to a CrowdStrike analysis, this would be less than the time period required for access expansion by many nation-state threat actors and criminals [54].

Ethical considerations. No-notice exercises simulating real-world attacks carry several potential risks: they may create unnecessary stress for participants or cause senior personnel to make unnecessary decisions based on a fictional threat. To mitigate these risks, researchers directed (in consultation with NDC leadership) the

trusted insider to immediately inform participants who detected the event that it was an exercise. All written and verbal communications from that point forward were prepended with “EXERCISE” to indicate that it was not a real event; this is a common exercise practice. Although participants were notified that the incident was an exercise, they were not informed that it was specifically connected to the playbook study.

Further, NDC leaders agreed not to consider participants’ performance in the exercises (for good or bad) in annual performance reviews, in order to treat the exercise as a learning opportunity to improve institutional practices. Finally, researchers understand that responding to a controlled event may detract from NDC’s ability to respond to an actual threat or security event. To mitigate this, events occurred only on days when NDC was fully staffed. Only 2-3 participants engaged in each response effort. As is standard in deception studies, after the final exercise, researchers debriefed participants, explained the true nature of the study, and provided them with an opportunity to withdraw their data from the study; no participants withdrew. This study was approved by an ethics review board.

Limitations. Best practices recommend interviewing 12-20 participants for thematic analysis data saturation [84]. Due to security and legal concerns, researchers were only able to conduct this case study with one partner organization and five participants. Researchers provide anecdotal observations from this unique opportunity to observe playbooks in use but do not attempt to generalize the findings. Ethically, it was necessary to inform participants (after initial threat detection) that

they were participating in an exercise. While this may somewhat degrade ecological validity, it does not impede our primary objective: evaluating playbook use without prior notice.

6.3.2 Results

Below are the results for the usability of playbook frameworks when (1) implementing new security controls based on a playbook and (2) utilizing a playbook during an incident response event. These results are based on observations, survey answers, and logged digital security artifacts throughout the network. Reported are participant demographics, participant feedback on implementing technical security controls, and participant feedback on using playbooks during three incident response events. These findings provide the first structured evaluation of playbook usability from within a live security operations center.

6.3.2.1 Recruitment

Overall, five people participated in the second case study: one who implemented security controls based on the selected playbooks and participated in two exercises, one who participated in three exercises, and three who participated in one exercise each (Table 6.2). Four participants had also participated in designing playbooks for the first case study; P13 joined NDC in late February 2020 and only participated in the final incident response event (IR3). As with the first case study, all participants knew about playbooks as an industry standard, but none had used

a playbook to respond to an incident. Participants averaged 8.2 years of digital security experience.

6.3.2.2 Playbook implementation

Participant P3 implemented the security controls called for by the two selected playbooks. P3 was the most experienced defender at NDC (18 years of hands-on and management experience). P3 did not design either of the two playbooks selected for implementation; they spent approximately 30 minutes becoming familiar with the playbook requirements before implementing the requisite security controls.

After assessing the playbooks, P3 determined that all of the requisite logging mechanisms (e.g., account login failures) and recorded network traffic already existed; all that was needed was a way to aggregate this data and correlate events to obtain meaningful information. After a three-week acquisition and change-oversight period (detailed below), P3 created within one hour a new alert dashboard and a data-analysis plan to populate the dashboard with events. For the first time, NDC had a real-time system to continually monitor the network for brute-force attacks and credential misuse. The dashboard is visible on a large monitor displayed in the front of the NDC workspace and is accessible from each analyst workstation. Once either of the two scenarios is detected by the automated system, the dashboard shows an alert that investigation is needed.

Oversight requirements, change control, and purchasing — mainly associated with buying new equipment capable of storing and processing required amounts of

network data — added about three weeks to the implementation process; potential delays of this kind should be taken into account when planning to adopt playbooks and new security controls. P3 suggested that playbooks explicitly include implementation requirements such as equipment specifications and change control procedures to make this more transparent.

Implementation feedback. Overall, the participant provided neutral responses as to whether or not the playbooks were useful. They somewhat agreed that playbooks improved quality of work, positively impacted productivity, and supported critical aspects of the job; however, they somewhat disagreed that playbooks allowed them to accomplish more work than would otherwise be possible. In particular, P3 felt that playbooks might be useful for more complex problems, but were not especially useful or time-saving for smaller-scale issues, like the ones in our scenarios.

P3 also reported needing to rely heavily on his security engineering background, as he found both playbooks too abstract to directly guide the development of new security controls. P3 slightly preferred the NIST playbook, citing previous familiarity with the framework (which he had seen but never used prior to the study). They reported spending more time with the IACD playbook to ensure an effective outcome, but attributed this primarily to lack of familiarity with the framework. P3 hypothesized that IACD's visual presentation would be easier for less experienced technicians to work with, but found the resulting playbook too generic for direct implementation. P3 reported making many notes to expand on each step and recommended adding complementary reference sheets providing detailed instructions

for each step.

6.3.2.3 Playbooks in use

Below are the results of the evaluation using incident response playbooks in an enterprise environment during a series of no-notice attacks. Reported are (1) how we conducted controlled incident response events through the use of a trusted agent to conduct simulated insider-threat attacks, (2) metrics for the efficacy of playbooks during incident response, (3) and general observations. Playbooks during the first two incident response events had mostly negative results: experienced security professionals did not feel the playbooks added much value to their response efforts, and junior analysts struggled with detecting incident-response events. After making modifications to the playbooks based on feedback from participants and our experts, as well as lessons learned from the first two incident response events, participants' perceived usability of playbook frameworks increased noticeably.

IR1 outcomes. During the first event, the trusted agent initiated a no-notice brute-force login attack against two user accounts. P4 was the first to detect the event, notifying his supervisor of a potential incident 10 minutes after attack execution. P3 independently detected the event two minutes later. The supervisor informed both participants that this was an exercise, and that they were to finish investigating the breach independently and without informing other technicians. Within one hour of detecting the threat, both participants successfully identified the point of origin for the attack, recommended removing the infected system from the network, identified

the person using the now-quarantined system, and notified the physical security team about the (notional) insider threat.

P2, however, did not detect the attack until 14 days later. P3 and P4 left all attack logs in place after their investigation concluded to provide P2 with more chances to detect the attack in the future (and allow us to assess P2's response decisions). According to the NDC supervisor, it is common for multiple technicians to check the same security logs and dashboard for alerts for redundancy. The brute-force alert appeared on P2's dashboard at least 19 different times during morning and evening checks, but P2 did not recognize it.

Once P2 realized an event had occurred, they made an initial report to a supervisor within 10 minutes. After that, it took P2 four hours to successfully identify the root cause of the attack (and the associated user) and submit an incident report to the physical security team. Altogether, 335 hours elapsed between the initiation of the attack and P2's report.

IR1 feedback. Participants P3 and P4 noted how the playbook contributed to their successful responses. Both said the brute-force playbook (IACD, written by P4) straightforwardly guided them to correct actions. However, both largely credited their past security experience rather than the playbook for the successful outcome. In particular, both said they relied on knowledge from past experience to make up for missing details, such as syntax for querying access logs to determine who was logged in at the system that initiated the attack.

P2 confirmed that he used the playbook, but nonetheless missed the alert as-

sociated with the brute-force attack all 19 times. They said, “the playbook did not have enough information for us to conduct a step-by-step walk-through. Because I was unfamiliar with the new [brand] dashboard, I didn’t know what the alerts would look like compared to normal data.” This suggests the 30-minute orientation session was insufficient for this novice defender to learn how and when to use the playbook. Because P2 missed the alerts, they never identified the initiating condition that requires a technician to use the incident response playbook. This aligns with participants’ comments in Section 6.2.3 that the triggering event is the most important step in a playbook design framework.

Further, P2 commented that since it was his “first time using [the playbook for an event], we needed to work out who to inform and when. Identifying critical information for each step and who needs to know it would have saved time.” P2’s supervisor rejected three reports during the 10-minute initial response window because they lacked sufficient detail to communicate what was going on.

P2 also noted that having two playbooks available delayed their response: faced with a stressful situation, P2 read through both playbooks to ensure they were using the correct one. There was no table of contents and no easily identifiable markings in the playbook headers (like bolded or colored text) to help an analyst quickly choose the correct playbook. “It would help to more clearly identify which playbook is for which event.” P2 commented that this problem could become worse with more playbooks for other kinds of incidents.

When analyzing post-utilization perceptions of playbooks, all three participants indicated an overall neutral sentiment (averaging 3.45 on 5-point Likert scale

questions adapted from SUS, $\sigma=0.69$) but all three strongly agreed that playbooks support critical aspects of their job.

IR2 outcomes and feedback. P2 and P3 participated in IR2, a credential-misuse event conducted in January 2020. (P4 was unavailable due to off-site training.)

P3 again successfully responded to the incident, performing nearly identically to their response in IR1 and resolving the situation in 65 minutes. P2 again failed to recognize the significance of alerts generated during the attack; the incident ended after 11 days with no recognition.

As with IR1, both participants noted that the NIST-framework credential-abuse playbook lacked sufficient detail, and P3 again relied heavily on past experience for their response. P2 provided two possible explanations for failing to detect the incident. Primarily, they said they had taken several weeks off from work for the holiday season, causing familiarity with the playbooks to atrophy. Second, P2 did not believe they would be evaluated with the playbook a second time. These comments align with findings from previous adult learning theories about the importance of continual, hands-on practice with new concepts [21, 118].

Resetting after failure. After IR1 and IR2, researchers worked with NDC to revise the designed playbooks, applying feedback we had received in the first case study and in this case study so far. In particular, researchers sought to address three interrelated concerns that had surfaced repeatedly: that playbooks contained insufficient detail for use during incident response, that too much experience was

required to use the playbooks properly (making things difficult for novices), and that identifying a playbook trigger was the most critical challenge.

First, all NDC participants collectively improved both playbooks by adding details appropriate for use by an entry-level technician, including click-by-click instructions for GUIs and specific text for command-line interfaces. As the playbooks expanded in detail, technicians documented lengthy processes by creating complementary guides alongside the playbooks. Technicians also made changes focused on recognizability: creating a table of contents for all playbooks, using bold-font titles on each playbook, and including summaries for what the playbooks are intended to help with.

Collectively, NDC participants walked through both scenarios in an ad-hoc tabletop exercise, annotated playbooks gaps, and later made updates accordingly. In one example, the tabletop exercise revealed that instructions for communication were not yet sufficiently detailed; after the exercise, technicians made cheat sheets documenting which information must be reported and to whom for each playbook step.

Next, researchers asked NDC to implement a more collaborative model in which technicians could work together while using playbooks. In particular, junior technicians were encouraged to ask questions and seek advice from senior leaders and technicians. After these changes, one month passed before initiating the final incident response event.

IR3 outcomes and feedback. The trusted insider initiated IR3 (credential abuse)

in March 2020. P2 and P13, both of whom had volunteered, were selected by NDC leadership as participants. P13 joined NDC two weeks prior to IR3 and completed all the on-board training related to playbooks that NDC had implemented. P3 and P4 were unavailable for IR3 due to other job obligations.

Both participants successfully detected (P2=3 min, P13=5 min) and responded to (P2=90 min, P13=104 min) the threat within our 140-minute threshold. P2 said that the more detailed steps added to the playbooks and the new mentorship program helped drastically with their understanding of how to respond to events and communicate more effectively with their supervisors. P13 said, “As a new employee, it helped me better understand our mission and how to do my job if a supervisor is not available.” By completing the on-boarding training using playbooks, P13 felt they more completely understood their role and responsibilities within NDC: “This is what I do, this is what is required of me.” This supports previous claims regarding the usefulness of playbooks for helping professionals learn new responsibilities and technologies [89].

After IR3, both participants strongly agreed that playbooks made their jobs easier, enhanced their effectiveness on the job, and allowed them to accomplish more work than would otherwise be possible. After IR1 and IR2, P2 had answered neutrally to these questions. Both participants also strongly disagreed that playbooks were confusing; P2 had answered neutrally after both IR1 and IR2.

While one cannot generalize from this experience, IR3 suggests that, when they include sufficient detail as well as additional practice and orientation, playbooks may be useful to help junior technicians with incident response. Future work is needed,

however, to investigate the extent to which different elements of the implemented improvements are useful.

6.3.3 Summary

Playbooks designed using both frameworks required significant modifications to be useful and usable, especially for more junior technicians. During IR1 and IR2, experienced technicians used the playbooks designed during the first case study — created within 45 minutes — successfully, but credited most of their success to prior experience rather than the playbooks. A junior technician was unable to respond within the expected time window in either case.

After updating the playbooks (and associated organizational processes) using lessons learned from both case studies, two junior technicians were able to use them to mitigate a credential misuse attack within 110 minutes.

These findings suggest that current playbook frameworks are not sufficient on their own, but may be useful as part of a larger process for developing and institutionalizing playbooks; further research is required for validation.

6.4 Playbooks in other domains

Playbooks, and guidelines for developing them, can be found in domains outside of digital security.

Business continuity plans (BCPs) help minimize financial losses, ensure the continuation of core functions, ensure resource availability, and train employees.

Many organizations are required by insurance or regulations to have BCPs. Numerous references provide reporting templates for communicating essential information, how-to guides for audits, and training scenarios for a vast array of situations that may cause damage to a business [27, 80, 94]. BCPs typically contain fine-grained detail to assist with implementation and auditing (similar to the playbooks used during IR3). BCP training varies, but typically involves intricate exercises [94].

U.S. government agencies maintain playbooks for natural disaster continuity and health emergency preparedness, among other crises [224]; libraries of pre-made disaster response playbooks are available for reference [67, 148, 227].

In the medical field, crisis resource management combines standard medicinal practices with non-technical skills to ensure exposure to best practices for likely emergency situations [41]. Studies found that simulated rehearsals with response action “playbooks” gave participants confidence that the lessons learned would transfer to real-life situations [175].

6.5 Discussion

Using two case studies, researchers provide the first structured evaluation of playbook framework usability within an enterprise environment. Overall, the findings suggest that playbook frameworks are moderately usable for technicians designing playbooks, but do have important areas for improvement. Playbooks generated using these frameworks may require significant modification to meet their goals of helping technicians implement the associated security controls and then respond to

security incidents. Perhaps the most significant drawback, observed in all phases of our evaluation, is that the frameworks do not elicit playbooks written in sufficient detail for real-world use. More experienced technicians were able to rely on their prior knowledge to fill in these gaps, but junior participants struggled to make use of the playbooks. Based on these results, we make several recommendations for playbook frameworks, playbooks themselves, and associated organizational processes.

Improvements to playbook frameworks. Technicians must understand the initiating condition for incident response and be able to detect it – everything that follows the initiating condition is irrelevant if defenders do not recognize the need for action. Engineers who implement detection mechanisms must generate meaningful alerts and should consider requiring a technician’s acknowledgement [53].

Playbooks must be usable during stressful situations. Minor changes such as using boldface titles, using a table of contents to organize multiple playbooks, and affixing summaries atop playbooks seemed to help technicians in our case study quickly select the appropriate playbook for a given situation.

Visualizations may support technicians in understanding the high-level approach and tasks required to respond to a given incident, but technicians of all experience levels indicated that highly-detailed instructions (based on best practices) are critical. Playbook designers should not assume playbook users are experts with various technology platforms or command-line interfaces. Instead, they should provide detailed instructions both for implementing required security controls and for responding to an incident.

The NIST playbook framework emphasizes communication throughout incident response, but IACD allows the designer to determine which communication is essential or optional. Both case studies suggest that playbooks should prompt technicians about what information to record as well as who to inform and when. Fill-in forms (such as those found in DHS bomb threat checklist [58]) could be useful for this purpose.

Playbook designers using IACD had difficulty grouping tasks together, in part because the instructions left the choice of groupings open-ended and provided little guidance for how to identify and label groups. Playbook frameworks might consider providing multiple-choice options for category selection, guides with more detailed prompts, or a large corpus of training examples annotated with explanations.

Playbook frameworks should prompt designers to plan for non-linear actions: accomplishing tasks in parallel and accounting for multiple scenarios that may occur during response actions. Parallel tasks allow for the execution of multiple automated response actions to expedite investigation. Offering best practices for a variety of likely encounters and adversarial actions could allow responders to maintain momentum during an investigation.

Playbooks should include the intent associated with every task. Helping users understand why a task is relevant may allow them to exercise initiative and improve overall response efforts [56]. Additionally, experts suggested that intent specification may help security engineers who are implementing automation solutions based on playbook design to better understand and meet requirements.

Finally, playbooks must carefully balance including all necessary information

without including too much information. Too much information could slow response time as technicians sift through details to determine appropriate next actions. One possible mitigation could be to include links and references to external resources such as best-practice repositories, allowing designers to convey important information without overly cluttering the playbook itself. Further research is needed to explore this tradeoff.

Improvements to organizational processes to support playbook adoption.

During the evaluation phase, all three experts recommended using tabletop exercises [225] to iteratively update each playbook until it is sufficiently detailed and tailored specifically for the local environment. These exercises can be used to validate that the playbook is complete and that all necessary policies and procedures are in place to support incident response. These exercises were perceived as helpful to the playbook revision process we observed.

Expectancy theory [18] suggests that if playbooks do not feel useful, it is unlikely they will be used. Improving playbooks themselves, as described above, will improve perceived usefulness, but organizational culture around playbooks may also play an important role. In the second case study, organizational improvements such as mentorship programs, peer partnering, and continual reinforcement of the playbook process were cited by our participants as helpful in improving their perception of playbook usefulness.

Finally, playbook designers must consider organizational concerns and processes. Understanding particular constraints, such as requiring approval to make

changes to a network or limiting hardware purchases to previously approved vendors, may shape an organization's incident response strategy and therefore its playbook design. Designing a playbook that meets best practices while conforming to local constraints may require significantly more effort and time than the averages shown in Section [6.2.3.2](#).

Chapter 7: Implementing Security: Complementing and Repairing Baseline Security

Proactive security measures require organizations to understand their own security posture, understand likely threats, and prioritize mitigation efforts. Chapters 4 and 5 explain that digital security compliance programs are comprised of technical controls and policies that establish a baseline of security practices in an organization. This baseline is mandatory for organizations to provide critical services or control access to sensitive data but also is laden with their own security issues. As a result, security professionals are left to assess the impact of compliance on their organization and to identify ways to extend security beyond compliance mandates to fill known security gaps.

This chapter reports on organizations' use of *complementary measures* — policies and technical controls enacted to mend known security gaps and exceed compliance requirements. To gain a better understanding of complementary measures, I surveyed 40 security professionals from across multiple essential-service sectors and representing several multi-million dollar organizations. Participants reported on which complementary measures their organizations use to address which security gaps, which complementary measures worked particularly well (or poorly), and how

their organizations prioritize and evaluate the complementary measures they adopt.

As expected, only 10 participants believed compliance programs are sufficient in themselves to establish baseline security, and 37 of 40 participants reported implementing complementary measures to mitigate risks unaddressed by compliance standards. Some of the most commonly reported complementary measures include multi-factor authentication, endpoint detection and response tools, periodic account-access reviews, physical access barriers, and threat-hunting processes.

Although the specifics of how and why organizations implement complementary measures vary, organizations often adopt complementary measures in response to security incidents, to reduce costs, when recommended by external experts, or requested by (sometimes non-technical) executives.

On the whole, participants found complementary measures to be beneficial, but far from perfect. Organizations know that gaps in compliance exist, and therefore solely relying on compliance to drive a defensive security posture exposes the organization and its users to risk of an attack. Therefore, organizations create complementary measures that go beyond compliance, however these efforts face organization inertia and risk. Participants reported numerous instances of poorly managed complementary processes, investments in unproven or incompatible “solutions,” information saturation, and difficulty keeping complementary measures up to date and relevant. The results can be used to improve compliance mandates that acknowledge their shortcomings and provide guidance on how to evolve a security posture against the threats of tomorrow.

7.1 Method

This section discusses survey design, participant recruitment, and the quantitative and qualitative analysis conducted on participant responses.

This study was reviewed and classified exempt by the UMD ethics-compliance office. Participants provided information about their professional experiences, perceptions, and background. Due to the sensitive nature of unmitigated security vulnerabilities, participants only disclosed information they were comfortable with sharing; additionally, the findings are generalized to protect organizations and systems.

7.1.1 Survey design

This study used a 21-question survey with a combination of open-ended and close-ended questions broken into four sections: introduction/screening, baseline understanding, assessment of complementary measures, and demographics (App. A.5). Research suggests that the quality of survey responses decreases over time, and excessively long surveys may result in a participant quitting the study [100]. To this end, I designed the surveys for experts to complete within 30 minutes of focused effort, in line with suggested best practices [73]. Actual completion time averaged 27.9 minutes ($\sigma = 0.024$). Participants were not compensated directly, but were invited to opt into a raffle for one of two \$50 gift cards.

First, participants answered screening questions (see Section 7.1.2) to ensure they were qualified to address our research questions.

Next, participants responded to two questions to better understand sentiment and baseline defensive practices: (1) If participants' organizations believed compliance is sufficient to protect their systems and data and (2) If participants' organizations employed proactive security controls to address threats not covered by compliance programs. Participants who indicated their current employer enacts defensive measures complementary to compliance controls were directed to the next section; otherwise they were directed to the demographics section. For participant-reported measures, our researchers independently verified that applicable compliance standards did not actually require its use; all reported measures were, in fact, complementary to compliance requirements.

The third section presented participants with a list of 18 proactive security controls, selected from a corpus of digital-security risk-mitigation literature [78, 207, 217] and previous research on applied security [60, 117, 191, 192]. Participants to selected all of the controls they employ that *complement* existing compliance controls at their organizations. Additionally, participant could describe and discuss other (unlisted) security paradigms they may employ.

If participants selected more than five complementary measures, the survey prompted them to select the five controls they were most interested in discussing. The survey back-end then randomized the order of the participants' five selections and asked six questions per control. Two questions were Likert-scale questions asking (1) how frequently the participant's organization assesses the control's effectiveness and (2) how well the control has worked out for their organization. Four of these six questions were open-ended and asked participants to describe in detail:

(1) why the control was implemented, (2) the aspects that worked (or did not work) well with implementation, (3) how participants ensured complementary measures were compatible with compliance standards, and (4) the key factors for prioritizing complementary measures.

Next, participants provided demographic information about their experiences and perspectives. These included specific work role, current business sector, years of experience, and information about their clientele.

Finally, the survey prompted participants' permission for further contact if response clarification was needed or for future studies.

Survey pilot. Prior to broadly distributing our survey, I asked two security professionals to complete the survey and provide feedback, specifically focused on question relevance, completeness, and clarity. I updated the survey based on pilot feedback and overall study flow; the final version of the survey is listed in Appendix [A.5](#).

7.1.2 Recruitment and Screening

Researchers leveraged personal contacts, email distribution lists, and social media outlets tailored towards multiple different business sectors to assist with response diversity. Specifically, researchers sought participants from the following sectors: government, healthcare, financial services, consumer services, information technology, and education. Researchers also employed snowball sampling, in which participants recommended other qualified professionals. Diversifying participants based on their current work role and business sector supports ecological validity and

ensures findings represent varying perspectives.

Researchers screened participants based on three factors: (1) they are actively employed by an organization that uses digital security compliance programs, (2) their current job involves compliance standards, and (3) their current work role.

The first two factors ensure participants are dealing with compliance currently. Additionally, researchers selected participants who serve as security managers, security analysts, security engineers, governance experts, or software developers to increase the likelihood participants provided responses from a technical perspective.

Researchers also screened participants to verify they were fluent in English, over 18 years old, and within the United States.

7.1.3 Data analysis

Researchers use both qualitative and quantitative analysis to identify themes and trends across participant responses.

Iterative open coding. Two researchers independently analyzed all participant open-ended responses using iterative open coding, creating a *codebook* to categorize responses based on labels [199]. For each response, coders may identify one or more applicable category labels. These categories are then aggregated into broader themes [218].

If a survey response was unclear, coders would request clarification or additional information from the participant via email (if the participant had consented to additional contact within their survey response); otherwise, researchers discarded

the response in question. Overall, only two responses were discarded.

To establish a baseline codebook, the two researchers jointly coded a random 10% of the data set ($n=4$). This established a working set of label definitions.

Next, each researcher independently coded a new subset of the data ($n=5$) and calculated the resulting Krippendorff's Alpha ($\alpha = 0.8594$) across the entire codebook. Krippendorff's Alpha measures inter-rater reliability — a measure of consistency among independent coders — while accounting for chance agreements [88]. An α value above 0.8 indicates high reliability [120, 126].

All disagreements during this iteration were associated with participants' use of technical jargon that could have multiple interpretations. All disagreements were fully resolved, the codebook was updated, and the researchers again independently coded a new subset of the data ($n=5$), with an $\alpha = 0.8229$. With two consecutive independent IRR scores above 0.8, the two researchers split the remaining 26 responses and independently coded them using a shared, collaborative codebook.

The two researchers iteratively updated the codebook as needed; when revisions were made, the researchers re-coded previously analyzed answers accordingly. Researchers repeated this process until we resolved all disagreements and the codebook was stable. Both coders attained thematic saturation in each of the five codebook subsets prior to exhausting the list of participant responses. The final codebook is given in Appendix B.5.5.

Statistical analysis. Researchers asked two Likert-scale questions about each complementary control participants described: their satisfaction with the control,

and how frequently that control is assessed.

To compare satisfaction across groups of controls, researchers used an ordinal logistic, mixed-model (random effect) regression [42]. This approach is appropriate for ordinal, non-continuous Likert data, while accounting for multiple answers from individual participants. Researchers added the adaptive Gauss-Hermite quadrature approximation with ten quadrature points to the model for better accuracy and fitting [174]. Full details of this regression are given in Appendix B.2.

To examine whether satisfaction is correlated with frequency of assessment, researchers used the non-parametric Kendall rank correlation coefficient, appropriate for ordinal data [1].

In both cases, we use $\alpha = 0.05$.

7.1.4 Limitations

All qualitative research should be interpreted in the context of its limitations.

For each finding, researchers provide counts for the number of participants who expressed that theme (and where relevant, the number of applicable security controls) to provide context. However, it is possible that participants may have omitted mentioning a specific concept when responding to open-ended questions rather than explicitly disagreeing with the concept. Therefore, statistical hypothesis tests are not used for these questions, nor is prevalence implied.

Researchers' recruitment messages and consent waiver explained the purpose of the study, which may lead to a self-selection bias such that personnel most in-

terested in the study were more likely to anonymously participate. However, this may also suggest that participants were prepared to think more critically about how compliance affects their security decisions.

All participants self-reportedly work directly with compliance standards and their experiences with compliance may reduce the possibility of demand characteristics — an experimental artifact in which participants unconsciously change their behavior to perform well within a study [167]. By allowing participants to complete anonymous online surveys, participants may be more likely to provide open-ended, candid feedback without fear of attribution or negative impacts from their employers [63].

In instances where participants indicated that their organization employs five or more complementary measures, participant selected five controls they were most interested in discussing. This may have introduce some bias into our results. Researchers felt this was acceptable in order to ensure participants were highly knowledgeable on the control and/or most willing to provide detailed responses about.

7.2 Results

Below are the study results on the use of complementary measures to address the shortcomings of digital security compliance programs within organizations that provide essential services. Below are participant demographics, the ways in which participants reported that compliance programs left their organizations exposed to risk, how and why organizations complement compliance controls, and the various

issues that arise when implementing complementary measures.

This section annotates prevalence by describing it with n the number of participants that reported a particular finding and with c the number of controls to which the particular finding was reported to apply.

7.2.1 Participants

Researchers recruited 100 participants for this study. In total, 41 responses were discarded due to a lack of participant qualification as well as 19 partial responses. Among the remaining 40 participants whose responses we analyze, researchers achieved data saturation by the fifteenth participant. These response rates, rejection rates, and population size were consistent with previously published studies with similar methods, participant types, and goals [14,29,38,210]. Table 7.1 describes the overall sample, and Table B.6 in Appendix B.5.4 details information about each participant.

The study participants included security managers (e.g., CIOs, CISOs, and SOC directors), specialists in compliance and governance, developers of security software, security engineers, and security analysts. Ten participants served as senior security officials for multi-million dollar organizations with client bases of more than 100,000 customers. These organizations represented six business sectors: consumer services, education, financial services, government, healthcare, and information technology. Twenty-six participants had more than 10 years of experience. Overall, the participants averaged 15.33 years of experience ($\sigma = 7.13$) working in information

Metric	Count	Metric	Count
Sector		Org Size	
Consumer services	3	0-50	5
Education	7	51-150	3
Financial services	1	151-500	9
Government	14	501-1000	3
Healthcare	4	1000+	20
Information tech	11	Clientele	
Job Role		1-500	6
Compliance	6	501-5000	10
Management	22	5001-10k	2
Security Analyst	6	10k-100k	9
Developer	4	100k+	13
Security Engineer	2	Exp (years)	
Education		2-5	5
Graduate degree	27	6-10	9
B.S.	11	11-15	8
Associates	1	16-20	7
PNTA	1	>20	11

Table 7.1: Participant demographics (n=40). First column highlights represented business sectors, current work roles, and educational background. Second column describes the number of employees at participants’ organizations, the size of participants’ clientele, and experience levels.

technology alongside compliance standards; median experience was 15 years.

Specific experiences may vary depending on the particular compliance standards in effect. The top 30% of compliance standards most frequently used by participants are: National Institute of Standards and Technology (NIST) Cybersecurity Framework (n=33), Health Insurance Portability and Accountability Act (HIPAA) (n=17), Payment Card Industry Data Security Standard (PCI DSS) (n=14), Federal Information Security Management Act (FISMA) (n=12), and at least one document from the International Organization for Standardization (ISO) (n=12). Figure 7.1 shows the distribution of compliance standards in greater detail. The full list is reported in App. B.5.3

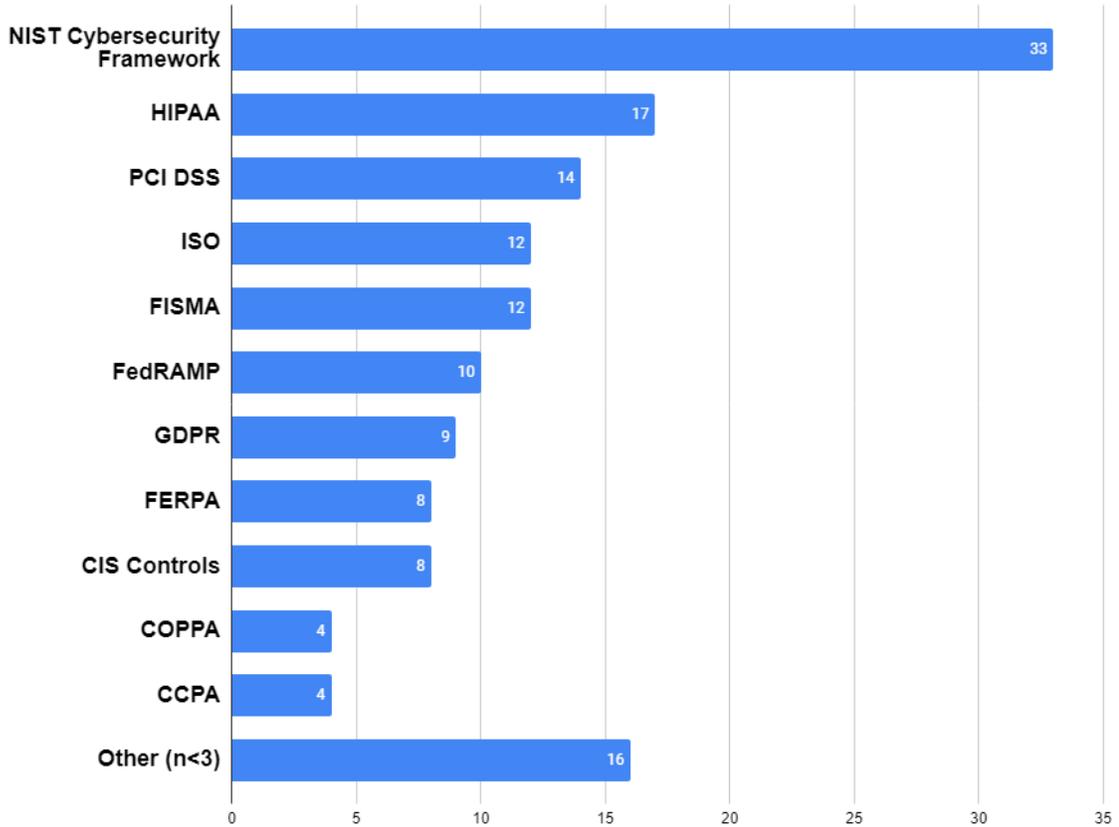


Figure 7.1: Distribution of standards used by participants. We aggregate standards used by three or fewer participants under “Other.” A complete list of standards and acronyms is given in Appendix B.5.3.

7.2.2 Compliance is insufficient

Many previous works detail complications with organizations implementing compliance programs. The following results provide further evidence from multiple business sectors that compliance programs are often insufficient for establishing baseline levels of security against common threats.

Overall, only 10 of 40 participants reported that their organizations believed compliance standards were sufficient for their security needs. Seven participants were unsure about the protection provided, and 23 indicated that compliance in-

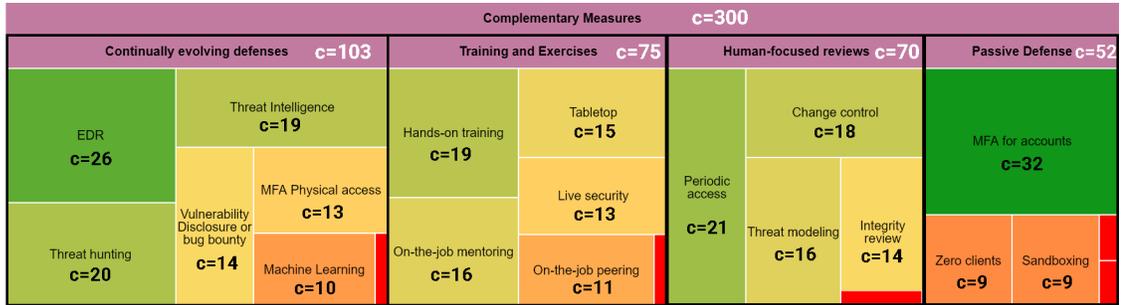


Figure 7.2: Distribution of complementary measures used by participants across four categories of controls. Block size and color indicate prevalence, with MFA reported as the most-used complementary measure (c=32) and five different controls mentioned only once each.

sufficiently protected their organizations and systems. Participant P17 stated that compliance failed to protect their organization from “nearly all threats. Compliance is so high-level and abstract it is nothing more than a ‘CYA’ [cover your ass] effort to make leaders invest in security.” Similar sentiment was shared by other participants, with 21 participants indicating that compliance was in some ways disconnected from addressing realistic threats faced by their respective organizations. Of note, this negative sentiment was shared by a majority of participants across all business sectors except for finance. Participant P11 — the only participant from the finance sector — offered their view of why their organization believed compliance standards were sufficient:

“Compliance standards are sufficient because there are SO many. The financial industry is literally choked with compliance standards. The real issue is whether the financial companies can implement those standards with enough flexibility to keep up with the changing threats, and that will depend upon the organization.” (P11)

This could be interpreted to mean that the plethora of (sometimes overlapping) compliance programs in the finance sector are sufficient to protect the organization, or perhaps sufficient in the sense that the organization was saturated and did not want to add more standards or controls.

In total, 37 participants indicated that they employed supplemental security controls (not required by compliance) to mitigate unaddressed threats. Here, seven of the 10 participants who previously indicated their organizations believed compliance standards sufficiently covered their security requirements explained that their organization faced specific threats and that compliance standards were too abstract to account for these threats. In these cases, compliance programs may have provided sufficient coverage for the industry at large but niche requirements remained. Participant P28 summarized the sentiment of these 37 participants, stating that:

“[Compliance is] a baseline to ensure you’re thinking about controls in many domains at a minimal/moderate [level]. Very often even the baseline controls are not even implemented well to begin with. [Compliance is] a starting point not a destination.”

Researchers then asked participants what specific threats were unaddressed by compliance programs. Researchers categorized these responses, and the largest category was emerging threats (n=10). Participant P19 stated that “published standards do not have sufficient flexibility and adaptability to changing threat types and methodologies. They serve only to resolve known or historic issues.” These attitudes align well with findings in prior work. Compliance standards vary in how

often they are updated, but nearly all fail to provide feedback opportunities after major version releases [192], and malicious exploit development surpasses the ability of compliance authors to modernize standards [2].

Twelve participants indicated that compliance fails against sophisticated attacks; three participants from the government sector indicated that nation-state actors are not deterred by compliance. P37 reported that compliance programs “only protect against 80% of threats (i.e. the low hanging fruit),” suggesting organizations are exposed to moderate and sophisticated attacks. Compliance was particularly concerning to P04: “we are a high profile target via name and reputation,” and because of that compliance leaves them “vulnerable to attacks.” An analysis of nation-state attack methods again highlights the gap between the speed and complexity of their attacks and the efficacy of compliance programs [54].

Seven participants stated that compliance does not adequately protect organizations from insider threats (n=7). This aligns with prior work suggesting insider threats possess privileged insight that allows them to bypass superficially-implemented defenses required by compliance [49, 102, 106, 208].

Other threats not covered by compliance but mentioned less frequently were: relying on self-reporting for security issues (n=1), denial of service attacks (n=1), phishing attempts (n=1), and untrained compliance auditors (n=1) who require modifications to security that have “no traceability to mission/business requirements.”

To proactively defend their organizations from these unaddressed threats, the 40 participants reported that they collectively employ 300 complementary mea-

asures to augment compliance. (As detailed in Section 7.1.1, participants were asked to select all applicable complementary measures from a predetermined list of 18, and offered space to report up to five additional measures under ‘other.’) After deduplicating the ‘other’ responses, researchers obtained a final list of 23 unique complementary measures that we bin within four different categories: (1) training and exercises, (2) human-focused reviews, (3) passive defense, and (4) continually evolving defenses. *Training and exercises* involve employees gaining exposure to defensive techniques interactively through hands-on training (n=19), formal mentorship programs (n=16), and tabletop “talk-through” exercises (n=15). *Human-focused reviews* are triggered by events and require human-in-the-loop interactions. Examples include change control boards that review and approve changes to digital systems (n=18), periodic account access reviews (n=21), or proactively assessing risks and developing mitigation strategies through threat modeling (n=16). *Passive defenses* involve technologies that infrequently require human interaction; examples include multi-factor authentication (MFA) for account protection (n=32) and zero-client hosts that provide users with new, pristine workstations for every use (n=9). Lastly, *continually evolving defenses* require extensive human-in-the-loop involvement to actively reduce threat exposure. The most used continually evolving defenses include endpoint detection and response tools (EDR) (n=26) which focus on detecting and investigating suspicious activities on endpoint systems such as workstations; implementing physical access controls (n=13) due to both physical and digital organizational changes (e.g., office swaps, new server rooms, or influx in hiring); threat hunting (n=20), where defenders attempt to identify and defeat

known or unknown threats that have already bypassed existing security; and threat intelligence (n=19), information feeds that inform defenders about emerging threats and recent events.

The full list of reported controls is included in Appendix B.5.2 and illustrated in Figure 7.2.

7.2.3 Going beyond compliance

Given that participants report compliance programs insufficiently address threats, researchers next explored how and why organizations choose to complement compliance controls. More than any other reported reason, compliant organizations implement complementary measures after they encounter a security incident. Other key factors include reducing overall costs and gaining better insight into network activity. Overall, participants generally have a positive outlook on complementary measures and the benefits they provide their respective organizations.

Security incidents lead organizations to adopt new controls. While researchers did not specifically ask if participants' organizations were the victims of a security breach, many offered that past incidents were a driving factor for implementing complementary measures (n=21, c=40). Security gaps not addressed by compliance programs compelled organizations to take action, reinforcing previous research that organizations make decisions based on risk exposure [99]. Participant P36 offered insight into their incident:

“We had a public data breach... a misconfigured database I think? There

was immediate pressure to prove to higher [management] that we were doing something to make sure it didn't happen again in the future.”

(P36)

Security incidents at already-compliant organizations inherently demonstrate that baseline compliance provides insufficient protection. To help mend security gaps exposed by incidents, 16 participants implemented continually evolving defenses (MFA, c=10), 14 implemented passive defenses (EDR, c=4), five implemented human-focused reviews (account access review, c=4), and five implemented training and exercises to help mitigate future incidents (tabletop, c=3). P27, a manager in healthcare, explained their reasoning for moving to EDR after a breach on top of compliance-mandated anti-virus:

“AV is just flatly insufficient. Attackers often use ”living off the land” tools, [EDR] helps to detect and prevent normal tools used in bad ways.”

(P27)

Participant P29 reported that their organization “does not embrace complementary measures, which has led to several incidents,” resulting in adoption of complementary defenses after the incidents. Participants (n=2) touched on the reactive inclinations of their respective organizations, with P23 stating that their organization waits until “incidents or threats appear, [then] prioritization changes.” It is important to note that even though organizations implemented complementary measures after an incident, not all of the new controls were directly related to the previous incident. By implementing new complementary measures after incidents

— whether or not related to the original problem — security teams signaled to their organization that they were dedicating resources (e.g., money and personnel) to improve overall security (n=3).

In addition to actual incidents, red teams — digital security professionals who act as an adversary to assess networks and systems — have a similar impact on implementing complementary measures, given that red teams are essentially controlled incidents. Three participants reported that they are more likely to initiate complementary improvements to compliance programs after a penetration test. This aligns with prior work suggesting that formal vulnerability reports can have a large impact [8, 223].

Organizations seek controls that reduce costs. Twenty participants indicated that budgetary constraints were key factors in deciding to implement controls not required by compliance. Five of these 20 said that if they were to complement compliance, the new complementary measures would need to reduce task completion times and overall costs. Participant P15 seeks “potential for asymmetric gains – [controls that let] a human do the same work 10x faster, or achieving quality/thoroughness that would be unachievable by any number of humans.” P40 looks for automation and “time-savings by reducing staff labor hours.” Business re-engineering researchers highlight these concepts as best practices, choosing to optimize the total effectiveness of employees rather than downsizing [82]. Similarly, P14 stated that some solutions “may be limiting if [they] are too time or labor intensive,” and their organization will avoid hiring new personnel to extend security beyond compliance.

When advocating for solutions that augment compliance, P09 and P10 had to appeal to senior management in terms of return on investment and getting the most “bang-for-buck.” However, Participant P24 indicates this is not always the case:

“Larger-budgeted enterprises can initiate security decisions at various levels based on what is needed,” while “organizations with low funding need to invest time into basic measures such as routine reviews of patch management and privileged account access or other inexpensive proactive measures like table-top reviews of incident response scenarios and in-person user training (that is actually engaging and informative).”

(P24)

The notion that larger-budgeted organizations permit lower echelons of decision makers to test various complementary controls is corroborated by P33, who said technicians at their organization are permitted to “[perform] pilots to determine if solutions were right for the need.”

Machine learning (ML) (n=3), on-the-job mentorship (n=4), and hands-on training (n=3) were other controls not required by compliance that participants selected to address skill shortages and overcome hiring limitations. P37, when discussing machine learning, stated that “humans do not scale and are in short supply, and security data is growing exponentially.” P26 similarly reported that they use ML because there is “not enough staff to keep up with human analysis” required to monitor compliance-mandated security platforms. P08 uses on-the-job mentoring at their organization because many of their employees are entry-level and have little-

to-no compliance experience; they said it is in their organization's best interest to "mentor our young employees to ensure they will be vigilant in the requirements of compliance and overall site security."

Some participants (n=2) cautioned about letting budgetary constraints drive security decisions when adopting complementary measures. P16 lamented their organization's decision to adopt EDR technologies: "we bought trash solutions from the lowest bidder."

There are unspoken benefits to having an incident. Budgeting constraints not only affect the adoption of complementary measures, but in some cases create perverse incentives. Participant P31 stated that the occurrence of incidents actually help security teams advocate for a higher budget prioritization. Participants P02 and P27 similarly discuss an often unspoken trade-off between security and budgeting. "If you have perfect security, you obviously don't need your whole budget so let's give it to someone else that needs things more," stated P02. For P27, "security breaches are a strong, public-facing signal that something is wrong and resources need to be applied to fix it. Embarrassment will continue until it is fixed." These comments fit with prior observations that security practitioners constantly compete for a slice of their organization's overall budget and must consistently demonstrate a return on security investments [12,187]. This reality may motivate security teams to roll out complementary measures over time, continually demonstrating to budget-controlling officials a need for growth beyond baseline compliance security.

Compliance measures do not provide requisite network insight. Participants often needed complementary measures to assist with decision-making because the insights provided by compliance controls were insufficient; this accords with findings from Kokulu et al. [117]. For example, participants reported that standard compliance controls lacked visibility into network traffic flows and user activities (c=35). Participant P36 explained one point of frustration with the NIST Cybersecurity Framework, General Data Protection Regulation, California Consumer Privacy Act, and other financial standards:

“It’s impossible to defend a network where I can’t tell you how many workstations are attached. How many belong to us? We can’t connect what users are visiting what sites, so how can I tell who downloaded malware?” (P36)

Complementary measures are intended to provide improved understanding, allowing managers and technicians to make better defensive decisions. In fact, participants (c=36) indicated they employed complementary measures to enhance the effectiveness of other digital defenses, some of which were mandated by compliance programs. Participant P28 uses EDR to support compliance-mandated anti-virus and post-incident reporting:

“Simple signature based AV is dead. EDR tooling gives a much richer vision of process execution that is valuable for both detection and forensics.” (P28)

Participant P32 decided to use MFA to complement their password policies and

provide “an additional level of security that assists in reducing the occurrence of gaining access to critical systems.”

Organizations rely on external recommendations. Sixteen participants indicated they rely heavily on the advice of experts from outside their organization or marketing to make security decisions not required by compliance (n=16). Four participants stated that the reputation of external experts plays a role in whether they adopt the recommendation or not (n=4). Participant P37 spends “a good bit of time finding vendors with truly useful technologies and not just well-marketed snake oil” when following up on external recommendations.

Executive-level decisions made in isolation are seen as harmful. Several participants reported that executives within organizations, some without technology backgrounds, make decisions to complement compliance without input from their technicians (n=5). This led to frustrations within the organization, with P22 feeling many decisions about complementary measures were based on “political pressure” to partner with a particular vendor, or “based entirely on whim of [the] CIO” without an operational need or threat model to justify the decisions. Similarly, P23 said complementary measures such as threat intelligence are “often seen as a ‘check the block’ [box] function for executives to claim they are doing things” to defend the network from threats not covered by compliance — often without a clear understanding of the expected outcomes.

P06 warned about the misalignment of resources based on these types of deci-

sions, recalling a time when their organization chose to buy a new security platform rather than exploring why their existing tools failed: “Interesting technology is great but if it doesn’t address a critical need, we end up working on less important needs.” P30 stated, “once it is determined what product is wanted [by managers], security is brought in to assess [the solution], which is slightly backwards.” P22 made a similar point discussing their company’s implementation of EDR: “It was a dumpster fire” because the purchased solution only worked on a fraction of the company’s systems.

Bottom-up recommendations to managers shaped implementation strategies. Eight participants who were managers said they implemented complementary measures based on technician-identified needs to address security gaps remaining despite compliance standards. In total, managers implemented 26 controls based on bottom-up requirements from security employees. Managers said they adopted these controls to reduce the time required to accomplish tasks, improve overall performance, and enhance shared ownership of the security situation.

Organizations generally have a positive outlook on complementary measures. Thirty participants indicated that their complementary measures had valuable outcomes for their organizations. More specifically, seven participants reported that their security investments made their organization’s overall attack surfaces smaller. For example, Participant P03 stated that their use of vulnerability disclosure programs, threat hunting, and live security exercises identified “numerous gaps that scoped tests required annually did not find.” Three participants stated that

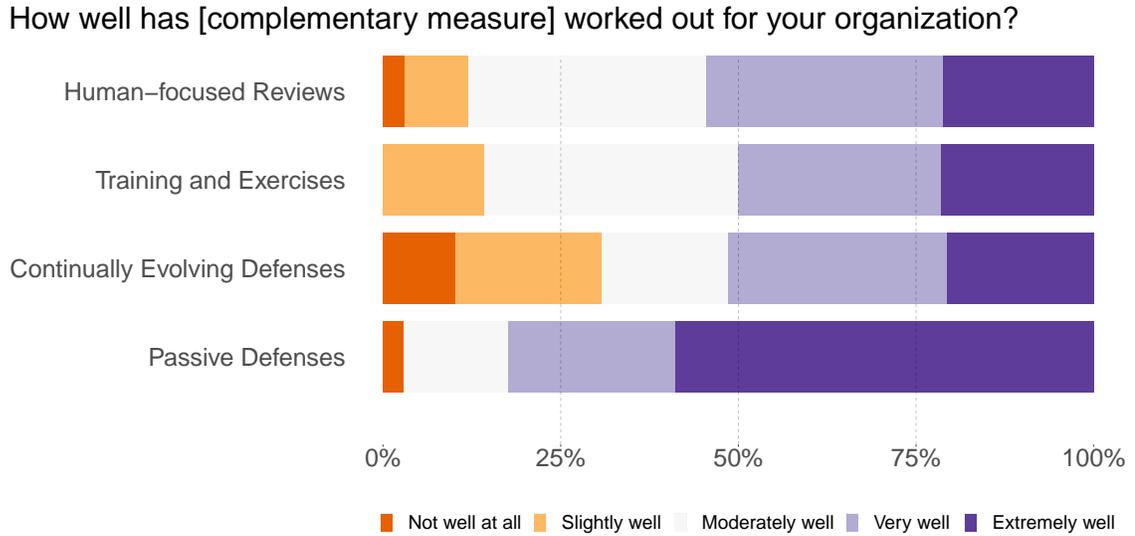


Figure 7.3: Visualization of participants’ sentiment towards complementary measures, per labeled category.

MFA counterbalances weak password policies, like those from the Internal Revenue Service [192]. Participant P13 enjoys not needing “to remember complex passwords, just need pin” for an MFA smart card, while P15 stated MFA removes attack vectors associated with passwords like pass-the-hash [163] or hash cracking.

Analysis revealed, using the quantitative analysis methods discussed in Section 7.1.3, that participants preferred complementary measures that did not require much effort to maintain. As depicted in Figure 7.3, the passive defense category — complementary measures that require minimal human-in-the-loop interaction — had the highest overall sentiment scores (averaging 4.35 out of 5, $\sigma = 0.95$) and served as the baseline for our ordinal logistic regression. Participants, as shown in Table 7.2, were significantly more satisfied with passive defenses than with any other category of complementary controls. In fact, the point estimates for the odds ratios indicate that participants were only 10-20% as likely to express higher satisfaction

Control Category	Odds Ratio	Conf. Int.	p-value
Passive Defenses (c=52)	–	–	–
Training and Exer. (c=75)	0.16	[0.06, 0.47]	0.0008*
Human-focused Rev. (c=70)	0.19	[0.07, 0.52]	0.0012*
Cont. Evolving Def. (c=103)	0.11	[0.04, 0.29]	<.0001*

*Statistically significant

Table 7.2: Summary of regression modeling participant satisfaction levels as a function of control category. Results demonstrate that passive defenses were preferred over other complementary measure categories

Proactive Control Group	τ	Correlation	p-value
Passive Defense	0.08	Weak (+)	0.6057
Training and Exer.	0.21	Weak (+)	0.2002
Human-focused Rev.	0.23	Weak (+)	0.1148
Cont. Evolving Def.	0.26	Weak (+)	0.0137*

*Statistically significant

Table 7.3: Ordinal association between participants’ satisfaction level with each control and their reported assessment frequency.

in the other categories as they were for passive defenses.

Additionally, participants are more satisfied with complementary measures when they are assessed frequently enough. The results from Kendall’s τ comparisons (Table 7.3) show a significant but weak positive correlation — indicating that satisfaction is higher when assessments are more frequent — for continually evolving defenses. Similar correlations are observed for the other three categories, but these trends do not reach statistical significance.

7.2.4 Additional measures are not a panacea

Despite participants generally having a favorable outlook, participants also warn that complementary measures are not one-size-fits-all. This section highlights

a range of challenges and complications associated with adoption of complementary measures.

Positive benefits come at a cost. Participants' efforts to fix security gaps not covered by compliance programs came at a cost to the organization — consuming time, money, or additional human capital (n=37).

Organizations typically weigh the costs of managing a compliance program [23], but also must consider the workforce costs for supporting complementary measures. While training and exercises invest in the technical competency of the workforce, they also require organizations to allocate time for employees to be away from their primary job. Participant P26 reported that live security exercises allowed for technicians to “discover what worked, and what didn't” during exercises — increasing the likelihood that mistakes are made during training and not real-life. P36 stated: “we usually see instant benefits after training... we typically cycle people through training in small groups so the overall security team still functions.” In addition, P02 warned about significant planning obligations leading up to exercises: “two planners from our SOC participated in [about] 100 hours of planning for 12 hours of training.” With complementary behaviors like on-the-job mentorship programs, P35 stated they provide the “ability to raise talent level. Strengthen internal rapport, structure, work product[s]. Shorten responses and knowledge transfer [during] emergencies.” At the same time that mentorship programs raise the talent level of organizations, P36 cautioned that organizations need to be prepared to dedicate time away from their job and focus deliberately on mentorship: “if you don't set

aside time for it, it isn't happening. But if you set aside time for it, plan for what isn't being done during that time. ”

Participants noted that human-focused reviews require organizations to trade speed for enhanced security. Participant P12 stated change control in their organization “slowed down change but increased reliability”; P20 made similar comments. According to P35, human-focused reviews in general should have mechanisms for “temporarily breaking beaucracy” under urgent circumstances and should optimize everyday timelines when possible — factors often not accounted for in compliance programs [192].

Participants also reported usability concerns with passive defenses (n=11) that diverted a significant amount of time away from other security tasks. MFA, the most commonly employed complementary measure, also had the most usability concerns. Six participants (n=6), representing each surveyed business sector, stated the security benefits of MFA came with usability challenges including lost smart cards, the migration of soft tokens to new phones, and lost hardware tokens (thus, corroborating complications discussed by Neware et. al [155]). Similarly, P17 and P19 highlight security-usability tradeoffs in the use of zero-client systems, which have no host operating system or storage and instead serve a clean virtual desktop that is erased after each use. Zero-client systems can create an “easy to establish ‘gold’ standard [that can be] updated as needed,” (P17), but the lack of “persistence or personalization of the operating environment” (P19) can inhibit required work.

Participant P25 indicated that some organizations implement controls without thinking about the “next step” of usability. For example, with sandboxing, their

organization suffers from usability “challenges [in] getting samples from the live environment to the [forensics] sandbox in a safe manner.” Mapping out end-to-end use cases may provide a benefit in adopting new technologies.

P21 noted tradeoffs in implementing end-to-end encryption, a passive defense measure: the security and privacy benefits “must be balanced with needs for logging, troubleshooting and forensics,” including creating challenges during incident investigation.

Complementary measures should not conflict with compliance. While complementary measures are intended to augment compliance controls, they are not always fully compatible with existing compliance standards. Participants reported that for 46 implemented complementary measures ($c=30$ from the government sector), there was no check for compatibility with compliance. P21’s comment about end-to-end encryption, for example, noted that this complementary measure may inhibit collection of logging data that is required under some compliance regimes.

Sometimes incompatibilities between compliance controls and complementary measures are more nuanced: P19 claimed that tabletop exercises do not have anything to do with compliance. However, considering compliance while executing tabletop exercises may help ensure participants practice compliant actions such as protecting sensitive information from improper disclosure [111, 204].

However, some participants did report that their organizations carefully consider compatibility when implementing complementary measures. Participant P28 described their organization’s methodical selection process for ensuring compatibil-

ity with existing programs. Five participants reported that they looked specifically for MFA solutions marketed as “compliance-ready” before buying, a trend previously identified by Julisch et al. [113]. When incompatible issues arise, P40 stated they “encourage self-report[ing]” when compliance may have been violated, which runs contrary to many zero-tolerance policies that enact financial sanctions for all infractions [158].

Poorly-managed measures provide reduced benefit. Eleven participants reported instances of complementary measures that provided reduced or even no benefit when poorly managed within their organization (n=11). Participant P27 indicated that their organization paid “six figures” for intrusion detection systems that remained in storage and were never set up (a pilot participant also reported a similar situation).

P25 said of their vulnerability disclosure program: “developing the program was great... informing everyone of its existence has been a struggle.” As a result, few vulnerabilities have been discovered or remediated. In contrast, P24 said marketing for their disclosure program yielded high participation with “over 100 vulnerabilities identified” despite their organization being compliant. However, uncovering vulnerabilities using a disclosure program may still not be sufficient if there is no plan in place to manage them: P01’s organization had issues reported, but they “go into a backlog where they don’t get remediated.” This finding accords with other examples of mismanagement of vulnerability disclosure programs [8].

Three participants reported financial losses when implementing threat hunt-

ing because their organization hired unqualified employees and did not adequately understand their own networks (n=3). Participant P17 stated that it “turns out finding unknown threats from systems that aren’t baselined is hard.” Since their organization’s compliance programs did not require up-to-date documentation (such as network maps), the organization paid hunters to sift through a network that its own administrators did not understand. P16 similarly said, “we go where we fear the threats are, rather than where they actually are,” that their organization often “ignores their [threat hunters’] findings,” and “fails to train, equip, or employ [threat hunters] properly.” P12 was “not convinced [their organization] brought in the right hunters.” As with vulnerability disclosure programs, complementary measures may fail if the organization is not prepared to use them effectively.

Three participants emphasized the importance of managing routine human-focused reviews, which is a known weakness in compliance standards themselves [192]. P12 complained that their organization does not perform account-access reviews as frequently as their internal policy requires; P24 offered that access reviews “don’t work well unless you commit to a routine schedule, and make time to conduct the review.” Additionally, P24 warned that change-control review boards “can become extremely bureaucratic and provide the opportunity for non-decision makers to become gatekeepers that slow down the process.” Of note, P24 indicated that missed change-control board response deadlines significantly delayed approvals for a new security platform. These comments suggest that complementary measure sometimes reify problems with baseline standards rather than alleviating them.

Keeping complementary measures relevant is difficult. A key weakness of compliance is staying up to date with current technology and best practices [192], but complementary measures often struggle with the same challenge. In total, 13 participants warned about the difficulties of keeping complementary measures relevant.

Participants argued that organizations should ensure training and exercises are congruent with the current threat landscape to maximize effectiveness (n=3). P16, expressing their frustrations with live security training, stated “We let morons design them. They are not grounded in reality and are at least five years behind [current] threats.”

Information does not equate to actionable intelligence. Ten participants reported that their organizations struggle to act on the information gained from implemented complementary measures (n=10), with four complaints specifically focused on threat intelligence (n=4).

Participant P16 lamented “information overload,” indicating that their organization’s implementation of threat intelligence “is neither timely, nor actionable. It is designed to give the illusion of insight, without forcing meaningful change.” P17’s organization likewise “struggles to quickly integrate paid vendor intel into our analysis systems,” and similarly, P13 has “yet to see any complementary measures taken based on threat intel.” P19 mentions that information overload of this kind can delay responses: “current models for developing and evaluating threat intelligence have been successful in timely development of information but have not provided

sufficient time to mitigate across the domain.” This sentiment corroborates prior findings that return-on-investment for threat intelligence varies [124].

Outside of threat intelligence, P24 warns that information is not sufficient when it’s not used properly: “Logs from agents may not be collected properly or reviewed by personnel with the proper training.”

7.3 Discussion

This paper examines how organizations overcome shortfalls with digital-security compliance programs. Security professionals rely on a wide range of complementary measures to address the threats their organization face, as discussed in Section 7.2.2 and depicted in App. B.5.2. While many security professionals described effective methods of complementing compliance, they also reported numerous inefficiencies and challenges that can occur when implementing complementary measures.

Based on these results, below are some recommendations for improving both compliance standards themselves and the ways that organizations supplement them.

Integrating complementary measures into compliance. This work echoes others in finding that compliance standards are insufficient on their own, leading many organizations to introduce complementary measures.

Section 7.2.3 shows that, in many cases, complementary measures can effectively reduce organizations’ attack surface. Complementary measures that gain significant adoption and acceptance are promising candidates for incorporation back into revised compliance standards as requirements. Standards authors should for-

malize mechanisms for audited organizations to provide feedback about the complementary measures they are using, why they are using them, and how well they are working. This would enable standards authors to observe trends at scale and identify generalizable benefits for participating organizations [16].

Documenting complementary measures' use cases also provides an opportunity to assist with compliance compatibility. Specifically, standards can be written to directly recognize that complementary measures are often desired or even necessary, akin to guidelines provided by the U.S. Food and Drug Administration [70]. Here, standards authors can create provisions that require organizations to document and carefully manage the complementary measures they implement, without spelling out what those measures might be. In particular, standards could require that organizations (and therefore auditors) ensure: (1) a holistic management program is in place, (2) complementary measures are being routinely monitored and/or adjusted, (3) employees are provided with requisite training to understand and implement complementary measures, and (4) incidents related to complementary measures are remediated.

Recommendations must come from reputable sources. Section 7.2.3 shows that organizations typically rely on advice from security tool vendors, external experts such as red teamers, or devise their own strategies for implementing complementary measures. But as this study shows, these sources may not always lead to security benefits: organizations are left to sift through vendors' "snake oil" solutions and some implementation strategies fail to address actual security problems.

Aggregating data that supports the efficacy of complementary measures, akin to the Building Security In Maturity Model [133], can help organizations make better implementation decisions and understand what worked well (or poorly) before investing. Anonymizing and making “success stories” publicly available may also help overcome common security secrecy [230].

Keeping pace with evolving threats and technologies. Compliance programs struggle with agility and responsiveness to evolving threats and technologies — often driving organizations to implement complementary measures. Our findings suggest, however, that some complementary measures suffer similarly from insufficient timeliness. There are opportunities to improve agility for both compliance and complementary measures.

Efforts to make compliance standards more responsive could reduce the need for complementary measures but, as others have noted, rapid compliance changes may have negative organizational impacts [135,192]. As seen in Section 7.2.4, there are examples of complementary issues failing to evolve when organizations invest in solutions that “are not grounded in reality” or focus on outdated, unlikely threats. It is important for decision-makers to understand that responsiveness is a systemic issue and compliance is not solely to blame. Organizational behaviors and cultural factors that reinforce responsiveness and agility also may have a significant impact on security.

Organizational factors are just as important. Compliance security gaps exist

and complementary measures can help. But when they are not planned for, as we see in Section 7.2.4, complementary measures may fail to fill security gaps or induce new problems. Complementary measures need to be routinely enforced, must align with actual security needs, and organizations need to prepare for end-to-end use cases prior to implementation.

Here is a slight modification to an old adage: if a thing is worth doing, it is worth doing well and routinely. Throughout Section 7.2.4, we find instances of organizations investing in complementary measures but not following through on required security tasks such as: buying new security platforms and failing to actually use them, missing critical security events because analysts failed to check logs, or delaying the approval of a much-needed security platform because a change review board missed their response deadline. Routine checkups, reassessments, and deadlines can help eliminate these problems.

Section 7.2.4 shows that organizations chose to spend money on new complementary controls rather than understand why their current security strategy failed. Participants highlighted instances where security platforms were misconfigured due to a lack of training and, instead of triaging the problem to identify the training or configuration deficiency, organizations were more likely to spend more money instead of optimizing the platforms already in use. Organizations should prioritize security requirements and expenditures against actual needs, based on business goals, requirements, and threats.

Lastly, organizations need to have support in place prior to deploying complementary measures. Section 7.2.4 reports instances of complementary measures

generating massive amounts of information (e.g., threat intelligence or security logs) without support in place, such as a place to store the data. There were also organizations that failed to train their employees on how to use new platforms or hired employees with mismatched skills for tasks such as threat hunting. Prior research suggests that organizations should sufficiently resource usable solutions, invest in employee training, and support the end-to-end use-case requirements for security measures [8, 60, 72, 87, 110, 117, 191, 202, 203].

While compliance is a critical aspect of an organization's digital security, it is far from a panacea. Organizations know that gaps in compliance exist, and therefore solely relying on compliance to drive a defensive security posture exposes the organization and its users to risk of an attack. Thus, organizations create complementary measures that go beyond compliance; however, these efforts face organization inertia and risk. These results may spur development of improved compliance mandates that acknowledge their shortcomings and provide guidance on how to evolve a security posture against emerging threats.

Chapter 8: Discussion

The work in this thesis provides insight into the effectiveness of proactive measures for security professionals in real-world environments. The proactive benefits of planning, baselining, and implementing security controls can provide organizations with measurably improved security, but it is essential for support to be integrated throughout an organization's culture.

8.1 Better together

The observations in this thesis suggest that proactive measures provide optimal benefits when employed in a complementary manner.

Proactive plans to mitigate risk provided tangible improvements to the security posture of NYC3, but gaps remained. Chapter 3 revealed that, prior to using threat modeling, compliance programs played a foundational role in how NYC3 and its security technicians implemented baseline security controls. However, in this study, participants did not consider security gaps caused by compliance standards into their threat model. Chapter 4 highlights the insight that organizations, like NYC3, could benefit from through systematic evaluations of their compliance programs and understanding how mandated baseline digital security measures may (or may not)

affect their overall security posture. Additionally, proactive planning through threat modeling allowed NYC3 security professionals to identify risk and develop mitigation strategies. Mitigation, however, does not mean organizations are now immune from attacks, as evident in Section 3.3.5. Once newly-implemented defensive measures detected on-going exploitation attempts, NYC3 incident responders from the SOC executed ad hoc remediation efforts — but failed to pre-develop response plans for likely threats. These findings suggest that the outputs generated from planning efforts (such as threat modeling) should become mandatory inputs for likely threats that need to be planned for through incident response playbooks. As Chapter 7 suggests, organizations must adopt mechanisms to support playbook development (e.g., allocating time, personnel, and requisite resources) as well as organizational behaviors that reinforce these habits. Section 8.2 expands on this notion below.

Baseline security mechanisms allow organizations to establish minimal security to perform essential business tasks, but alone are insufficient to protect organizations from targeted attacks and proper baselining may actually induce their own security issues. Chapters 4 and 5 revealed particular issues that may arise through perfect compliance and also several threat models not considered by an exemplar compliance program. This suggests that organizations need to process discovered security gaps and compliance complications through a planning framework such as threat modeling to account for newly-identified issues. Additionally, Chapter 7 details the impact of outdated policies and technical controls and how they can inhibit security benefits that could be gained from compliance programs and complementary measures. For this reason, both organizations and compliance authors should employ

more agile mechanisms that allow for periodic reviews and updates of policies in ways that promote security and minimize disruption [178].

Proactively implementing security controls and policies through incident response playbooks and other complementary measures may provide valuable boosts to security readiness, but they must be organizationally anchored and focused against realistic threats. Participants in Chapter 7 discuss the numerous frustrations, wasted resources, and overall inefficiencies associated with organizational emphasis on security platforms that do not fulfill actual requirements and siphon needed resources from other security efforts for support. Similarly, incident response playbooks require significant organizational change and emphasis to support tailored security requirements, as seen in Section 6.3.2.3. Organizations should avoid developing playbooks for low-impact scenarios that are either unlikely or negligible in severity. Planning frameworks increase the likelihood that complementary measures and incident response playbooks meet actual organizational requirements. Additionally, organizations must ensure complementary measures (including playbooks) are compatible with mandated compliance standards, otherwise conflicting, non-compliant actions could inadvertently be reinforced (as seen in Section 7.2.4). As further discussed in Section 8.2 below, organizations should strive to implement security solutions that enhance usability, provide shared understanding for technicians and leaders, and allocate time for training and rehearsals. As Chapter 6 shows, playbooks were only able to benefit security technicians of all experience levels after multiple no-notice events, training sessions, and revision sessions that required dedicated emphasis from the SOC.

Understanding that proactive measures are better when used in harmony together, Section 8.3 recommends an end-to-end workflow based on observations throughout this thesis.

8.2 Impacts of organizational culture

This thesis elucidates the required, holistic integration of security within organizational culture to help ensure proactive measures take hold. This is rooted in foundational knowledge from behavioral psychology and organizational culture [178], and Chapter 7 anecdotally confirms the sub-optimal consequences of non-integration.

Organizations should reinforce security tasks and habits through peer collaboration and mentorship programs. Peering, as described in Chapters 3, 6, and 7, allows small teams to challenge one another and assist with task completion. A clear benefit of peer collaboration is that similarly-experienced employees can share knowledge, reinforce their own understanding of security tasks, and build improved bonds within their security team [178]. Similarly, mentorship programs provide opportunities for junior employees to partner with mid- to senior-level security professionals of a similar work role to gain a new perspective on the significance of task completion (e.g., how failure can affect the organization), learn complementary skills, and help shape junior professionals' career goals [178]. Whether peer collaboration and mentorship programs are informal or part of a formal initiative, organizations must clearly communicate program expectations to their employees, preserve time for employees to participate in the program, establish incentives for

participation, and, when necessary, establish deterrence for non-participation. Partnered organizations in Chapters 3 and 6 both integrated peer collaboration and mentorship into periodic performance reviews and assessments. Monthly and quarterly periodic performance reviews allowed supervisors to monitor participation and nudge closer collaboration before annual assessments; annual assessments served as “report cards” for program participation as well as overall performance.

Perceived efficacy and perceived ownership shape adoption rates. The theory of planned behaviour (TPB) should also be considered when implementing proactive measures, specifically when attempting to improve adoption rates. For proactive measures to take hold within an organization, they must decrease the difficulty of performing security tasks and organizations must actively reduce as many barriers as possible that may impede task accomplishment. These two factors directly shape employees’ perceived behavioral control and the likelihood they will continue using the proactive measures [6], as seen in Chapters 3, 6, and 7. In Chapter 3, participants that used threat modeling reported high levels of perceived usability: threat modeling improved their job performance, it was easy to use, it was useful for their job. Additionally, senior leaders at NYC3 deliberately set aside time for employees to perform threat modeling and established a system for tracking known security issues derived from threat modeling, which allowed employees to feel like their recommendations mattered and helped improve the organization’s security. Employees at NDC had similar sentiments in Chapter 6, understanding that dedicating time towards refining playbooks would have positive outcomes and allow them

to accomplish tasks during stressful events more quickly. In Chapter 7, it became evident that organizations needed to streamline the use of proactive measures for employees, especially for platforms that required significant financial and/or personnel investments. Poorly-managed programs suffered while well-resourced proactive measures were embraced by employees and organizations alike.

On-boarding training and periodic skill assessments proved to be valuable for security professionals and organizations. Introducing threat modeling and playbook development to new employees allowed them to better understand their role in the organization, their responsibilities, and the impact of failing to perform their security tasks. Periodic skill assessments, such as live security exercises, allowed employees to reinforce skills learned from on-the-job or formal training means and allow organizations to evaluate both the effectiveness of their personnel and training programs [22].

Organizations should be proactive and avoid constant reactive changes. Quickly realigning security priorities can have severe second- and third-order impacts on organizations. As seen in Chapter 7, participants discussed the varying ways and reasons that organizations chose to implement complementary measures after a security incident occurred — often with urgency. Prior research suggests it is in organizations’ best interests to be as proactive as possible; constant, reactive realignments in security priorities can negatively impact overall business performance [37, 82]. Funds have to be reallocated for unforecasted security expenditures

and may harm daily business operations elsewhere [12, 187]. Embracing proactive measures and minimizing significant shifts in security priorities may counter these negative impacts.

Bi-lateral communication between leaders and technicians is key. Dis-jointed or poorly communicated defensive plans are resented by those that have to implement them. Chapter 7 shows that uninformed decisions by executives (made in isolation away from security teams) often result in the implementation of proactive measures that do not adequately address actual requirements or do not support the usability needs of the workforce that will actually be using the platform. These observations accord with previous behavioral psychology research and shows the importance of bidirectional communication with front-line to supervisors [229]. Conversely, open communication and transparency between supervisors and security professionals lead to trust and the ability to adjust proactive measures to best fit organizational requirements — thus, enhancing effectiveness [221]. In Chapters 3, participants' self-value to the organization improved because their organization incorporated their insight into defending New York City. Threat modeling provided an open-communication model that allowed professionals to identify problems, recommend solutions, and watch with full transparency as their organization implemented solutions to address the problems. In Chapter 6, playbooks initially provided little benefit to senior or junior technicians. The organization reviewed failures, implemented internally-recommended changes, and built comprehensive response plans that met their specific requirements. In Chapter 7, organizations that encouraged

open-communication were more likely to implement well-fit solutions to meet specific security requirements; if problems arose, communication between leaders and technicians allowed the organization to quickly adapt solutions for a better fit.

Cultural reinforcements need to consider how adults learn best. Organizations should consider best-practices from behavioral psychology and behavioral change when considering security training. The experiential learning theory outlines that adults can better learn new concepts by integrating new concepts into existing ones, experiencing concepts first-hand, reflecting on new behaviors, and repetition of new concepts [118]. Social learning theory helps explain particular learning practices in digital security, showing that professionals draw linkages between behaviors and consequences [122, 208]. Organizations need to maximize this, allowing employees to understand the direct benefits of new behavior patterns [21] — specifically, that employing proactive security measures has positive outcomes on their organization and quality of work. Chapters 3 and 6 highlight the effectiveness of these learning strategies. Employees at New York City Cyber Command were able to quickly learn threat-modeling techniques given a short educational intervention and an opportunity to apply their new knowledge in a one-on-one setting. Junior National Defense Center employees required repeated opportunities to apply new concepts and understand the positive impacts that quick-response incident response playbooks can have on their organization. In both instances, tailoring training strategies towards adult learning helped reinforce security training.

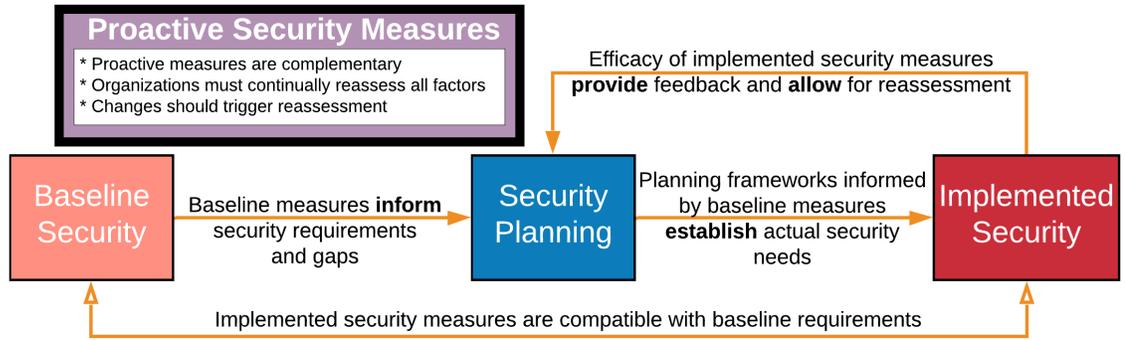


Figure 8.1: Visualization of the recommended workflow that organizations should follow to benefit from proactive measures

8.3 Proposed workflows for proactive measures

The anecdotes and observations from Chapter 7, when coupled from lessons learned from Chapters 3, 4, 5, and 6, provide a blueprint for how organizations can effectively employ proactive measures in a complementary manner. This section describes the proposed workflow visualized in Figure 8.1 in greater detail.

Upfront, organizations should integrate and emphasize sound security practices throughout all business operations. Senior organizational leaders, mid-level managers, and security professionals should all have active roles throughout the employment of proactive measures.

Organizations must analyze baseline requirements to first understand digital constraints (the things they cannot do) and restraints (the things they must do). The analysis of mandatory policies and technical controls should then be used to inform organizations' planning efforts.

Next, organizations should execute planning with an emphasis on which threats

are mitigated by baseline requirements, which threats may be created by baseline requirements, and unaddressed threats. By using planning frameworks such as threat modeling, organizations can begin to understand their actual security requirements and develop strategies for improving their overall security posture — as well as how the organization can best holistically support these strategies.

After prioritizing their security needs with consideration to organizational goals, organizations can begin implementing security measures that adequately address threat exposure. These security measures are meant to mitigate threats and organizations should have plans in place for when threat actors attack. Incident response playbooks based on actual threat exposure should be developed to enhance readiness. Playbooks help security professionals establish required organizational support structures and technical systems before they are needed. Readily-available and well-rehearsed playbooks support security professionals with step-by-step prompts to help establish momentum during stressful events against threats that they are likely already familiar with — because they participated in the baseline analysis and planning process. Organizational leaders are prepared because they were integrated in the planning and response efforts. Whether the incident was triggered by friendly red teams or actual adversaries, organizations should evaluate the efficacy of implemented security measures and iteratively integrate that feedback into follow-on planning efforts. Negative feedback can be helpful in updating threat models and better tailoring security mechanisms towards the organization’s actual security requirements.

Proactive measures should not be a “one-and-done” effort. Any changes in key

factors such as compliance updates, emerging threats, or performance feedback for proactive measures should trigger a re-evaluation of each factor. Proactive measures, if adequately supported and emphasized by organizations, should be continual efforts to enhance the robustness of their security program.

Further research is needed to confirm the overall efficacy of this proposed workflow. Various factors such as organizational size, budget, and maturity should be evaluated to determine the effect, if any, on the overall process.

Chapter 9: Conclusion

Proactive measures can provide organizations with the tools necessary to improve their digital security posture before threats manifest.

As I show in Chapter 3 and 7, organizations — with the best of intentions — may employ security measures ineffectively without anchoring requirements against likely threat exposure. Planning frameworks such as threat modeling can provide much-needed structure when codifying threat exposures and developing methods to mitigate risk.

In Chapters 4, 6, and 7, I find that organizations may need to dedicate additional effort towards engineering solutions that adequately address their specific security requirements. Additionally, organizations that sufficiently allocate resources, emphasize training, and encourage open communication may have an agile advantage when responding to threats or iteratively improving their security posture.

In this thesis, I establish that organizations that provide essential services such as utilities or financial transactions have additional hardships placed upon them. In Chapters 4, 5, and 7, I find that mandatory compliance programs enforced by governments or regulatory entities may be outdated, lack sufficient detail for implementation, offer conflicting information, or may create other sub-optimal

security conditions when followed “by-the-letter.” Organizations are left to evaluate the impacts of compliance programs and develop their own strategies for becoming more secure.

Taken together, these results suggest that proactive measures can be optimized to provide greater benefits to organizations’ security operations, but conditions are still far from perfect. Organizations, governments, and regulatory entities all need to adapt their standard business practices to support relevant, responsive security programs that adequately protect against the constantly evolving threat landscape.

Appendix A: Study Instruments

A.1 Survey Questions from Chapter 3

A.1.1 Pre-intervention survey

This survey will ask for information about your current work role and your personal assessment of IT security. Please be as candid and detailed as possible.

1. Please provide your DoITT email address. We will use a SHA1 hash of your email address as your unique identifier throughout the study. We will make every effort to protect your privacy and keep your responses confidential. We will report data in the aggregate, thus no individual will be identified. (short answer)

2. What division do you work for? [Citywide Cybersecurity, Other DoITT Division]

3. What is your division's mission or objective? [long answer]

4. What group do you work for? [OPS, Engineering, Identity Management, Other within Citywide Cybersecurity, Other within DoITT]

5. What is your group's mission or objective? [long answer]

6. How many people work in your group? [number answer]

7. What is your position / duty title? [short answer]

8. What are your responsibilities? [long answer]

9. What technology assets are involved in accomplishing your group's mission or objective? An asset can be a device you defend, a system, a service, a tool, or any form of technology relevant to your work, not just security assets. Please list at least three. [form-style short answers]

10. Prioritize the importance of these assets in descending order based on your personal assessment. The #1 item should be the asset you perceive to be the most important. [drag-and-drop answer]

11. Consider what you ranked as the Top 3 critical assets. How are these three most critical assets defended? [form-style long answers]

12. Do you feel that the current defensive measures are sufficient for #Loop Q11 answers# [(Definitely yes) 1-5 (Definitely No)]

13. When building a defense plan, how do you determine how to defend assets? [long answer]

14. Is there a program/guideline/checklist for assessing your digital security posture? If yes, please describe these items and describe how frequently they are used, if relevant.

15. Currently, how effective is Citywide Cybersecurity at defending against, mitigating, and responding to digital security threats? [(Not effective at all) 1-5 (Extremely effective)]

16. Currently, how effective is your group at defending against, mitigating, and responding to digital security threats? [(Not effective at all) 1-5 (Extremely effective)]

17. Do you feel there are single points of failure for IT defense within DoITT?
[(Definitely yes) 1-5 (Definitely No)]
18. #If yes# Please describe any possible points of failure within DoITT.
[long answer]
19. Please rate your knowledge of center of gravity before this study. [(Extremely knowledgeable) 1-5 (Not knowledgeable at all)]
20. Please rate how certain you are that you can do the digital security tasks below by using the sliding bars. Rate your degree of confidence by recording a number from 0 to 100 using the scale given below [(Cannot do at all) 0-50 (Moderately can do) 50-100 (Highly certain can do)]: Monitor critical assets, Identify Threats, Mitigate Threats, Proactively Respond to Intrusions and Eliminate Threats
21. How many years of IT experience do you have in large organizations?
[number answer]
22. How many years of IT experience do you have in small organizations?
[number answer]
23. How many years formal or informal training have you received? This includes schooling, industry training programs, and work-specific training programs.
[number answer]
24. Do you have any industry certifications? [Yes, Yes but expired, No]
25. # If yes# Which certifications do you currently possess? [long answer]
26. What is the highest level of school you have completed or the highest degree you have received? [Less than high school degree, High school graduate (high school diploma or equivalent including GED), Some college but no degree,

Associate degree in college (2-year), Bachelor's degree in college (4-year), Master's degree, Doctoral degree, Professional degree (JD, MD), Prefer not to answer]

27. Please specify the gender with which you most closely identify. [Male, Female, Other, Prefer not to answer]

28. Please specify your age. [18-29, 30-39, 40-49, 50-59, 60-69, Over 70, Prefer not to answer]

29. Please specify your ethnicity. [Hispanic or Latino; Black or African American; White; American Indian or Alaska Native; Asian, Native Hawaiian, or Pacific Islander; Other; Prefer not to answer]

A.1.2 Post-intervention survey

This survey will ask for information about your current work role and will require you to conduct a threat-based assessment on security using the center of gravity framework. Please be as candid and detailed as possible.

1. Please provide your DoITT email address. We will use a SHA1 hash of your email address as your unique identifier throughout the study. We will make every effort to protect your privacy and keep your responses confidential. We will report data in the aggregate, thus no individual will be identified. [short answer]

2. What technology assets are involved in accomplishing your group's mission or objective? An asset can be a device you defend, a system, a service, a tool, or any form of technology relevant to your work, not just security assets. Please list at least three. [form-style short answers]

3. Based on your list for the previous question, how did the list of technology assets that accomplish your mission or objective change from the initial survey you took? [More answers, Same number of answers, Fewer answers, I cannot remember]
4. #If more or less# Why did your answers change? [long answer]
5. Of these previously listed assets, which is your group's Center of Gravity (CoG) for accomplishing its mission? As a reminder, the center of gravity is the hub in which other assets derive their power or capability. [short answer]
6. What are the critical capabilities associated with #COG from Q5#? List at least three.
7. What are the critical requirements of #Loop Q6 answers#? [form-style short answers]
8. Please list any critical vulnerabilities or vulnerable conditions associated with #COG from Q5#. Please use thresholds when appropriate. An example threshold would be "packet loss consistently exceeding 50%." [form-style short answers]
9. What are the critical requirements for an adversary or threat actor to exploit #Loop Q8 answers#? [form-style short answers]
10. Does Citywide Cybersecurity currently have sufficient active defenses or sensors defending against #Loop Q8 answers#? [Yes, No, Unsure]
11. Does Citywide Cybersecurity currently have sufficient passive defenses or sensors defending against #Loop Q8 answers#? [Yes, No, Unsure]
12. #If yes for Q10 or Q11# You indicated #Loop Q8 answers# did not have a sufficient active or passive defense in place. How would you recommend defending

against it. [long answer]

13. Did you consider this defense plan before conducting the CoG exercise?

[Yes, No]

14. Think back to your previous assessment of Citywide Cybersecurity's effectiveness in defending against, mitigating, and responding to digital security threats from the first survey. After learning about CoG, how effective is Citywide Cybersecurity at defending against, mitigating, and responding to digital security threats currently? [(Not effective at all) 1-5 (Extremely effective)]

15. Think back to your previous assessment of your group's effectiveness in defending against, mitigating, and responding to digital security threats from the first survey. After learning about CoG, how effective is your group at defending against, mitigating, and responding to digital security threats currently? [(Not effective at all) 1-5 (Extremely effective)]

16. The questions in this section will gauge how certain you are that you currently can perform various digital security tasks. Answer these questions without applying any concepts of center of gravity. Please rate how certain you are that you can do the things below by using the sliding bars. Rate your degree of confidence by recording a number from 0 to 100 using the scale given below [(Cannot do at all) 0-50 (Moderately can do) 50-100 (Highly certain can do)]: Monitor critical assets, Identify Threats, Mitigate Threats, Proactively Respond to Intrusions and Eliminate Threats

17. In general, did your responses about how certain you are that you can perform various digital security tasks change since the first survey? [Yes, Unsure,

No]

18. #If yes# Why did your response change after learning about CoG?

19. Please rate your current knowledge of center of gravity. [(Extremely knowledgeable) 1-5 (Not knowledgeable at all)]

20. The questions in this section will gauge how certain you are that you can perform various digital security tasks using center of gravity. Answer these questions as if you completely redesigned, prioritized, and implemented defensive measures only based on the center of gravity framework. Please rate how certain you are that you can do the things below by using the sliding bars. Rate your degree of confidence by recording a number from 0 to 100 using the scale given below [(Cannot do at all) 0-50 (Moderately can do) 50-100 (Highly certain can do)]:
Monitor critical assets, Identify Threats, Mitigate Threats, Proactively Respond to Intrusions and Eliminate Threats

21. Please identify the center of gravity node in this undirected graph. [0-10 from image]

22. Identify the center of gravity for US Pacific Theater Forces within the following passage: [US amphibious forces (Marines/transport craft), Logistic resupply assets, US naval aircraft carrier]

23. Which of the following is an example of a critical capability for a local area network? [A router, Providing internode connectivity, An administrator]

24. Which of the following is an example of a critical requirement for a local area network? [Blocking unwanted connections, Providing connections to other networks, Network interface cards]

25. Which of the following is an example of a critical vulnerability for a local area network? [Users without cyber awareness training, Forwarding emails to co-workers, Loss of network connectivity due to a denial of service attack]

26. To your knowledge, was #COG from Q5# ever targeted during a cyber attack or intrusion? [Yes, No, Possibly]

27. #If yes# Did this attack or intrusion impact your choice to select #COG from Q5# as the CoG? Please describe in an UNCLASSIFIED manner. Include details about attack frequency, impact, or any other applicable details. [long answer]

A.1.3 Follow-up survey

This survey will assess your understanding and integration of Center of Gravity within your work. Please be as candid and detailed as possible.

1. Please provide your DoITT email address. We will use a SHA1 hash of your email address as your unique identifier throughout the study. We will make every effort to protect your privacy and keep your responses confidential. We will report data in the aggregate, thus no individual will be identified. [short answer]

If you would like to access a digital copy of the center of gravity worksheet, it is available for viewing here: <https://goo.gl/icVMLX>

2. Please rate your current knowledge of center of gravity. [(Extremely knowledgeable) 1-5 (Not knowledgeable at all)]

3. The questions in this section will gauge how certain you are that you

can perform various digital security tasks using center of gravity. Answer these questions as if you completely redesigned, prioritized, and implemented defensive measures only based on the center of gravity framework. Please rate how certain you are that you can do the things below by using the sliding bars. Rate your degree of confidence by recording a number from 0 to 100 using the scale given below [(Cannot do at all) 0-50 (Moderately can do) 50-100 (Highly certain can do)]: Monitor critical assets, Identify Threats, Mitigate Threats, Proactively Respond to Intrusions and Eliminate Threats

4. In general, did your responses about how certain you are that you can perform various digital security tasks using CoG change since the second survey? [Yes, No, Unsure]

5. #If yes# Why did your response change after 30 days? [long answer]

6. Have you applied concepts of center of gravity within your work? If yes, how? If no, why not? [long answer]

7. In Part 2 of this study, you completed a center of gravity worksheet that was tailored to your specific work role. You identified critical vulnerabilities, potential threat capabilities against those vulnerabilities, and proposed an ideal defense plan. Have you used any of the information from this worksheet since completing it? If yes, how? If no, why not? [long answer]

8. Have you applied concepts from center of gravity anywhere else in your life? If yes, how? If no, why not? [long answer]

9. Please rate the following statements based on your current understanding of CoG [Strongly agree Agree, Somewhat agree, Neutral, Somewhat disagree, Disagree,

Strongly disagree]: Using center of gravity improves the quality of the work I do; Using center of gravity gives me greater control over my work; Center of gravity enables me to accomplish tasks more quickly; Center of gravity supports critical aspects; Center of gravity increases my productivity; Center of gravity improves my job performance; Center of gravity allows me to accomplish more work than would otherwise be possible; Center of gravity enhances my effectiveness on the job; Center of gravity makes it easier to do my job; Overall, I find center of gravity useful in my job; Center of gravity enhances the way I think about digital security

10. Which of the following is an example of a critical capability for a the center of gravity website content server? [A database, Providing users with high website availability, An administrator]

11. Which of the following is an example of a critical requirement for the critical capability providing communication with partnered agencies? [Email servers, Enforcing multi-factor authentication, Blocking external access to FTP servers]

12. Which of the following is an example of a critical vulnerability for the critical requirement local area network connectivity? [Cyber awareness training, Loss of network connectivity due to a denial of service attack, Forwarding emails to co-workers]

13. Please identify the center of gravity node in this undirected graph. [0-10 from image]

A.1.4 NYC leadership panel questions

We asked our panel of NYC3 leaders to answer the following questions for each participants' post-training survey results.

1. How likely is the identified asset the critical enabler for the participant's responsibilities? Please use a scale from 0 to 5, with 0 being "extremely unlikely" and 5 being "extremely likely"
2. How likely would the identified vulnerabilities stop the participant from fulfilling their responsibilities? Please use a scale from 0 to 5, with 0 being "extremely unlikely" and 5 being "extremely likely"
3. How likely would the identified threats exploit the vulnerabilities and prevent mission fulfillment? Please use a scale from 0 to 5, with 0 being "extremely unlikely" and 5 being "extremely likely"
4. How likely would the plan of action mitigate threats from exploiting the critical vulnerabilities? Please use a scale from 0 to 5, with 0 being "extremely unlikely" and 5 being "extremely likely"
5. Is the proposed defense plan sufficiently detailed to implement? Please respond with yes, no, or unsure.

A.2 Expert Survey from Chapter 4

Participant is presented with consent form; Please check all that apply (you may choose any number of these statements): I am age 18 or older; I have read this consent form or had it read to me; I voluntarily agree to participate in this research

and I want to continue to the survey.

Introduction: This survey will ask for you to assess the validity of an independent evaluation of [standard name]. Please be as candid and detailed as possible.

For each issue, please confirm:

1. If your organization followed the standard as written and nothing else, would your organization be vulnerable to this issue? (Yes/No/Possibly)
2. If yes or possibly \Rightarrow In your opinion, what is the likelihood of this vulnerability being exploited if standard is followed as written and nothing else? (Frequent - Occurs often, continuously experienced; Likely - Occurs several times; Occasional - Occurs sporadically; Seldom - Unlikely, but could occur at some time; Unlikely - Can assume it will not occur)
3. If yes or possibly \Rightarrow In your opinion, what is the severity associated with exploitation if standard is followed as written and nothing else? (Catastrophic - Complete system loss, major property damage, full data breach, corruption of all data; Critical - Major system damage, significant property damage, significant data breach, corruption of sensitive data; Moderate - Minor system damage, minor property damage, partial data breach; Negligible - Minor system impairment)
4. If yes or possibly \Rightarrow Is there past evidence of this vulnerability within your organization? (Yes/No/Maybe)
5. If yes or possibly \Rightarrow What would you recommend, based on your experience,

to remedy this issue? (Open response)

6. If no \Rightarrow What additional policies, procedures, or defensive techniques does your organization use to mitigate this issue? (Open response)

End of survey: Does your organization allow waivers to the compliance standard? If yes, how frequently are they used? If no, does frequently does this create issues for your organization?

Demographics: What is the highest level of school you have completed or the highest degree you have received? Please estimate the number of years experience you have in the compliance and information technology fields. Please describe your work role and your interaction with compliance standards. Please estimate the organization size that you work in.

A.3 Survey instruments for Chapter 6

A.3.1 Design phase

After using the FRAMEWORK playbook design framework, please answer the following:

(Matrix of options for Likert scale: Strongly Disagree, Somewhat Disagree, Neutral, Somewhat Agree, Strongly Agree)

I think that I would like to use this framework frequently.

I found this framework unnecessarily complex.

I thought this framework was easy to use.

I sufficiently provided all information to complete this step.

I think that I would need assistance to be able to use this framework.

I found the various functions in this framework were well integrated.

I thought there was too much inconsistency in this framework.

I would imagine that most people would learn to use this framework very quickly.

I found this framework very cumbersome/awkward to use.

I felt very confident using this framework.

I needed to learn a lot of things before I could get going with this framework.

I think other experts could use this framework easily.

I think other non-experts could understand products from this framework easily.

I think other non-experts could use this framework easily.

Rate each step in order of importance for completing the playbook, with #1 being the most important step. [Drag and drop list of steps based on framework]

Please explain why you ranked [TOP CHOICE] step most important. [open response]

Please explain why you ranked [LOWEST CHOICE] step least important. [open response]

Please provide any positive feedback you may have on using the [FRAMEWORK] playbook design framework.

Please provide any negative feedback you may have on using the [FRAMEWORK] playbook design framework, especially any parts that you felt were confusing or needed additional information.

Please provide any neutral feedback you may have on using the [FRAMEWORK] playbook design framework. Do you feel anything was missing? Anything that could be better designed?

Demographics:

What is the highest level of school you have completed or the highest degree you have received?

Please estimate the number of years experience you have in the digital security and information technology fields:

Please indicate which role most accurately reflects your current position:

Please estimate the organization size that you work in:

A.3.2 Evaluation phase

Is this playbook sufficiently detailed to implement and actually detect the event? [Yes, no, unsure]

How likely is the playbook to adequately respond to the scenario event (with 1 being least likely)?

Please explain why this playbook would or would not adequately respond to the event. [open response]

Are there errors in the provided playbook that would hinder response efforts? [Yes, no, unsure]

Are there critical elements of a response plan missing from the playbook? [open

response]

Do you have any other feedback for this playbook? Explain. [open response]

Demographics:

What is the highest level of school you have completed or the highest degree you have received?

Please estimate the number of years experience you have in the digital security and information technology fields:

Please indicate which role most accurately reflects your current position:

Please estimate the organization size that you work in:

A.3.3 Implementation phase

Based on your experiences, please indicate which framework was more useful for each task:

(Matrix ranging from NIST much better, NIST better, no difference, IACD better, IACD much better)

Identifying assets at risk:

Identifying required response tasks:

Building a comprehensive plan:

Being easily understandable:

Being easily implementable:

Please provide any positive feedback you may have for NIST with respect to imple-

menting a playbook. [open response]

Please provide any negative feedback you may have for NIST with respect to implementing a playbook. [open response]

Please provide any positive feedback you may have for IACD with respect to implementing a playbook. [open response]

Please provide any negative feedback you may have for IACD with respect to implementing a playbook. [open response]

Were there any unexpected modifications you had to make to implement your plan using NIST? [open response]

Were there any unexpected modifications you had to make to implement your plan using IACD? [open response]

Please rate the following statements based on your current overall understanding of digital security playbooks:

(Matrix of options for Likert scale: Strongly Disagree, Somewhat Disagree, Neutral, Somewhat Agree, Strongly Agree)

Using digital security playbooks improves the quality of the work I do.

Using digital security playbooks gives me greater control over my work.

Digital security playbooks enables me to accomplish tasks more quickly.

Digital security playbooks supports critical aspects of my job.

Digital security playbooks increases my productivity.

Digital security playbooks improves my job performance.

Digital security playbooks allows me to accomplish more work than would otherwise

be possible.

Digital security playbooks enhances my effectiveness on the job.

Digital security playbooks makes it easier to do my job.

Overall, I find digital security playbooks useful in my job.

Digital security playbooks are confusing or unintuitive.

A.3.4 Utilization phase

From intrusion event to detection, how much time do you assess passed? How did you determine an event occurred? [open response]

Were there any unexpected issues associated with detecting the event? What decisions did you have to make during detecting the event that were not covered in the playbook? [open response]

From detection to initial response using a playbook, how much time do you assess passed? [open response]

Were there any unexpected issues associated with initial response? [open response]

What decisions did you have to make during responding to the event that were not covered in the playbook?[open response]

From initial response to threat neutralization, how much time do you assess passed? [open response]

How did you determine the event was stabilized/quarantined to a sufficient level?[open response]

Were there any unexpected issues associated with threat neutralization?[open response]

Please rate the following statements based on your current overall understanding of digital security playbooks:

(Matrix of options for Likert scale: Strongly Disagree, Somewhat Disagree, Neutral, Somewhat Agree, Strongly Agree)

Using digital security playbooks improves the quality of the work I do.

Using digital security playbooks gives me greater control over my work.

Digital security playbooks enables me to accomplish tasks more quickly.

Digital security playbooks supports critical aspects of my job.

Digital security playbooks increases my productivity.

Digital security playbooks improves my job performance.

Digital security playbooks allows me to accomplish more work than would otherwise be possible.

Digital security playbooks enhances my effectiveness on the job.

Digital security playbooks makes it easier to do my job.

Overall, I find digital security playbooks useful in my job.

Digital security playbooks are confusing or unintuitive.

A.4 Interview guide

For each survey response across all phases that required follow-up questions:

In your survey, you indicated [TOPIC]. Could you please explain more information about [TOPIC]?

For each expert evaluation survey response required follow-up questions:

In your response, you indicated [TOPIC]. Could you please explain more information about [TOPIC]? How would you handle this in your organization? Have you encountered this situation in your organization before? Do you have any insight that would not necessarily be intuitive for people following playbook frameworks?

For our trusted insider during the utilization phase:

What time did you initiate your attack?

Were there any special considerations when you conducted the attack?

What times did participants report detecting the attack?

How long until they initiated initial response actions?

How long did they take to neutralize the threat?

Were there any observations that stuck out to you?

For each participant during the utilization:

In your survey response, you indicated [TOPIC]. Could you please explain why you felt [TOPIC] presented a unique challenge? Was there anything that could have

prepared you more for [TOPIC]?

A.5 Survey instruments for Chapter 7

- Does your organization adhere to any form of mandatory compliance standard or regulatory controls?
- Which compliance standards does your organization deal with?
- Does your organization believe compliance is sufficient to protect your systems and data? (Yes, maybe no)
- Does your organization employ proactive security controls to address threats not covered by compliance programs?
- If yes, please describe unaddressed threats.
- Please select which of the following proactive controls your organization uses to complement (in addition to) compliance programs. Please do not select controls required by compliance programs that your organization follows. (Multiple selection options: Vulnerability disclosure or bug bounty programs Machine learning and other statistical analysis, Threat modeling, Tabletop security training exercises, Live security training exercises, Threat intelligence, Threat hunting (regular searches to ascertain the presence of a previously undetected adversary or compromise), Endpoint Threat Detection and Response solutions, Change control reviews/panels, Sandboxing, Zero clients/one-time-use systems, Integrity review of data and application updates, Periodic access

review, Multi-factor authentication, Multi-factor physical access, Hands-on training, On-the-job mentorship security training, On-the-job peer partnering training, Others: (fill-in-the-blank), None of the above)

- Enter loop for each item:
 - Is this security control required by compliance?
 - Why did you / your organization decide to implement this control?
 - On a scale from 1 to 5, with 5 being the highest rating, how well has this control worked out for your organization? (1-5)
 - What worked (or did not work) well about this control?
 - How do you ensure this measure is compatible with compliance controls (or if not, why not)?
 - How often do you reassess this control's effectiveness? (Daily, Once a week, One a month, Every few months, Yearly, Never)
- How does your organization prioritize which proactive measures you are going to invest in? What are the key factors?
- Demographics
 - What is the highest level of school you have completed or the highest degree you have received?
 - Please select the option that best categorizes your organization (Government/Defense, Entertainment, Financial services (payments, credit

cards), Consumer services (hotels, retail, sales), Critical services (power, water, etc), Healthcare, Agriculture/mining, Information technology, Education/Research)

- Please specify the job role that most closely reflects your employment position (Security Engineer, Security Analyst, Management, Compliance/Governance SME, Developer)
- Please estimate the number of years experience you have in the compliance and information technology fields.
- Please specify the estimated size of your organization.
- Please specify the estimated size of your constituency or clientele.
- How many organizations do you support?

Appendix B: Additional Data

B.1 Additional data from Chapter 3

We used the following two scenarios during our educational intervention training to communicate CoG analysis concepts to participants.

B.1.1 Star Wars walkthrough

The educational intervention instructor guided participants through this scenario, explaining the CoG analysis for the Galactic Empire. The Galactic Empire's desired end state is to provide peace and stability throughout the galaxy. To do this, their objective is to eliminate rebel forces. The Empire has many assets available for destroying the rebel scum to include: TIE fighters, stormtroopers, Darth Vader, and the Death Star. Of these assets, we know that the most powerful means for destroying planets and eradicating sources of rebellion is the Death Star; thus, it is the CoG analysis for the Empire. Critical capabilities for the Death Star include the ability to destroy planets. Critical requirements for this capability include Kyber crystals, engineers, and the superlaser. A critical vulnerability against the superlaser is accessible via a thermal exhaust port with an exterior opening. Threat

capabilities include the ability to fire weapons into the exhaust port and threat requirements include X-wing fighter aircraft. Given this scenario, an actionable defense plan for the Death Star would be concealing the thermal port or installing anti-aircraft turrets near the opening.

B.1.2 E-commerce scenario

In the second scenario, groups of participants applied CoG analysis without instructor assistance. The following examples are not exhaustive but include actual responses from the groups. This scenario was the first and only time participants completed CoG analysis analysis in a group setting.

We consider a small e-commerce business with the primary objective of maximizing profit and secondary objectives of customer satisfaction and website availability. We focus on defending assets that maximize our profits. The e-commerce business relies on a front-end webserver, a back-end database, redundant servers with load balancers, software developers, and a banking institution. Of the previously identified assets, the back-end database is the CoG analysis it conducts transactions with customers (the primary means for accomplishing our primary objective) and because of its interconnectedness with other assets. Critical capabilities for our business back-end database include (1) conducting atomic, consistent, isolated, and durable transactions, (2) permitting responsive queries from the front-end webserver, and (3) providing security safeguards for inventories and customer data. Critical requirements for providing security safeguards for inventories and customer

data would be (1) encrypted communication between customers, the front-end web-server, and the database; (2) encrypted sensitive data within the database; and (3) compliance with regulatory guidelines for business transactions. Examples of critical vulnerabilities would be continued use of software without periodically checking for updates and patching, such as continued use of OpenSSL 1.0.1 which is vulnerable to Heartbleed [176]. Threat capabilities against a vulnerable version of OpenSSL include conducting reconnaissance and network scans of vulnerable systems. Threat requirements include a valid exploit and payload against OpenSSL. A simple actionable defense plan for our running example includes (1) upgrading OpenSSL to a version that is patched against Heartbleed and (2) validating system performance post-upgrade.

B.1.3 Participant P17 example

Understand the end state and objective. Participant P17 is a security analyst who works within the NYC Security Operations Center (SOC). The SOC's defensive end state is maintaining an environment that is resilient and responsive to known and unknown threats. Based on P17's work role in NYC3, his personal objective is to defend workstations and respond to threats against the NYC3 environment.

Identify assets. P17 relies on network traffic inspectors, endpoint detection and response (EDR) solutions, and log aggregators to accomplish his objective. EDRs are tools for investigating suspicious activities throughout networks, hosts, and other

endpoints [43].

Identify the CoG. Of the previously identified P17 assets, the EDR is the CoG analysis because of its inherent ability to thoroughly protect systems across the enterprise, using input from network traffic inspectors and feeding log aggregators.

Identify critical capabilities (CC). P17's critical capabilities for EDR include blocking intrusion attempts, sending alerts, conducting queries, and quarantining infected systems.

Identify critical requirements (CR). CRs for P17 to block intrusion attempts include possessing updated indicators of compromise (IOCs) (i.e., threat signatures) and having the EDR agent installed on workstations.

Identify critical vulnerabilities (CV). P17 examples of critical vulnerabilities would be corrupted IOCs or workstation operating systems that are incompatible with a particular EDR application.

Enumerate threat capabilities (TC). With respect to our running example, representative TCs against corrupted updates include the ability to tamper with or man-in-the-middle IOC updates.

Enumerate threat requirements (TR). For P17, TRs include physical access or remote access to an update mechanism.

Develop an actionable defense plan (ADP). One mitigation strategy in P17's ADP verifies the integrity of updates from vendors before applying them to the EDR.

B.1.4 Visualizing Center of Gravity

Center of Gravity Worksheet

<p>Please state your work section's objective/mission:</p> <p>1</p> <p>What assets are used to accomplish this mission?</p> <p>2</p> <p>What is your center of gravity?</p> <p>3</p>	<p>Critical Capabilities</p> <p>4</p>
<p>Critical Requirements</p> <p>5</p>	<p>Critical Vulnerabilities</p> <p>6</p>
<p>Threat Capabilities</p> <p>7</p>	<p>Threat Requirements</p> <p>8</p>
<p>Defense Plan</p> <p>9</p>	

Figure B.1: Depiction of CoG analysis tabular method

Each participant received a printed version of the worksheet shown in Figure B.1 to help guide them through CoG analysis. Numbers indicate the order in which participants completed the form, as described in Section 3.1. Ad-

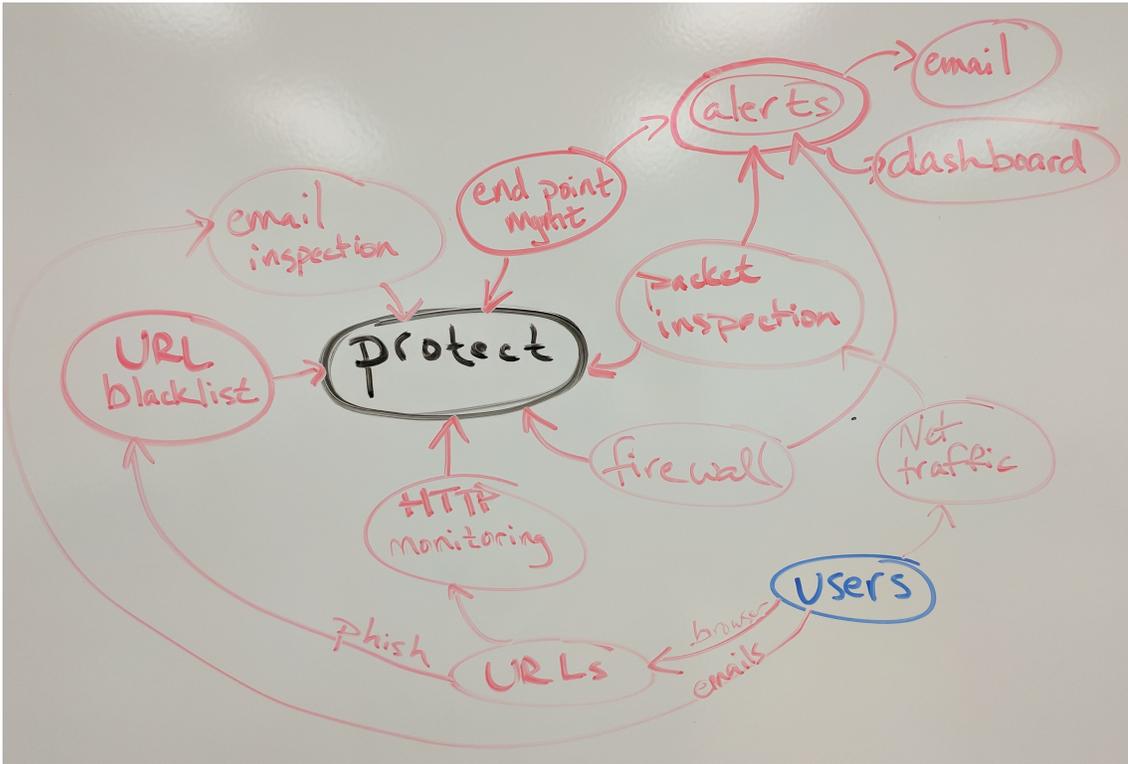


Figure B.2: Depiction of P18 visualizing his CoG analysis

ditionally, we provided participants with a digital version of this worksheet during all online surveys. A more detailed version of the worksheet is available at: <https://goo.gl/icVMLX>.

Some participants opted to use a whiteboard to visually depict their thought processes and building heterogeneous, relational linkages between nodes. As shown in Figure B.2, P18 began by writing his objective to “protect” networks. P18 then mapped how firewalls, EDRs, deep-packet inspection tools, and other defensive techniques support this objective. The commonality among all of these tools is that the defender uses cues from alerts to respond to incidents; thus, “alerts” are P18’s CoG.

B.1.5 CoG Identification Accuracy Regression

Variable	Value	Odds Ratio	CI	p-value
IT Exp.	0-5 yrs	–	–	–
	6-10 yrs	0.17	[0, 11.36]	0.408
	11-15 yrs	3.82	[0.26, 55.28]	0.325
	16-20 yrs	0.74	[0.04, 12.16]	0.83
	21-25 yrs	0.39	[0.01, 20.26]	0.643
	26+ yrs	0.26	[0, 60.44]	0.626
Edu.	Some College	–	–	–
	Associates	3.02	[0.03, 289.4]	0.634
	Bachelors	3.51	[0.25, 49.43]	0.352
	Graduate	4.64	[0.21, 100.14]	0.327
*Significant effect		– Base case (OR=1, by definition)		

Table B.1: Summary of regression over participants’ accuracy at identifying centers of gravity with respect to their years of experience and education. McFadden and Nagelkerke Pseudo R^2 measures are given for each regression.

B.2 Additional data from Chapter 4

B.2.1 Overall risk distribution from Chapter 4

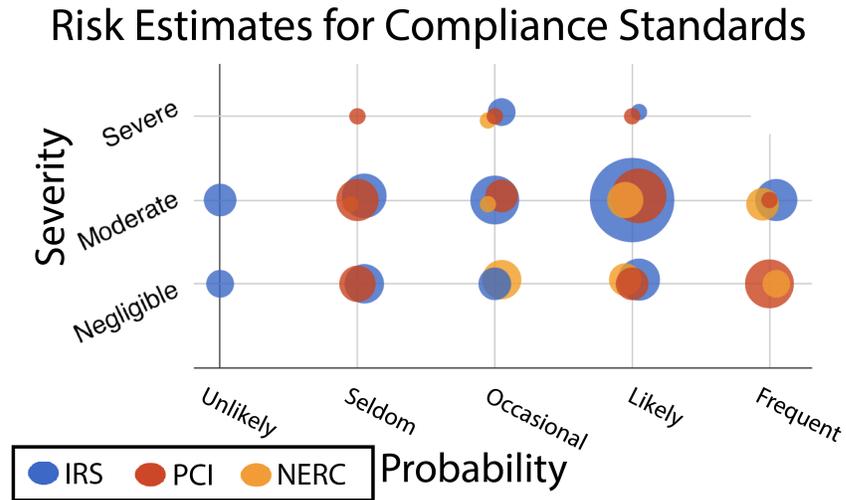


Figure B.3: Distribution of risk estimates by compliance standard

B.2.2 Compliance audit findings from Chapter 4

The following pages include the audit findings for IRS P1075, PCI DSS, NERC CIP 007-6, and FedRAMP.

Compliance Audit Results for IRS P1075

Section	Text of concern	Concern	Issue	Probability	Severity	Impact
1.3.2 Mailbox	The Safeguards Mailbox is a repository for information and communication to the Office of Safeguards relative to Safeguarding requirements and Publication 1075. The Mailbox is located at SafeguardReports@irs.gov. Below are items that are appropriate for submission to the Mailbox. Safeguards Reports and Extension Requests 45 Day Notifications Publication 1075 Technical Inquiries Re-Disclosure Agreements	Mailbox communication permits IRS to receive unencrypted sensitive information. IRS permits Secure Data Transfer (SDT) between gov agencies but offers nothing for common users. Example: https://www.dhs.gov/sites/default/files/publications/privacy/Guidance/handbookforsafe-guardingsensitivePII_march_2012_webversion.pdf http://www.law.com/sites/almstaff/2017/09/26/deloitte-hack-reveals-email-vulnerabilities-and-regulatory-gaps/?sreturn=20170909183037 Recommendation: Provide PGP public keys, use Keybase, and adopt another secure message transmission mode. Require encrypted attachments.	data vulnerability	Frequent	Moderate	Low
1.4.1 Federal Tax Information (FTI)	Safeguarding FTI is critically important to continuously protect taxpayer confidentiality as required by IRC 6103. FTI is a term of art and consists of federal tax returns and return information (and information derived from it) that is in the agency's possession or control which is covered by the confidentiality protections of the IRC and subject to the IRC 6103(p)(4) safeguarding requirements including IRS oversight. FTI is categorized as Sensitive But Unclassified information and may contain personally identifiable information (PII). FTI includes return or return information received directly from the IRS or obtained through an authorized secondary source, such as Social Security Administration (SSA), Federal Office of Child Support Enforcement (OCSE), Bureau of the Fiscal Service (BFS), or Centers for Medicare and Medicaid Services (CMS), or another entity acting on behalf of the IRS pursuant to an IRC 6103(p)(2)(B) Agreement. FTI includes any information created by the recipient that is derived from federal return or return information received from the IRS or obtained through a secondary source. FTI may not be masked to change the character of information to circumvent IRC 6103 confidentiality requirements.	As I continue reviewing, I find that the definition of FTI is missing several key items: 1. The definition doesn't specify if or how FTI can lose that designation. At no time will information cease to be FTI, as currently defined, but perhaps it should. Should protections be relaxed after 50 years? After the person dies? 2. Is there a point at which data be obfuscated and no longer require the FTI protections? Is it sufficient, for example, to sanitize 1000 records by removing SSN or giving each record a random identifier, and therefore no longer require FTI protections?	verbiage issue	Unlikely	Moderate	Low
1.4.3 Personally Identifiable Information	FTI may include Personally Identifiable Information (PII). FTI may include the following PII elements: Name of a person with respect to whom a return is filed Taxpayer mailing address Taxpayer identification number E-mail addresses Telephone numbers Social Security Numbers Bank account numbers Date and place of birth Mother's maiden name Biometric data (e.g., height, weight, eye color, fingerprints) Any combination of the above	*Name of a person with respect to whom a return is filed* should be modified to include any name of any person associated within tax document. The IRS tends to keep records indefinitely and as we've seen with attacks against OPM or Equifax, these records allow attackers to derive security content through PII history. Example: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.147.2471&rep=rep1&type=pdf Recommend: Expanding definition of PII to include any named individual in tax documents.	data vulnerability	Frequent	Moderate	Low
1.4.4 Information Received From Taxpayers or Third Parties	Copies of tax returns or return information provided to the agency directly by the taxpayer or his/her representative (e.g. W-2's, Form 1040, etc.) or obtained from public information files (e.g. federal tax lien on file with the county clerk, Offers in Compromise available for public inspection; court records, etc.) is not protected FTI that is subject to the safeguarding requirements of IRC 6103(p)(4). If the agency independently verifies FTI provided by the IRS or a secondary source with the taxpayer or a third party source, the verified information is not FTI as long as the IRS source information is replaced or overwritten with the newly provided information.	This section essentially states that the IRS is not responsible for securing information they receive from customers or from third-party companies. This is **NOT** okay.	data vulnerability	Frequent	Critical	Low
1.4.5 Unauthorized Access	*Unauthorized access occurs when an entity or individual knowingly or due to gross negligence receives or has access to FTI without authority*	This is excessively wordy. If the entity or individual has access, without authority, then it is **unauthorized access**. It should matter if it was done knowingly or due to negligence.	verbiage issue	Likely	Moderate	Low
1.4.6 Unauthorized Disclosure	An unauthorized disclosure occurs when an entity or individual with authorization to receive FTI knowingly or with gross negligence discloses FTI to another entity or individual who does not have authority, as defined in IRC 6103 and IRC 6104(c). An unauthorized disclosure has occurred when FTI is knowingly or due to gross negligence provided to an individual who does not have the statutory right to have access to it under the IRC. Even without willfulness or gross negligence FTI is not to be disclosed to entities or individuals who are not authorized by IRC 6103 to have it. Subject to the disclosure provisions of IRC 6103, agencies may need to disclose FTI to outside entities (e.g., for prosecution, appeals, or collection processes) as long as the receiving entity has a need-to-know and the individual recipient has authority under IRC 6103 to receive it. If the individual does not have a need-to-know, this constitutes an unauthorized disclosure.	> FTI is knowingly or due to gross negligence provided to an individual This section does not account for data breaches due to digital system vulnerabilities.	insufficient process	Likely	Critical	Low
1.4.7 Need to Know	Under need-to-know restrictions, even if an entity or an individual has the authority to access FTI, one would not be given access to such information if it were not necessary to perform his or her official duties with regard to the purpose for which IRC 6103 provides the FTI is to be used. Limiting access to individuals on a need-to-know basis reduces opportunities to "browse" or improperly view FTI. Restricting access to designated personnel minimizes improper access or disclosure. When FTI must be provided to clerical, computer operators, or others, these should only be provided the FTI that is essential to accomplish their official duties	This (and other similar sections) are written as very human centric and do not consider machines, automated analyses, or artificial intelligence. Section 2.4 does consider "statistical analysis, tax modeling, [and] revenue projections". Maybe this is a non-issue.	verbiage issue	Occasional	Negligible	Low
1.4.7 Need to Know - (Technical Controls)	Under need-to-know restrictions, even if an entity or an individual has the authority to access FTI, one would not be given access to such information if it were not necessary to perform his or her official duties with regard to the purpose for which IRC 6103 provides the FTI is to be used. Limiting access to individuals on a need-to-know basis reduces opportunities to "browse" or improperly view FTI. Restricting access to designated personnel minimizes improper access or disclosure. When FTI must be provided to clerical, computer operators, or others, these should only be provided the FTI that is essential to accomplish their official duties	In section 1.4.7 Pub 1075 mentions only providing information to users with a valid need to know. The document does not levy a formal requirement to implement technical controls to enforce need to know. The requirement is worded to be a goodwill based requirement only. If an agency is storing large quantities of FTI in a single system, then I would argue that an individual would not typically need access to the entirety of the data-set to perform their official job duties. I recommend instituting a formal requirement for receiving agencies to implement technical controls with granular entitlements to enforce need-to-know in digital systems.	data vulnerability	Likely	Moderate	Low
12.1 General	> Statistical tabulations prepared at the state level may not be released for cells containing data for fewer than 10 returns. Data for geographic areas below the state level such as county may not be released with cells containing data from fewer than 20 returns. In addition for tabular data at the ZIP Code level, additional procedures must be employed. Individual ZIP Codes areas with fewer than 100 returns cannot be shown. Additionally, any cell in the ZIP Code table based on fewer than 20 returns cannot be shown. Finally, individual returns that represent a large percentage of the total of a particular cell must be excluded from the data	Requires data to be tracked at an uncommon granular level. Additionally, this provides no guarantees that the statistics protect privacy.	data vulnerability	Unlikely	Moderate	Low
2.3 Secure Data Transfer	Only the following types of documents will be accepted via SDT: Control File (.txt) Adobe (.pdf) Word Document (.doc or .docx) Excel Document (.xls or .xlsx) Zipped File (.zip) Contact the SafeguardReports@irs.gov mailbox for specific details on how to submit information via SDT.	The "types of documents [that] will be accepted via SDT" doesn't specify if/how file types are validated. It appears that it allow malicious files to be submitted if renamed to an allowed type (e.g. rename .exe to .txt). It does not specify if submitted files are virus scanned.	data vulnerability	Likely	Moderate	Low

2.3 Secure Data Transfer ACES Digital Certificates	<p>In addition to installing the SDT software, each agency must also have an IdenTrust Certificate installed. After the initial installation, agencies are required to renew the IdenTrust Certificate every two years. Refer to the ACES (Access Certificates for Electronic Services) IdenTrust website for additional information.</p>	<p>Section 2.3 references IdenTrust as the issuer of digital certificates for data transmission to IRS. I briefly visited the IdenTrust website recommended in the documentation and discovered that the system allows automated trusted certificate purchase. I reached a search page which listed different organizations to choose from when requesting a digital certificate. I did not proceed further because at that time I was presented with a warning banner: "You have accessed a U.S. Government sponsored computer system. Unauthorized use may be punished by fines or imprisonment."</p> <p>Once I encountered the warning banner I discontinued use of the website immediately.</p> <p>I believe this represents a major vulnerability for two reasons: **1.** It appears to be trivial to request and purchase a digital certificate with a root authority that the IRS trusts. As I said, I did not continue further to verification stages due to the warning banner. **2.** The search form can enumerate other agencies that are working directly with the IRS. These agencies may present softer targets for malicious actors to exploit.</p>	data vulnerability	Occasional	Critical	Low
2.7.2 Computer Security Review Process	<p>The Office of Safeguards will assess agency compliance with the computer security requirements identified in this publication as part of the on-site review process. Requirements are assessed as they relate to NIST SP 800-53 security controls as outlined in this publication. To ensure a standardized computer security review process, the following techniques will be used to evaluate agency policies, procedures, and IT equipment that receive, process, store, or transmit FTI: Automated Compliance and Vulnerability Assessment Testing</p> <p>Computer Security Reviewers will use a combination of compliance and vulnerability assessment software tools to validate the adequate protection of FTI on agency and contractor owned equipment. These automated tools will be launched from either IRS-issued flash drives or laptop computers. Profiles used with these tools can be downloaded from the Office of Safeguards' website.</p> <p>SCSEM Documents and tests hardening requirements for specific technologies used to receive, process, store, or transmit FTI. SCSEMs can and should be downloaded from the Office of Safeguards' website.</p> <p>MOT Documents will be requested in advance and expected to be provided prior to the review process, no later than the opening conference.</p> <p>Agencies should be prepared for the Computer Security Reviewers to use the preceding resources as part of the on-site review. As necessary, agency management approval must be obtained prior to the on-site review, if agency policies and procedure contradict any of these methods.</p>	<p>I'm uneasy with this section, but need help articulating what feels wrong. The "automated compliance and vulnerability assessment testing" seems to ignore testing human education. I'm also uncomfortable with "IRS-issued flash drives or laptop computers" which is oddly specific.</p>	data vulnerability	Likely	Moderate	Low
2.9.1 Termination Documentation	<p>When an agency no longer requires FTI, notify Safeguards at SafeguardReports@irs.gov by providing the following:</p> <ol style="list-style-type: none"> 1. Copies of notifications to all agencies from which FTI is received, that FTI will no longer be requested, and 2. Letter from the Head of Agency certifying that all residual FTI has been destroyed. (See Section 8.0 Disposal of FTI – IRC 6103(p)(4)(F)) <p>Once documentation is reviewed, the Office of Safeguards will send an acknowledgement of the agency's termination, instructions on Safeguard reporting and on-site review obligations. Instructions for reinstatement will be included in the acknowledgement letter.</p>	<p>There is no indication that senders are authorized or validated. It appears that a spoofed email could trigger the termination process.</p> <p>Would recommend a pre-established list of authorizing officials, and email digital signatures.</p>	data vulnerability	Seldom	Moderate	Low
3.2 Electronic and Non-Electronic FTI Logs	<p>The agency must establish a tracking system to identify and track the location of electronic and non-electronic FTI from receipt until it is destroyed. The FTI log may include tracking elements, such as:</p> <ul style="list-style-type: none"> Taxpayer Name or other identifier* Tax year(s) Type of information (e.g., revenue agent reports, Form 1040, work papers) The reason for the request Date requested Date received Exact location of the FTI Who has had access to the data If disposed of, the date and method of disposition <p>*To the extent possible, do not include FTI in the log. If FTI is used, the log must be secured in accordance with all other safeguards requirements.</p>	<p>There is a great deal of emphasis on "location" without defining that term, especially granularity. Country? State? Building? Room? Server?</p> <p>This issue may be exacerbated with the use of cloud storage where data may be stored in more than one location, and constantly being moved.</p>	verbiage issue	Seldom	Negligible	Low
3.2 Electronic and Non-Electronic FTI Logs	<p>The agency must establish a tracking system to identify and track the location of electronic and non-electronic FTI from receipt until it is destroyed. The FTI log may include tracking elements, such as:</p> <ul style="list-style-type: none"> Taxpayer Name or other identifier* Tax year(s) Type of information (e.g., revenue agent reports, Form 1040, work papers) The reason for the request Date requested Date received Exact location of the FTI Who has had access to the data If disposed of, the date and method of disposition <p>*To the extent possible, do not include FTI in the log. If FTI is used, the log must be secured in accordance with all other safeguards requirements.</p>	<p>Issue: reliance to self-reported transmissions of FTI. Mechanism cannot account for undisclosed transmission or access to FTI.</p> <p>Recommendation: Tripwire-style file monitoring for electronic systems. https://www.raymond.cc/blog/3-portable-tools-monitor-files-folders-changes/ Similar input elicitation requirements for physical access -> if FTI is accessed, generate a digital ticket that necessitates follow-up information.</p>	data vulnerability	Likely	Moderate	Low
3.3 Converted Media	<p>Conversion of FTI from paper to electronic media (scanning) or from electronic media to paper (print screens or printed reports) also requires tracking from creation to destruction of the converted FTI. All converted FTI must be tracked on logs containing the fields detailed in Section 3.2, depending upon the current form of the FTI, electronic or non-electronic.</p>	<p>Issue: unclear how to track converted instances</p> <p>Recommendation: treat each conversion of the media as a new, unique instance for tracking.</p>	insufficient process	Likely	Moderate	Low
3.4 Recordkeeping of Disclosures to State Auditors	<p>> In instances where auditors read large volumes of records containing FTI, whether in paper or electronic format, the state tax agency need only identify bulk records examined. This identification will contain the approximate number of taxpayer records, the date of inspection, a description of the records, and the name of the individual(s) making the inspection.</p>	<p>Issue: unacceptable level of control over bulk record access. Especially with digital access, there should be precise means for tracking which records were accessed.</p> <p>Recommendation: Require precise control over bulk records access.</p>	data vulnerability	Likely	Moderate	Low
4.2 Minimum Protection Standards	<p>> Security Container A security container is a storage device (e.g., turtle case, safe/vault) with a resistance to forced penetration, with a security lock with controlled access to keys or combinations.</p>	<p>Issue: no specifics on grade of locks</p> <p>Recommendation: Specify high-grade locks conducive to protecting moderate-risk data</p>	data vulnerability	Seldom	Moderate	Low
4.2 Minimum Protection Standards	<p>> Badged Employee During business hours, if authorized personnel serve as the second barrier between FTI and unauthorized individuals, the authorized personnel must wear picture identification badges or credentials. The badge must be clearly displayed and worn above the waist.</p>	<p>Issue: Badges are not a barrier. Enforcement is the barrier.</p> <p>Recommendation: Require card readers or guards that check badge validity.</p> <p>> A security guard, custodial services worker, or landlord may have access to a locked building or a locked room if FTI is in a locked security container.</p> <p>Issue: This makes a single barrier.</p> <p>Recommendation: Require an escort or CCTV monitoring</p>	data vulnerability	Likely	Moderate	Low

4.2 Minimum Protection Standards - Fire Safety	The Perimeter is enclosed by slab-to-slab walls constructed of durable materials and supplemented by periodic inspection. Any lesser-type partition must be supplemented by electronic intrusion detection and fire detection systems	This indicates to me that fire detection systems are optional **IF** you are using slab-to-slab perimeter walls. Is employee safety not a concern? I believe employees working in facilities with slab-to-slab walls should still be afforded fire detection systems...	verbiage issue	Seldom	Negligible	Low
4.3.1 Use of Authorized Access List	To facilitate the entry of employees who have a frequent and continuing need to enter a restricted area, but who are not assigned to the area, an Authorized Access List (AAL) can be maintained so long as MPS are enforced (see Section 4.2, Minimum Protection Standards). Agency Employees: The AAL must contain the following: Name of individual Agency or department name Name and phone number of agency POC Address of agency POC Purpose for access The AAL for agency employees must be updated at least annually or when employee access changes. Vendors and Non-Agency Personnel: The AAL must contain the following information: Name of vendor/contractor/non-agency personnel Name and phone number of agency Point of Contact authorizing access Name and address of vendor POC Address of vendor/contractor Purpose and level of access Vendors, contractors, and non-agency personnel AAL must be updated monthly. If there is any doubt of the identity of the individual, the security monitor must verify the identity of the vendor/contractor individual against the AAL prior to allowing entry into the restricted area.	Issue: No mention of validating AAL. Recommendation: Monthly auditing AAL for need-to-access. Revoke access before transfer or termination.	insufficient process	Occasional	Moderate	Low
4.4 FTI in Transit	Handling FTI must be such that the FTI does not become misplaced or available to unauthorized personnel. Any time FTI is transported from one location to another, care must be taken to provide appropriate safeguards. When FTI is hand-carried by an individual in connection with a trip or in the course of daily activities, it must be kept with that individual and protected from unauthorized disclosures.	Issue: This section does call out electronic FTI and protects it with documentation and double-sealed envelopes. Recommended: All electronic records in transit must be encrypted.	data vulnerability	Likely	Moderate	Low
4.5 Physical Security of Computers, Electronic, and Removable Media	Computers and electronic media that receive, process, store, or transmit FTI must be in a secure area with restricted access. In situations when requirements of a secure area with restricted access cannot be maintained, such as home work sites, remote terminals or other office work sites, the equipment must receive the highest level of protection practical, including full disk encryption. All computers and mobile devices that contain FTI and reside at an alternate work site must employ encryption mechanisms to ensure that FTI may not be accessed if the computer is lost or stolen	First paragraph requires full disk encryption for systems that process, store or transmit FTI, but do nothing to address the common pitfalls of full disk encryption including: - Weak users passwords that render full disk encryption practically worthless - Should provide users with training to lock their machines when stepping away to encrypt disk. Files are unencrypted for valid active logon session.	data vulnerability	Likely	Moderate	Low
4.5 Physical Security of Computers, Electronic, and Removable Media	> All computers and mobile devices that contain FTI and reside at an alternate work site must employ encryption mechanisms to ensure that FTI may not be accessed if the computer is lost or stolen	Mobile devices that store FTI should also be configured with a remote wipe mechanism. Ideally a mobile device management software suite like Maas 360 should be used, allowing IT to remotely disable access if the employee is fired, or if the device is lost. Device should be able to be remotely locked and wiped.	insufficient process	Likely	Moderate	Low
4.7 Telework Locations	"The agency must conduct periodic inspections of alternative work sites during the year to ensure that safeguards are adequate."	This seems to include employee homes. I can't think of all the possible concerns here, but it seems fraught with possible issues.	unenforceable	Frequent	Moderate	Low
4.7.3 Other Safeguards	> The agency must provide specialized training in security, disclosure awareness, and ethics for all participating employees and managers. This training must cover situations that could occur as the result of an interruption of work by family, friends, or other sources.	Issue: does not account for vendors / contractors with regular access to FTI. > 6.3 Disclosure Awareness Training. Employees and contractors must maintain their authorization to access FTI through annual training and recertification. Prior to granting an agency employee or contractor access to FTI, each employee or contractor must certify his or her understanding of the agency's security policy and procedures for safeguarding IRS information. Recommendation: extend verbiage to consistently include anyone that may access FTI	verbiage issue	Seldom	Negligible	Low
5.3 Access to FTI via State Tax Files or Through Other Agencies	> FTI cannot be accessed by agency employees, agents, representatives, or contractors located offshore/outside of the United States territories, embassies or military installations. Further, FTI may not be received, processed, stored, transmitted, or disposed of by information technology (IT) systems located offshore.	The Internet makes this essentially unenforceable. Recommend at least rewording to prefer data be kept in the United States, but requiring strong data protections at rest and in transit.	unenforceable	Frequent	Moderate	Low
5.4.2.2 Consolidated Data Centers	Agencies using consolidated data centers must implement appropriate controls to ensure the protection of FTI, including a service level agreement (SLA) between the agency authorized to receive FTI and the consolidated data center. The SLA must cover the following: The agency with authority to receive FTI is responsible for ensuring the protection of all FTI received. The consolidated data center shares responsibility for safeguarding FTI. The SLA provides written notification to the consolidated data center management that they are bound by the provisions of Publication 1075, relative to protecting all FTI within their possession or control. The SLA shall detail the IRS' right to inspect consolidated data center facilities and operations accessing, receiving, storing or processing FTI under this agreement to assess compliance with requirements defined in IRS Publication 1075. The SLA shall specify that IRS' right of inspection includes the use of manual and/or automated scanning tools to perform compliance and vulnerability assessments of information technology (IT) assets that access, store, process or transmit FTI. The SLA shall detail the consolidated data center's responsibilities to address corrective action recommendations to resolve findings of noncompliance identified by IRS inspections. The agency will conduct an internal inspection of the consolidated data center every 18 months, as described in Section 6.4, Internal Inspections. Multiple agencies sharing a consolidated data center may partner together to conduct a single, comprehensive internal inspection. However, care must be taken to ensure agency representatives do not gain unauthorized access to other agencies' FTI during the internal inspection.	SLA should add a provision to physically separate or distinctly label all devices that process FTI. Racks should be locked or in a cage and only accessible by people authorized to access FTI.	data vulnerability	Occasional	Moderate	Low
6.1 General	> IRC 6103(p)(4)(D) requires that agencies receiving FTI to provide other safeguard measures, as appropriate, to ensure the confidentiality of the FTI.	I can't figure out why this is limited to confidentiality and doesn't include integrity. Integrity does show up in Section 9 (Computer System Security). Recommend changing to "...ensure the confidentiality and integrity of the FTI."	verbiage issue	Seldom	Moderate	Low
6.4.2 Secure Storage	> FTI (including tapes, cartridges, or other removable media) must be stored in a secure location, safe from unauthorized access.	This section is simply redundant. It should reference other areas within the document that provide specific guidance otherwise it only clarifies that "Secure storage must be secure"	data vulnerability	Likely	Moderate	Low

6.5 Plan of Action and Milestones	The agency must implement a process for ensuring that a Plan of Action and Milestones (POA&M) is developed and monitored. The POA&M must include the corrective actions identified during the internal inspections and will identify the actions the agency plans to take to resolve these findings.	Issue: no mention of sanctions or action against non-compliant agencies. Recommendation: revoke access to FTI is non-compliant EDIT: Exhibit 3 USC Title 26, CFR 301.6103(p)(7)-1 says IRS can terminate or suspend access to FTI. There is no enforcement mechanism for making them purge existing FTI.	data vulnerability	Likely	Moderate	Low
7.1 General	"FTI to report on procedures established and used for ensuring the confidentiality of FTI that is received"	Recommendations: Introduce language to protect both confidentiality and integrity	verbiage issue	Seldom	Moderate	Low
7.1.2 Encryption Requirements	>Communicate the password or pass phrase with the Office of Safeguards through a separate email or via a telephone call to your IRS contact person. Do not provide the password or passphrase in the same email containing the encrypted attachment	Passwords for documents should not be transmitted over the same communications mechanism, even if they are sent in a separate correspondence. Recommend: Send passwords to attachments in secondary communications mechanism. i.e. Phone call or SMS. Better yet, IRS should publish a Public PGP key and all attachment should be encrypted using IRS's public key, eliminating the need for password transmittal entirely.	data vulnerability	Likely	Moderate	Low
7.1.2 Encryption Requirements	> The Office of Safeguards recommends that all required reports, when sent to the Office of Safeguards via email, be transmitted using IRS-approved encryption methods to protect sensitive information.	This should not be optional. Recommend changing "recommends" to "requires".	insufficient process	Likely	Moderate	Low
7.2.3 Annual SSR Update Submission Instructions	_" The agency must updated and submit the SSR **annually** to encompass any changes that impact the protection of FTI. Example changes include, but are not limited to:_" - New Computer equipment, systems or applications (hardware or software)	This says to me that an organization can have a security assessment and then completely change all of their hardware, networking, and software stack and still be good without a security review for up to another year. I would argue that any changes to hardware, networking, software outside of the evaluated base can introduce serious vulnerabilities and should be performed using strict change management with a well defined threshold that would trigger the requirement of a new security assessment.	insufficient process	Likely	Moderate	Low
7.3.1 CAP Submission Instructions and Submission Dates	When extenuating circumstances exist, agencies may request an extension for no more than 30 days. Extension requests should be submitted not later than (NLT) 30 days prior to the scheduled CAP due date. Request for extensions will not be considered after the scheduled CAP due date. Extension requests should be sent to the Office of Safeguards via Secure Data Transfer (SDT) or email to SafeguardReports@irs.gov, with the subject CAP extension request and reasons for the request. All extension requests will be evaluated on a case by case basis. Safeguards will provide an email response, approving or disapproving the request within 5 work days after receipt of the request.	No sanctions or consequences for noncompliance EDIT: Exhibit 3 USC Title 26, CFR 301.6103(p)(7)-1 says IRS can terminate or suspend access to FTI. There is no enforcement mechanism for making them purge existing FTI.	insufficient process	Occasional	Moderate	Low
7.4.1 Cloud Computing	Receiving, processing, storing, or transmitting FTI in a cloud environment requires prior notification to the Office of Safeguards. Refer to Section 9.4.1, Cloud Computing Environments, for guidance and details on 45-day notification requirements	Although this is the reporting sections, this may be the place to address cloud computing requirements i.e. That all FTI stored on IaaS must be in a US region.	insufficient process	Likely	Moderate	Low
7.4.8 Virtualization of Information Technology Systems	> No prior notification is required when an agency is planning to receive, process, store, or transmit FTI in virtualized environments.	Most cloud environments use virtualization. This section could therefore contradict 7.4.1 which states that "Receiving, processing, storing, or transmitting FTI in a cloud environment requires prior notification..." Recommend removing this section. (Maybe someone can figure out how to narrowly articulate non-shared non-cloud virtualization)	verbiage issue	Unlikely	Negligible	Low
8.2 Returning IRS Information to the Source	>Agencies electing to return IRS information must use a receipt process and ensure that the confidentiality is protected at all times during transport (see Section 4.4, FTI in Transit).	Recommendations: Add confidentiality and integrity	verbiage issue	Likely	Moderate	Low
8.3 Destruction and Disposal	> When using either method for destruction, every third piece of physical electronic media must be checked to ensure appropriate destruction of FTI.	This doesn't make any sense to me, and sounds excessive. Why every third piece of media? Recommend changing to a percentage of the total amount of data storage being destroyed (e.g. check 1GB for every 100GB being destroyed)	insufficient process	Likely	Moderate	Low
9.3.1.1 Access Control Policy and Procedures	_"Review and update the current access control policy every three years (or if there is a significant change)_"	There is no formal definition of a threshold which indicates a "significant change." This is vague and if left up to the subjective view of other agencies, this will be abused.	verbiage issue	Likely	Moderate	Low
9.3.1.1.2 Remote Access	_"Authorize and **document** the execution of privileged commands and access to security-relevant information via remote access for compelling operational needs only"	Documenting that an action occurred does not lead to discovery of compromise unless the logs are reviewed frequently. The document lays out many requirements like this to "document" accesses, but does not formally define requirements for reviewing these logs. What good is documenting an event without a system/process formally established to review events?	insufficient process	Likely	Moderate	Low
9.3.1.1.4 Access Control for Mobile Devices (AC-19)	The agency must: a. Establish usage restrictions, configuration requirements, connection requirements, and implementation guidance for agency-controlled mobile devices b. Authorize the connection of mobile devices to agency information systems c. Employ encryption to protect the confidentiality and integrity of information on mobile devices (e.g., smartphones and laptop computers) (CE5) b. Purge/wipe information from mobile devices based on 10 consecutive, unsuccessful device logon attempts (e.g., personal digital assistants, smartphones and tablets). Laptop computers are excluded from this requirement (AC-7, CE2)	Does not specify logon requirements for the device (e.g. pin, fingerprint, facial recog, password, etc) Additionally, this section is inconsistent with 9.3.1.1.2 Remote Access (AC-17) without mention of MFA when used off prim.	data vulnerability	Likely	Moderate	Low
9.3.1.1.5 Use of External Information Systems (AC-20)	> The agency may allow the use of personally-owned devices, without notification, only > for the following purposes: > a. Bring Your Own Device (BYOD) used to access e-mail, where all requirements > in Section 9.4.8 Mobile Devices are met	In other sections they stated that FTI can be transferred over email, so this waiver would allow FTI to be viewed on a personal device.	data vulnerability	Likely	Moderate	Low
9.3.1.1.7 Publicly Accessible Content (AC-22)	The agency must: a. Designate individuals authorized to post information onto a publicly accessible information system b. Train authorized individuals to ensure that publicly accessible information does not contain FTI c. Review the proposed content of information prior to posting onto the publicly accessible information system to ensure that FTI is not included d. Review the content on the publicly accessible information system for FTI, at a minimum, quarterly and remove such information, if discovered	Does not specify approval authorities for review or how review should be conducted. Is the janitor a sufficient reviewer?	insufficient process	Occasional	Moderate	Low

9.3.1.2 Account Management (AC-2)	> Notify account managers when accounts are no longer required, when users are terminated or transferred, or when individual information system usage or need-to-know permission changes > Review accounts for compliance with account management requirements at a minimum of annually for user accounts and semi-annually for privileged accounts	Notifying is wholly insufficient. This should trigger a mandatory revocation of access. Otherwise: They could maintain access for damn-near a year. _Note: these are taken from 9.3 NIST SP 800-53 Control Requirements_	insufficient process	Occasional	Moderate	Low
9.3.1.2 Account Management (AC-2) - shared accounts	> Establish a process for reissuing shared/group account credentials (if > deployed) when individuals are removed from the group.	Nope, just nope. While I agree there should be a process for removing accounts from a group. There should be **NO** shared account access to FTI, this defeats any auditing actions in place.	data vulnerability	Frequent	Moderate	Low
9.3.1.2 Account Management: User Monitoring	"Monitor the use of information system accounts"...	Monitor is not well defined here. Are we talking about email monitoring? Keystroke logging? Event logs? PCAPS/ Network Monitoring?	insufficient process	Likely	Moderate	Low
9.3.1.9 Session Lock (AC-11)	> Prevent further access to the system by initiating a session lock after 15 minutes of inactivity or upon receiving a request from a user	This is pretty secure, but also pretty unusable. Not really a vulnerability, but could be one if users try to develop mechanisms to circumvent this protection. (Remind me to tell you about our novel use of a mouse and a clock sometime...)	insufficient process	Likely	Moderate	Low
9.3.11.5 Access Control for Output Devices (PE-5)	> The agency must control physical access to information system output devices to prevent unauthorized individuals from obtaining the output. Monitors, printers, copiers, scanners, fax machines, and audio devices are examples of information system output devices.	What about network taps? Rogue access points?	data vulnerability	Seldom	Moderate	Low
9.3.11.6 Monitoring Physical Access (PE-6)	> Review physical access logs annually	A yearly review of physical access logs is ineffective.	insufficient process	Likely	Negligible	Low
9.3.13.3 Personnel Screening (PS-3)	The agency must: a. Screen individuals prior to authorizing access to the information system b. Rescreen individuals according to agency-defined conditions requiring rescreening	> Screen individuals prior to authorizing access to the information system Screen how? Federal database? Local?	insufficient process	Frequent	Critical	Low
9.3.13.4 Termination (PS-4) [and 9.3.13.5 Personnel Transfer (PS-5)]	The agency, upon termination of individual employment must: a. Disable information system access b. Terminate/ revoke any authenticators/credentials associated with the individual c. Conduct exit interviews, as needed d. Retrieve all security-related agency information system-related property e. Retain access to agency information and information systems formerly controlled by the terminated individual f. Notify agency personnel upon termination of the employee	How is this verified? Where is the check and balance? Additionally, there is no time period associated with revocation or transfer. Recommendation: Revoke access immediately before notifying employee of termination to prevent access to sensitive data.	data vulnerability	Likely	Critical	Low
9.3.14.3 Vulnerability Scanning (RA-5)	Remediate legitimate vulnerabilities in accordance with an assessment of risk	No fixed requirement for changing / patching / fixing ID'd vulns (d.) says "Remediate legitimate vulnerabilities in accordance with an assessment of risk" but whose assessment? Without a timeline suspense, this may never get fixed.	data vulnerability	Seldom	Moderate	Low
9.3.15.4 Acquisition Process (SA-4)	The agency must include the following requirements, descriptions, and criteria, explicitly or by reference, in the acquisition contract for the information system, system component, or information system service in accordance with applicable federal laws, Executive Orders, directives, policies, regulations, standards, guidelines, and agency mission/business needs: a. Security functional requirements b. Security strength requirements c. Security assurance requirements e. Security-related documentation requirements f. Requirements for protecting security-related documentation g. Description of the information system development environment and environment in which the system is intended to operate h. Acceptance criteria When applicable, the agency must require the developer of the information system, system component, or information system service to provide a description of the functional properties of the security controls to be employed (CE1)	No mention of supply chain security or authorized purchase locations. Example: [https://arstechnica.com/information-technology/2016/11/chinese-company-installed-secret-backdoor-on-hundreds-of-thousands-of-phones/]([https://arstechnica.com/information-technology/2016/11/chinese-company-installed-secret-backdoor-on-hundreds-of-thousands-of-phones/])	data vulnerability	Likely	Critical	Low
9.3.15.8 Developer Configuration Management (SA-10)	> Track security flaws and flaw resolution within the system, component, or service > and report findings to designated agency officials	I don't know what this means, but developers shouldn't note security vulnerabilities in the systems the same system that's vulnerable.	insufficient process	Likely	Negligible	Low
9.3.16.5 Boundary Protection	> The agency must limit the number of external network connections to the information system. (CE3)	Recommend that a maximum number or other quantifiable limit be specified.	unenforceable	Likely	Negligible	Low
9.3.17.3 Malicious Code Protection (SI-3)	Malicious code protection includes antivirus software and antimalware and intrusion detection systems. The agency must: a. Employ malicious code protection mechanisms at information system entry and exit points to detect and eradicate malicious code b. Update malicious code protection mechanisms whenever new releases are available in accordance	What happens when the AV is the point of infection? Nowhere is there a forcing function for vulns to be fixed in a timely manner. Recommendation: Whitelist applications, prevent chained installations of child applications from whitelisted apps (e.g., Chrome cannot install anything nor can Norton AV)	data vulnerability	Occasional	Critical	Low
9.3.17.6 Spam Protection (SI-8)	> Malicious code protection includes antivirus software and antimalware and intrusion detection systems.	I don't think that _Spam_ is the issue here. They should be more concerned with phishing attacks.	verbiage issue	Seldom	Negligible	Low
9.3.17.7 Information Input Validation (SI-10)	The information system must check the validity of information inputs.	Incredibly vague. How is this check performed? Fuzzing? Manual audit? Additionally, if org is using a third-party application and a vuln is found through pentesting or fuzzing, what then?	verbiage issue	Likely	Negligible	Low
9.3.2.3 Role-Based Security Training (AT-3)	> Note: Training conducted under this section is distinct from Section 6.3, Disclosure Awareness, and Section 9.3.2.2, Security Awareness Training (AT-2).	Just to pick a nit. This role based training should encompass the concepts in 9.3.2.2. **DO NOT** make this role based training an addendum. It will be seen as punitive and reduces the effectiveness, since employees will try everything to reduce their mandatory training burden. (<- I know from experience.)	insufficient process	Seldom	Negligible	Low
9.3.3.2 Audit Events (AU-2)	> a. Determine that the information system is capable, at a minimum, of auditing the following event types: > 4. Changes made to an application or database by a batch file	I find the list unsatisfying. What does this mean, and why batch files? What about printing? Large network transfers? Network connections in general (ie netflow)? At a minimum, I think the list needs to cover more general areas (user behavior, system behavior, data access and modification, data transfer, etc)	data vulnerability	Occasional	Moderate	Low
9.3.3.5 Response to Audit Processing Failures (AU-5)	> Provide a warning when allocated audit record storage volume reaches a > maximum audit record storage capacity (CE1)	Should provide a warning **BEFORE** reaching audit record storage capacity.	insufficient process	Likely	Moderate	Low

9.3.3.6 Audit Review, Analysis, and Reporting (AU-6)	> a. Review and analyze information system audit records at least weekly or more frequently at the discretion of the information system owner for indications of unusual activity related to potential unauthorized FTI access	This isn't horrible, but I'd prefer they do continual monitoring and identify indications of unusual activity more frequently than weekly.	insufficient process	Unlikely	Moderate	Low
9.3.4.2 Security Assessments (CA-2)	The agency must: a. Develop a security assessment plan that describes the scope of the assessment, including: 1. Security controls and control enhancements under assessment 2. Assessment procedures to be used to determine security control effectiveness 3. Assessment environment, assessment team, and assessment roles and responsibilities b. Assess the security controls in the information system and its environment at a minimum on an annual basis to determine the extent to which the controls are implemented correctly, operating as intended, and producing the desired outcome with respect to meeting established security requirements	Does not specify vuln assessment, pentest, or expectations of scope for annual assessment. Wholly insufficient.	insufficient process	Seldom	Moderate	Low
9.3.4.3 System Interconnections (CA-3)	> a. Authorize connections from the information system to other information systems through the use of interconnection Security Agreements > Document, for each interconnection, the interface characteristics, security requirements, and the nature of the information communicated	This should also require a validation/audit requirement. These connection documents get outdated very fast and details are often missed. Need a mechanism to test whether systems have additional interconnects not stipulated in the documentation (or that currently documented interconnects still exist).	insufficient process	Occasional	Moderate	Low
9.3.4.4 Plan of Action and Milestones (CA-5)	a. Develop a POA&M for the information system to document the agency's planned remedial actions to correct weaknesses or deficiencies noted during the assessment of the security controls and to reduce or eliminate known vulnerabilities in the system b. Update the existing POA&M on a quarterly basis, at a minimum, based on the findings from security controls assessments, security impact analyses, and continuous monitoring activities	No sanctions or consequence for non-compliance. EDIT: 9.3.13.8 Personnel Sanctions (PS-8) talks about personnel sanctions, not organization sanctions EDIT: Exhibit 3 USC Title 26, CFR 301.6103(p)(7)-1 says IRS can terminate or suspend access to FTI. There is no enforcement mechanism for making them purge existing FTI.	data vulnerability	Likely	Moderate	Low
9.3.4.5 Security Authorization (CA-6)	> Assign a senior-level executive or manager as the authorizing official for the information system	No consequences outlined for this individual.	insufficient process	Likely	Negligible	Low
9.3.5.10 Software Usage Restrictions (CM-10)	>Control and document the use of peer-to-peer file sharing technology to ensure that this capability is not used for the unauthorized distribution, display, performance, or reproduction of copyrighted work	Should update to include cloud sharing services (Google Drive, Dropbox, etc.)	data vulnerability	Likely	Moderate	Low
9.3.5.11 User-Installed Software (CM-11)	The agency must: a. Establish policies governing the installation of software by users b. Enforce software installation policies through automated methods c. Monitor policy compliance on a continual basis	Don't allow users to install software. Or at least institute application whitelisting.	data vulnerability	Likely	Critical	Low
9.3.5.8 Information System Component Inventory (CM-8)	The agency must: a. Develop and document an inventory of information system components that: 1. Accurately reflects the current information system 2. Includes all components that store, process, or transmit FTI 3. Is at the level of granularity deemed necessary for tracking and reporting 4. Includes information deemed necessary to achieve effective information system component accountability b. Review and update the information system component inventory through periodic manual inventory checks or a network monitoring tool that automatically maintains the inventory c. Update the inventory of information system components as an integral part of component installations, removals, and information system updates (CE1)	No temporal requirement for maintaining updated inventories. This lends itself to rogue systems being on the network for an enduring amount of time. Furthermore, the "or" statement of manual inspection or automated assessments need to be tied to a ground truth --> asset inventory through supply systems ("do we own the devices on our network?"). Google SRE talks about this being a core tenet of their security. Failure to ID ground truth could permit rogue computers to be whitelisted and treated as legit systems.	data vulnerability	Occasional	Moderate	Low
9.3.6.8 Information System Recovery and Reconstitution (CP-10)	The agency must provide for the recovery and reconstitution of the information system to a known state after a disruption, compromise, or failure.	Reverting an unpatched system to a previous unpatched state helps nothing.	data vulnerability	Likely	Moderate	Low
9.3.7.2 Identification and Authentication (Organizational Users) (IA-2)	The information system must: a. Uniquely identify and authenticate agency users (or processes acting on behalf of agency users) b. Implement multi-factor authentication for all remote network access to privileged and non-privileged accounts for information systems that receive, process, store, or transmit FTI. (CE1, CE2) c. Implement multi-factor authentication for remote access to privileged and non-privileged accounts such that one of the factors is provided by a device separate from the system gaining access. NIST SP 800-63 allows the use of software tokens. (CE11)	Why only MFA for a subset of systems? Why not MFA for any system that touches FTI?	data vulnerability	Likely	Moderate	Low
9.3.7.5 Authenticator Management (IA-5)	a. Verifying, as part of the initial authenticator distribution, the identity of the individual, group, role, or device receiving the authenticator b. Establishing initial authenticator content for authenticators defined by the agency c. Ensuring that authenticators have sufficient strength of mechanism for their intended use d. Establishing and implementing administrative procedures for initial authenticator distribution, for lost/compromised or damaged authenticators, and for revoking authenticators e. Changing default content of authenticators prior to information system installation f. Establishing minimum and maximum lifetime restrictions and reuse conditions for authenticators g. Changing/refreshing authenticators h. Protecting authenticator content from unauthorized disclosure and modification i. Requiring individuals to take, and having devices implement, specific security safeguards to protect authenticators j. Changing authenticators for group/role accounts when membership to those accounts changes The information system must, for password-based authentication: a. Enforce minimum password complexity of: 1. Eight characters 2. At least one numeric and at least one special character 3. A mixture of at least one uppercase and at least one lowercase letter 4. Storing and transmitting only encrypted representations of passwords b. Enforce password minimum lifetime restriction of one day c. Enforce non-privileged account passwords to be changed at least every 90 days d. Enforce privileged account passwords to be changed at least every 60 days e. Prohibit password reuse for 24 generations	Shitty password change policies that NIST no longer recommends.	data vulnerability	Frequent	Moderate	Low
9.3.8.3 Incident Response Testing (IR-3)	Agencies entrusted with FTI must test the incident response capability at least annually. a. Agencies must perform tabletop exercises using scenarios that include a breach of FTI and should test the agency's incident response policies and procedures. b. A subset of all employees and contractors with access to FTI must be included in tabletop exercises. c. Each tabletop exercise must produce an after-action report to improve existing processes, procedures, and policies.	A tabletop exercise is the only requirement? This goes back to [security assessments](https://github.com/rstevens70/hackingcompliance/issues/44) being horribly defined. You evaluate SOPs and policies through live training against a pentest team.	insufficient process	Likely	Moderate	Low
9.3.8.4 Incident Handling (IR-4)	> Implement an incident handling capability for security incidents that includes preparation, detection and analysis, containment, eradication, and recovery	Does not cover evidence preservation/forensics.	insufficient process	Likely	Negligible	Low

9.3.8.9 Information Spillage Response (IR-9)	The agency must respond to information spills by: a. Identifying the specific information involved in the information system contamination b. Alerting authorized incident response personnel of the information spill using a method of communication not associated with the spill c. Isolating the contaminated information system or system component d. Eradicating the information from the contaminated information system or component e. Identifying other information	This section doesn't say anything about documenting information spillage or providing an after-action step to learn from mistakes.	insufficient process	Occasional	Moderate	Low
9.4.1 Cloud Computing Environments	> Data is not stored in an agency-managed data center	This is true only of public clouds. Private clouds may reside in an agency-managed data center.	verbiage issue	Occasional	Negligible	Low
9.4.18 Wireless Networks	Requirements To use FTI in an 802.11 WLAN, the agency must meet the following mandatory requirements: a. The agency should have WLAN management controls that include security policies and procedures, a complete inventory of all wireless network components, and standardized security configurations for all components. b. WLAN hardware (access points, servers, routers, switches, firewalls) must be physically protected in accordance with the minimum protection standards for physical security outlined in Section 4.0, Secure Storage—IRC 6103(p)(4)(B). c. Each system within the agency's network that transmits FTI through the WLAN is hardened in accordance with the requirements in this publication. d. The WLAN is architected to provide logical separation between WLANs with different security profiles and from the wired LAN. e. WLAN infrastructure that receives, processes, stores, or transmits FTI must comply with the Institute of Electrical and Electronic Engineers 802.11i wireless security standard and perform mutual authentication for all access to FTI via 802.1X and extensible authentication protocol f. Vulnerability scanning should be conducted as part of periodic technical security assessments for the organization's WLAN. g. Wireless intrusion detection is deployed to monitor for unauthorized access, and security event logging is enabled on WLAN components in accordance with Section 9.3.3, Audit and Accountability. h. Disposal of all WLAN hardware follows media sanitization and disposal procedures in Section 9.3.10.6, Media Sanitization (MP-6), and Section 9.4.7, Media Sanitization.	No mention of rogue access points. No mention of logical segmentation for each user session (logical switch vs hub).	data vulnerability	Likely	Critical	Low
9.4.8 Mobile Devices	> Access to hardware, such as the digital camera, global positioning system > (GPS), and universal serial bus (USB) interface, must be disabled to the extent > possible	I disagree with the GPS portion of this statement. Using location data to geo-fence, located a device, or perform and out of bounds remote wipe are all strategies that should be used for mobile.	insufficient process	Likely	Negligible	Low
Table 5 - Evidentiary Requirements for SSR approval before release of FTI	AC-17, Remote Access Screenshot of authentication screens Document how multi-factor authentication is deployed for all remote network access to systems containing FTI and the tokens used for authentication Section 5.2, Comingling and Labeling Screenshots of database schemas that show electronic FTI labeling Sample output (report/notice) that shows how FTI is labeled	The use of screenshots to prove evidentiary requirements feels squishy to me. These could be easily forged/photoshopped. They are not strong evidence of safeguards.	data vulnerability	Likely	Moderate	Low

Compliance Audit Results for PCI DSS						
Section	Text of concern	Concern	Issue	Probability	Severity	Impact
Network Segmentation	Network segmentation of, or isolating (segmenting), the cardholder data environment from the remainder of an entity's network is not a PCI DSS requirement. However, it is strongly recommended as a method that may reduce the scope of the PCI DSS assessment	Excluding segments of network outside of CDE from protections puts CDE at risk of VLAN hopping. Could allow attackers to use less secure systems to pivot into CDE	data vulnerability	Probability Seldom	Severity Critical	Low
1.1.1.a	Examine documented procedures to verify there is a formal process for testing and approval of all:	Does not specify how to safely perform change control... Should recommend a review process or something similar. It feels strange to allow someone to makeup their own processes. We don't let them roll their own crypto, so why let them roll their own procedures for collecting sensitive data?	insufficient process	Probability Frequent	Severity Moderate	Low
1.1.2.b	Interview responsible personnel to verify that the diagram is kept current.	What does kept current mean? how often? right before the inspection or more regularly?	insufficient process	Probability Frequent	Severity Negligible	Low
1.3.7.a	> Examine firewall and router configurations to verify that methods are in place to prevent the disclosure of private IP addresses and routing information from internal networks to the Internet	What about other public records like DNS that would allow enumeration?	insufficient process	Probability Frequent	Severity Negligible	Low
1.4	Install personal firewall software or equivalent functionality on any portable computing devices (including company and/or employee-owned) that connect to the Internet when outside the network	Do not allow personal systems within CDE. Firewalls can't stop all threats.	data vulnerability	Probability Likely	Severity Critical	Low
2.1.a	Choose a sample of system components, and attempt to log on (with system administrator help) to the devices and applications using default vendor-supplied accounts and passwords, to verify that ALL default passwords have been changed.	Automate a full scan. 1 default account is all it takes.	insufficient process	Probability Seldom	Severity Moderate	Low
3.2.1	For a sample of system components, examine data sources including but not limited to the following, and verify that the full contents of any track from the magnetic stripe on the back of card or equivalent data on a chip are not stored after authorization.	Do not sample. Develop regex and automate.	insufficient process	Probability Frequent	Severity Moderate	Low
3.5.1	Interview responsible personnel and review documentation to verify that a document exists to describe the cryptographic architecture, including:	Do not allow home-rolled crypto. Mandate a minimum strength for approved standards	insufficient process	Probability Occasional	Severity Moderate	Low
4.1.1	> Weak encryption (for example, WEP, SSL) is not used as a security control for authentication or transmission	Short or easily guessable WPA2 PSK within an unsegmented LAN provides is a danger as well	insufficient process	Probability Occasional	Severity Moderate	Low
5.2.a	Examine policies and procedures to verify that anti-virus software and definitions are required to be kept up to date.	Define periodic, 24 hours and auto install?	insufficient process	Probability Frequent	Severity Moderate	Low
Requirement 5: Protect all systems against malware and regularly update anti-virus software or programs	Requirement 5: Protect all systems against malware and regularly update anti-virus software or programs	This entire section relies solely on antivirus to prevent malware infections. This is wholly insufficient and should mandate more, such as: application whitelisting, block installations in areas that permit persistence, etc.	insufficient process	Probability Frequent	Severity Moderate	Low
PCI DSS Applicability Information	Cardholder data and sensitive authentication data are defined as follows:	- Email address - Password - Social Security Number - Data of Birth	data vulnerability	Probability Likely	Severity Moderate	Low
3.6.1.b	Observe the procedures for generating keys to verify that strong keys are generated.	"Observation" is insufficient for verifying key strength.	verbiage	Probability Seldom	Severity Moderate	Low
4.2	Never send unprotected PANs by end-user messaging technologies (for example, e-mail, instant messaging, SMS, chat, etc.)	Should similarly apply to passwords.	insufficient process	Probability Likely	Severity Moderate	Low
5.1	Deploy anti-virus software on all systems commonly affected by malicious software (particularly personal computers and servers)	I don't know what to recommend, but "commonly affected" seems like a bad specification.	data vulnerability	Probability Likely	Severity Moderate	Low
Requirement 7: Restrict access to cardholder data by business need to know	Requirement 7: Restrict access to cardholder data by business need to know	This section is missing requirements for re-review of access, and revocation of access.	insufficient process	Probability Frequent	Severity Moderate	Low
11.1.c and 11.1.d	> 11.1.c If wireless scanning is utilized, examine output from recent wireless scans to verify that: > 11.1.d If automated monitoring is utilized (for example, wireless IDS/IPS, NAC, etc.)	Why have additional provisions for more secure implementations? This might de-incentivize security and encourage weaker implementations to bypass more work/auditing/review.	insufficient process	Probability Seldom	Severity Negligible	Low
11.2.1.b	Review the scan reports and verify that all high risk vulnerabilities are addressed and the scan process includes rescans to verify that the high risk vulnerabilities (as defined in PCI DSS Requirement 6.1) are resolved.	As written, high-risk vulns can go unaddressed for 11 months as long as they're addressed before an inspection.	insufficient process	Probability Occasional	Severity Critical	Low
11.3.3	Examine penetration testing results to verify that noted exploitable vulnerabilities were corrected and that repeated testing confirmed the vulnerability was corrected.	How soon after ID?	verbiage	Probability Likely	Severity Moderate	Low
A1.1	All CGI scripts used by an entity must be created and run as the entity's unique user ID.	CGI scripts is too narrow, should be "applications"	verbiage	Probability Frequent	Severity Moderate	Low
PCI DSS Applicability Information	> The primary account number is the defining factor for cardholder data. If cardholder name, service code, and/or expiration date are stored, processed or transmitted with the PAN, or are otherwise present in the cardholder data environment (CDE), they must be protected in accordance with applicable PCI DSS requirements.	Why not protect PII when PAN is not present?	insufficient process	Probability Frequent	Severity Negligible	Low
1.1.3	> Examine data-flow diagram and interview personnel to verify the diagram: Is kept current and updated as needed upon changes to the environment.	What does kept current mean? how often? right before the inspection or more regularly?	insufficient process	Probability Frequent	Severity Negligible	Low
1.1.5	> Description of groups, roles, and responsibilities for management of network components	what are the minimum set of requirements for responsibilities?	insufficient process	Probability Frequent	Severity Negligible	Low
1.1.6.b	> Identify insecure services, protocols, and ports allowed; and verify that security features are documented for each service.	why are insecure services allowed? who is approval authority?	data vulnerability	Probability Likely	Severity Moderate	Low
2.1.1.a	> Interview responsible personnel and examine supporting documentation to verify that:	Why is this different than previous section? What not verify technical implementation that creds are changed instead of interviewing?	insufficient process	Probability Seldom	Severity Moderate	Low
2.2	> Assure that these standards address all known security vulnerabilities and are consistent with industry-accepted system hardening standards.	Use current versions.	verbiage	Probability Seldom	Severity Negligible	Low
2.2.3	> Implement additional security features for any required services, protocols, or daemons that are considered to be insecure.	This relies on homebrewed wrappings to secure insecure protocols. Just don't use them if insecure. See #8	data vulnerability	Probability Likely	Severity Moderate	Low
2.4.b	> Interview personnel to verify the documented inventory is kept current.	Again, what does current mean? Action-based updating? Updated every X days? How are updated inventories verified? See #5	insufficient process	Probability Frequent	Severity Negligible	Low
3.2.2	> For a sample of system components, examine data sources, including but not limited to the following, and verify that the three-digit or four-digit card verification code or value printed on the front of the card or the signature panel (CVV2, CVC2, CID, CAV2 data) is not stored after authorization.	Do not sample. Develop regex and automate.	insufficient process	Probability Frequent	Severity Negligible	Low
3.2.3	> For a sample of system components, examine data sources, including but not limited to the following and verify that PINs and encrypted PIN blocks are not stored after authorization.	Do not sample. Develop regex and automate.	insufficient process	Probability Frequent	Severity Negligible	Low
3.6.4.b	> Interview personnel to verify that keys are changed at the end of the defined cryptoperiod(s).	Physically verify key change. Do not take someone's word for it	insufficient process	Probability Seldom	Severity Negligible	Low
For Assessors: Sampling of Business Facilities/System Components	> The sample must be large enough to provide the assessor with reasonable assurance that all business facilities/system components are configured per the standard processes.	A specific confidence level and interval would be better, and would allow assessors to calculate the sample size. For example, the standard could say 90% confidence +/- 5%.	verbiage	Probability Likely	Severity Moderate	Low
2.1	> Always change vendor-supplied defaults and remove or disable unnecessary default accounts before installing a system on the network.	Change them to new "unique" and "strong" passwords. Maybe password policies come up later in the doc?	verbiage	Probability Seldom	Severity Moderate	Low
3.4.1.c	> Note: If disk encryption is not used to encrypt removable media, the data stored on this media will need to be rendered unreadable through some other method.	This is insufficient. The data should be encrypted, even if disk encryption is not used. "Unreadable" is not the same as encrypted, and is unacceptable	data vulnerability	Probability Occasional	Severity Moderate	Low
Requirement 4: Encrypt transmission of cardholder data across open, public networks	Encrypt transmission of cardholder data across open, public networks	I don't understand why this requirement applies only to "open, public networks" and not all networks?	verbiage	Probability Likely	Severity Moderate	Low
5.2.d	> Examine anti-virus configurations, including the master installation of the software and a sample of system components, to verify that: Anti-virus software log generation is enabled, and logs are retained in accordance with PCI DSS Requirement 10.7	Missing a requirement to review the logs.	insufficient process	Probability Likely	Severity Moderate	Low
6.1	> Establish a process to identify security vulnerabilities, using reputable outside sources for security vulnerability information, and assign a risk ranking (for example, as high, medium, or low) to newly discovered security vulnerabilities.	Missing a frequency requirement. Should you review every day? Week? Section 6.2 requires installation of critical patches within one month of release, so at least that frequent!	verbiage	Probability Likely	Severity Negligible	Low
8.2.1	> Using strong cryptography, render all authentication credentials (such as passwords/phrases) unreadable during transmission and storage on all system components.	I don't like the implication that encryption renders data "unreadable." The data should be unintelligible to unauthorized parties.	verbiage	Probability Frequent	Severity Negligible	Low
8.2.3	> Passwords/passphrases must meet the following...	NIST guidance changed in 2017: https://pages.nist.gov/800-63-3so800-63b.html	insufficient process	Probability Likely	Severity Moderate	Low
9.1	> Observe a system administrators attempt to log into consoles for randomly selected systems in the cardholder data environment and verify that they are locked to prevent unauthorized use.	Automate.	insufficient process	Probability Likely	Severity Moderate	Low
9.4.4.b	> 9.4.4.b Verify that the log contains: The visitors name. The firm represented, and The onsite personnel authorizing physical access.	Missing date and time.	insufficient process	Probability Likely	Severity Negligible	Low
12.7	> Inquire with Human Resource department management and verify that background checks are conducted (within the constraints of local laws) prior to hire on potential personnel who will have access to cardholder data or the cardholder data environment.	Needs to require correlation between local, state, and federal background checks, to include municipalities and states of previous residence.	insufficient process	Probability Seldom	Severity Moderate	Low
12.6.2	> Require personnel to acknowledge at least annually that they have read and understood the security policy and procedures.	This is good, but perhaps add testing to validate understanding?	insufficient process	Probability Likely	Severity Negligible	Low
Appendix A2	> Additional PCI DSS Requirements for Entities using SSL/TLS	What defines "early TLS"? Version 1.1 and below? Should specify what versions are acceptable.	verbiage	Probability Seldom	Severity Moderate	Low

5.1.1 antivirus programs	"Ensure that anti-virus programs are capable of detecting, removing, and protecting against all known types of malicious software..."	"Ensure that anti-virus programs are capable of detecting, removing, and protecting against all known types of malicious software..." I noted that they should also check that AV signatures are updated regularly. Also might want metrics on <u>refining false positives</u>	insufficient process	Probability Likely	Severity Moderate	Low
2.5 Security Policy	"Ensure that security policies and operational procedures for managing vendor defaults and other security parameters are documented, in use, and known to all affected parties..."	HTRUST has 2 requirements for policies & procedures that are annoying to audit against but sort of help make policies & procedures stronger: - need annual review of policies; they should be updated as at least reviewed annually - need the users to acknowledge they received the policies and read them This control looks for policies & procedures in place, but doesn't have the review/acknowledgment enforcement parts in place, which could help make sure the policies aren't written once and forgotten about	insufficient process	Probability Seldom	Severity Negligible	Low
9.3.1.2 Account Management (AC-2) - shared accounts	> Establish a process for reissuing shared/group account credentials (if > deployed) when individuals are removed from the group.	Nope, just nope. While I agree there should be a process for removing accounts from a group. There should be "NO!" shared account access in FTL. This defeats any auditing actions in place	data vulnerability	Frequent	Moderate	Low
9.3.1.2 Account Management: User Monitoring	"Monitor the use of information system accounts..."	Monitor is not well defined here. Are we talking about email monitoring? Keystroke logging? Event logs? PCAPS/ Network Monitoring?	insufficient process	Likely	Moderate	Low
9.3.1.9 Session Lock (AC-11)	> Prevent further access to the system by initiating a session lock after 15 minutes > of inactivity or upon receiving a request from a user	This is pretty secure, but also pretty unusable. Not really a vulnerability, but could be one if users try to develop mechanisms to circumvent this protection. (Remind me to tell you about our novel use of a mouse and a clock sometime...)	insufficient process	Likely	Moderate	Low
9.3.11.5 Access Control for Output Devices (PE-5)	> The agency must control physical access to information system output devices to prevent unauthorized individuals from obtaining the output. Monitors, printers, copiers, scanners, fax machines, and audio devices are examples of information system output devices.	What about network taps? Rogue access points?	data vulnerability	Seldom	Moderate	Low
9.3.11.6 Monitoring Physical Access (PE-6)	> Review physical access logs annually	A yearly review of physical access logs is ineffective.	insufficient process	Likely	Negligible	Low
9.3.13.3 Personnel Screening (PS-3)	The agency must: a. Screen individuals prior to authorizing access to the information system b. Rescreen individuals according to agency-defined conditions requiring rescreening The agency, upon termination of individual employment must: a. Disable information system access b. Terminate/revoke any authenticators/credentials associated with the individual c. Conduct exit interviews, as needed d. Retrieve all security-related agency information system-related property e. Retain access to agency information and information systems formerly controlled by the terminated individual f. Notify agency personnel upon termination of the employee	> Screen individuals prior to authorizing access to the information system Screen how? Federal database? Local? How is this verified? Where is the check and balance? Additionally, there is no time period associated with revocation or transfer. Recommendation: Revoke access immediately before notifying employee of termination to prevent access to sensitive data.	insufficient process	Frequent	Critical	Low
9.3.13.4 Termination (PS-4) [and 9.3.13.5 Personnel Transfer (PS-5)]			data vulnerability	Likely	Critical	Low
9.3.14.3 Vulnerability Scanning (RA-5)	Remediate legitimate vulnerabilities in accordance with an assessment of risk	No fixed requirement for changing / patching / fixing ID'd vulns (d.) says "Remediate legitimate vulnerabilities in accordance with an assessment of risk" but whose assessment? Without a timeline suspense, this may never get fixed.	data vulnerability	Seldom	Moderate	Low
9.3.15.4 Acquisition Process (SA-4)	The agency must include the following requirements, descriptions, and criteria, explicitly or by reference, in the acquisition contract for the information system, system component, or information system service in accordance with applicable federal laws, Executive Orders, directives, policies, regulations, standards, guidelines, and agency mission/business needs: a. Security functional requirements b. Security strength requirements c. Security assurance requirements d. Security-related documentation requirements e. Requirements for protecting security-related documentation f. Description of the information system development environment and environment in which the system is intended to operate g. Acceptance criteria h. When applicable, the agency must require the developer of the information system, system component, or information system service to provide a description of the functional properties of the security controls to be employed (CE1)	No mention of supply chain security or authorized purchase locations. Example: (https://arstechnica.com/information-technology/2016/11/chinese-company-installed-secret-backdoor-on-hundreds-of-thousands-of-phones/)(https://arstechnica.com/information-technology/2016/11/chinese-company-installed-secret-backdoor-on-hundreds-of-thousands-of-phones/)	data vulnerability	Likely	Critical	Low
9.3.15.8 Developer Configuration Management (SA-10)	> Track security flaws and flaw resolution within the system, component, or service > and report findings to designated agency officials	I don't know what this means, but developers shouldn't note security vulnerabilities in the systems the same system that's vulnerable.	insufficient process	Likely	Negligible	Low
9.3.16.5 Boundary Protection	> The agency must limit the number of external network connections to the information system. (CE3)	Recommend that a maximum number or other quantifiable limit be specified	unenforceable	Likely	Negligible	Low
9.3.17.3 Malicious Code Protection (SI-3)	Malicious code protection includes antivirus software and anti-malware and intrusion detection systems. The agency must: a. Employ malicious code protection mechanisms at information system entry and exit points to detect and eradicate malicious code b. Update malicious code protection mechanisms whenever new releases are available in accordance with the agency's security policy	What happens when the AV is the point of infection? Nowhere is there a forcing function for vulns to be fixed in a timely manner. Recommendation: Whitelist applications, prevent chained installations of child applications from whitelisted apps (e.g., Chrome cannot install anything nor can Norton AV)	data vulnerability	Occasional	Critical	Low
9.3.17.6 Spam Protection (SI-5)	> Malicious code protection includes antivirus software and anti-malware and intrusion detection systems	I don't think that Spam_ is the issue here. They should be more concerned with phishing attacks.	verbiage issue	Seldom	Negligible	Low
9.3.17.7 Information Input Validation (SI-10)	The information system must check the validity of information inputs.	Incredibly vague. How is this check performed? Fuzzing? Manual audit? Additionally, if org is using a third-party application and a vuln is found through pentesting or fuzzing, what then?	verbiage issue	Likely	Negligible	Low
9.3.2.3 Role-Based Security Training (AT-3)	> Note: Training conducted under this section is distinct from Section 6.3, Disclosure > Awareness, and Section 9.3.2.2, Security Awareness Training (AT-2).	Just to pick a nit. This role based training should encompass the concepts in 9.3.2.2. "DO NOT" make this role based training an addendum. It will be seen as punitive and reduces the effectiveness, since employees will try everything to reduce their mandatory training burden. (-: know from experience.)	insufficient process	Seldom	Negligible	Low
9.3.3.2 Audit Events (AU-2)	> a. Determine that the information system is capable, at a minimum, of auditing the following event types: > 4. Changes made to an application or database by a batch file	I find the list unsatisfying. What does this mean, and why batch files? What about printing? Large network transfers? Network connections in general (ie netflow)? At a minimum, I think the list needs to cover more general areas (user behavior, system behavior, data access and modification, data transfer, etc)	data vulnerability	Occasional	Moderate	Low
9.3.3.5 Response to Audit Processing Failures (AU-5)	> Provide a warning when allocated audit record storage volume reaches a > maximum audit record storage capacity (CE1)	Should provide a warning "BEFORE" reaching audit record storage capacity.	insufficient process	Likely	Moderate	Low
9.3.3.6 Audit Review, Analysis, and Reporting (AU-6)	> a. Review and analyze information system audit records at least weekly or more frequently at the discretion of the information system owner for indications of unusual activity related to potential unauthorized FTL access	This isn't horrible, but I'd prefer they do continual monitoring and identify indications of unusual activity more frequently than weekly	insufficient process	Unlikely	Moderate	Low
9.3.4.2 Security Assessments (CA-2)	The agency must: a. Develop a security assessment plan that describes the scope of the assessment, including: 1. Security controls and control enhancements under assessment 2. Assessment procedures to be used to determine security control effectiveness 3. Assessment environment, assessment team, and assessment roles and responsibilities b. Assess the security controls in the information system and its environment at a minimum on an annual basis to determine the extent to which the controls are implemented correctly, operating as intended, and producing the desired outcome with respect to meeting established security requirements	Does not specify vuln assessment, pentest, or expectations of scope for annual assessment. Wholly insufficient.	insufficient process	Seldom	Moderate	Low
9.3.4.3 System Interconnections (CA-3)	> Authorize connections from the information system to other information systems > through the use of Interconnection Security Agreements > Document, for each interconnection, the interface characteristics, security > requirements, and the nature of the information communicated	This should also require a validation/audit requirement. These connection documents get outdated very fast and details are often missed. Need a mechanism to test whether systems have additional interconnects not stipulated in the documentation (or that currently documented interconnects still exist).	insufficient process	Occasional	Moderate	Low
9.3.4.4 Plan of Action and Milestones (CA-5)	a. Develop a POA&M for the information system to document the agency's planned remedial actions to correct weaknesses or deficiencies noted during the assessment of the security controls and to reduce or eliminate known vulnerabilities in the system b. Update the existing POA&M on a quarterly basis, at a minimum, based on the findings from security controls assessments, security impact analyses, and continuous monitoring activities	No sanctions or consequence for non-compliance. EDIT: 9.3.13.8 Personnel Sanctions (PS-8) talks about personnel sanctions, not organization sanctions EDIT: Exhibit 3 USC Title 26, CFR 301.6103(p)(7)-1 says IRS can terminate or suspend access to FTL. There is no enforcement mechanism for making them purge existing FTL.	data vulnerability	Likely	Moderate	Low
9.3.4.5 Security Authorization (CA-6)	> Assign a senior-level executive or manager as the authorizing official for the > information system	No consequences outlined for this individual.	insufficient process	Likely	Negligible	Low
9.3.5.10 Software Usage Restrictions (CM-10)	>Control and document the use of peer-to-peer file sharing technology to ensure > that this capability is not used for the unauthorized distribution, display, > performance, or reproduction of copyrighted work	Should update to include cloud sharing services (Google Drive, Dropbox, etc.)	data vulnerability	Likely	Moderate	Low
9.3.5.11 User-Installed Software (CM-11)	The agency must: a. Establish policies governing the installation of software by users b. Enforce software installation policies through automated methods c. Monitor policy compliance on a continual basis	Don't allow users to install software. Or at least institute application whitelisting.	data vulnerability	Likely	Critical	Low

9.3.5.8 Information System Component Inventory (CM-8)	The agency must: a. Develop and document an inventory of information system components that: 1. Accurately reflects the current information system 2. Includes all components that store, process, or transmit FTI 3. Is at the level of granularity deemed necessary for tracking and reporting 4. Includes information deemed necessary to achieve effective information system component accountability b. Review and update the information system component inventory through periodic manual inventory checks or a network monitoring tool that automatically maintains the inventory c. Update the inventory of information system components as an integral part of component installations, removals, and information system updates (CE1)	No temporal requirement for maintaining updated inventories. This lends itself to rogue systems being on the network for an enduring amount of time. Furthermore, the "or" statement of manual inspection or automated assessments need to be tied to a ground truth → asset inventory through supply systems ("do we own the devices on our network?"). Google SRE talks about this being a core tenet of their security. Failure to ID ground truth could permit rogue computers to be whitelisted and treated as legit systems.	data vulnerability	Occasional	Moderate	Low
9.3.6.8 Information System Recovery and Reconstitution (CP-10)	The agency must provide for the recovery and reconstitution of the information system to a known state after a disruption, compromise, or failure.	Reverting an unpatched system to a previous unpatched state helps nothing.	data vulnerability	Likely	Moderate	Low
9.3.7.2 Identification and Authentication (Organizational Users) (A-2)	The information system must: a. Uniquely identify and authenticate agency users (or processes acting on behalf of agency users) b. Implement multi-factor authentication for all remote network access to privileged and non-privileged accounts for information systems that receive, process, store, or transmit FTI. (CE1, CE2) c. Implement multi-factor authentication for remote access to privileged and non-privileged accounts such that one of the factors is provided by a device separate from the system gaining access. NIST SP 800-63 allows the use of software tokens. (CE11)	Why only MFA for a subset of systems? Why not MFA for any system that touches FTI?	data vulnerability	Likely	Moderate	Low
9.3.7.5 Authenticator Management (A-5)	The agency must manage information system authenticators by: a. Verifying, as part of the initial authenticator distribution, the identity of the individual, group, role, or device receiving the authenticator b. Establishing initial authenticator content for authenticators defined by the agency c. Ensuring that authenticators have sufficient strength of mechanism for their intended use d. Establishing and implementing administrative procedures for initial authenticator distribution, for lost/compromised or damaged authenticators, and for revoking authenticators e. Changing default content of authenticators prior to information system installation f. Establishing minimum and maximum lifetime restrictions and reuse conditions for authenticators g. Changing/refreshing authenticators h. Protecting authenticator content from unauthorized disclosure and modification i. Requiring individuals to take, and having devices implement, specific security safeguards to protect authenticators j. Changing authenticators for group/role accounts when membership to those accounts changes The information system must, for password-based authentication: a. Enforce minimum password complexity of: 1. Eight characters 2. At least one numeric and at least one special character 3. A mixture of at least one uppercase and at least one lowercase letter 4. Storing and transmitting only encrypted representations of passwords b. Enforce password minimum lifetime restriction of one day c. Enforce non-privileged account passwords to be changed at least every 90 days d. Enforce privileged account passwords to be changed at least every 60 days e. Prohibit password reuse for 24 generations f. Allow the use of a temporary password for system logon requiring an immediate change to a permanent password g. Password-protect system initialization (boot) settings	Shitty password change policies that NIST no longer recommends.	data vulnerability	Frequent	Moderate	Low
9.3.8.3 Incident Response Testing (R-3)	Agencies entrusted with FTI must test the incident response capability at least annually. a. Agencies must perform tabletop exercises using scenarios that include a breach of FTI and should test the agency's incident response policies and procedures. b. A subset of all employees and contractors with access to FTI must be included in tabletop exercises. c. Each tabletop exercise must produce an after-action report to improve existing processes, procedures, and policies.	A tabletop exercise is the only requirement? This goes back to [security assessments](https://github.com/stevens70/tracking-compliance-issues#44) being horribly defined. You evaluate SCPs and policies through live training against a pentest team.	insufficient process	Likely	Moderate	Low
9.3.8.4 Incident Handling (R-4)	> Implement an incident handling capability for security incidents that includes > preparation, detection and analysis, containment, eradication, and recovery	Does not cover evidence preservation/forensics.	insufficient process	Likely	Negligible	Low
9.3.8.9 Information Spillage Response (R-9)	The agency must respond to information spills by: a. Identifying the specific information involved in the information system contamination b. Alerting authorized incident response personnel of the information spill using a method of communication not associated with the spill c. Isolating the contaminated information system or system component d. Eradicating the information from the contaminated information system or component e. Identifying other information	This section doesn't say anything about documenting information spillage or providing an after-action step to learn from mistakes.	insufficient process	Occasional	Moderate	Low
9.4.1 Cloud Computing Environments	> Data is not stored in an agency-managed data center	This is true only of public clouds. Private clouds may reside in an agency-managed data center.	verbiage issue	Occasional	Negligible	Low
9.4.18 Wireless Networks	Requirements To use FTI in an 802.11 WLAN, the agency must meet the following mandatory requirements: a. The agency should have WLAN management controls that include security policies and procedures, a complete inventory of all wireless network components, and standardized security configurations for all components. b. WLAN hardware (access points, servers, routers, switches, firewalls) must be physically protected in accordance with the minimum protection standards for physical security outlined in Section 4.0, Secure Storage—IRC 6103(p)(4)(B). c. Each system within the agency's network that transmits FTI through the WLAN is hardened in accordance with the requirements in this publication. d. The WLAN is architected to provide logical separation between WLANs with different security profiles and from the wired LAN. e. WLAN infrastructure that receives, processes, stores, or transmits FTI must comply with the Institute of Electrical and Electronic Engineers 802.11i wireless security standard and perform mutual authentication for all access to FTI via 802.1X and extensible authentication protocol f. Vulnerability scanning should be conducted as part of periodic technical security assessments for the organization's WLAN. g. Wireless intrusion detection is deployed to monitor for unauthorized access, and security event logging is enabled on WLAN components in accordance with Section 9.3.3, Audit and Accountability. h. Disposal of all WLAN hardware follows media sanitization and disposal procedures in Section 9.3.10.6, Media Sanitization (MP-6), and Section 9.4.7, Media Sanitization	No mention of rogue access points. No mention of logical segmentation for each user session (logical switch vs hub).	data vulnerability	Likely	Critical	Low
9.4.8 Mobile Devices	> Access to hardware, such as the digital camera, global positioning system > (GPS), and universal serial bus (USB) interface, must be disabled to the extent > possible	I disagree with the GPS portion of this statement. Using location data to geo-fence, located a device, or perform and out of bounds remote wipe are all strategies that should be used for mobile.	insufficient process	Likely	Negligible	Low
Table 5 - Evidentiary Requirements for SSR approval before release of FTI	ATC-17, Remote Access Screenshot of authentication screens Document how multi-factor authentication is deployed for all remote network access to systems containing FTI and the tokens used for authentication Section 6.2, Corriging and Labeling Screenshots of database schemas that show electronic FTI labeling Sample output (report/notification) that shows how FTI is labeled	The use of screenshots to prove evidentiary requirements feels squishy to me. These could be easily forged/photoshopped. They are not strong evidence of safeguards.	data vulnerability	Likely	Moderate	Low

Compliance Audit Results for NERC CIP 007-6

Section	Text of concern	Concern	Issue	Probability	Severity	Impact
2.1	A patch management process for tracking, evaluating, and installing cyber security patches for applicable Cyber Assets.	Recommend a test environment for patching before applying to live system	insufficient process	Probability Frequent	Severity Critical	Low
1.1	If a device has no provision for disabling or restricting logical ports on the device then those ports that are open are deemed needed.	This document should mandate that hardware facilitates this requirement or remove it from the standard completely (because it's essentially a recommendation)	insufficient process	Probability Frequent	Severity Moderate	Low
3.1	> Deploy method(s) to deter, detect, or prevent malicious code.	Reference specific procedures like whitelisting, etc. Don't leave it up to the operator to reinvent the wheel.	insufficient process	Probability Frequent	Severity Moderate	Low
2.1 / 2.2	The tracking portion shall include the identification of a source or sources that the Responsible Entity tracks. 2.2 references this language in 2.1 as such > evaluate security patches for applicability that have been released since the last evaluation from the source or sources identified in Part 2.1	It appears that this control only applies to sources the Responsible Entity identifies, which opens a possible loophole that if the RE doesn't identify any sources, they are not required to do patch management. Also, if a security is identified, but does not come from an "source" the RE identifies, what is the requirement for patching?	insufficient process	Probability Occasional	Severity Critical	Low
4.3	"_Where technically feasible, retain applicable event logs identified in Part 4.1 for at least the last 90 consecutive calendar days except under CIP Exceptional Circumstances_" Additionally, section 1.2 under compliance states the following: " Each Responsible Entity shall retain evidence of each requirement in this standard for three calendar years_"	Dwell time for adversaries without detection is usually much longer than 90 days. I know of several public compromises of high profile companies where the only way the mechanisms of compromise were determined were through forensic evaluation of months of event log data. Event log data should be retained and backed up / duplicated frequently. Acceptable retention periods can vary, but I would recommend keeping event log data in an indexable, easily accessible location for up to a year with 2 to 5 year of long term storage. This data will be critical to your response team in the event of a compromise. 1.2 seems a bit hypocritical... RE's are responsible for keeping proof of their compliance for 3 years, yet we are rolling all of the useful event log data after 90 days...	insufficient process	Probability Frequent	Severity Moderate	Low
5.4	> Documentation in system manuals or other vendor documents showing default vendor passwords were generated pseudo-randomly and are thereby unique to the device.	> Documentation in system manuals or other vendor documents showing default vendor passwords were generated pseudo-randomly and are thereby unique to the device. Pseudorandom also might be predictable. Do not rely on vendors to set your pws.	insufficient process	Probability Seldom	Severity Moderate	Low
5.5.1	> Password length that is, at least, the lesser of eight characters or the maximum length supported by the Cyber Asset	> Password length that is, at least, the lesser of eight characters or the maximum length supported by the Cyber Asset Pin code devices are easy to brute and should not be allowed	insufficient process	Probability Occasional	Severity Moderate	Low
5.1	> Have a method(s) to enforce authentication of interactive user access, where technically feasible.	This should be a hardware/software requirement. "Feasible" would allow obsolete things to remain in use	data vulnerability	Probability Likely	Severity Moderate	Low
5.7	> Where technically feasible, either: Limit the number of unsuccessful authentication attempts; or Generate alerts after a threshold of unsuccessful authentication attempts.	> Where technically feasible, either: Limit the number of unsuccessful authentication attempts; or Generate alerts after a threshold of unsuccessful authentication attempts. Additionally, alert abnormalities like 2am logins or logins from foreign nations	insufficient process	Probability Frequent	Severity Negligible	Low
3.3	> For those methods identified in Part 3.1 that use signatures or patterns, have a process for the update of the signatures or patterns. The process must address testing and installing the signatures or patterns.	> For those methods identified in Part 3.1 that use signatures or patterns, have a process for the update of the signatures or patterns. The process must address testing and installing the signatures or patterns. What are the timeframe requirements?	insufficient process	Probability Occasional	Severity Negligible	Low
4.1.X	Log events at the BES Cyber System level (per BES Cyber System capability) or at the Cyber Asset level (per Cyber Asset capability) for identification of, and after-the-fact investigations of, Cyber Security Incidents that includes, as a minimum, each of the following types of events:	Although likely covered in general as 4.1.1 'log all access logins', specific logging should be applied to any accounts elevating privs or switching users. It's possible that this information is logged if the log includes enough information on the source of the login attempt/access	insufficient process	Probability Likely	Severity Negligible	Low
1.2	> Protect against the use of unnecessary physical input/output ports used for network connectivity, console commands, or Removable Media.	what constitutes necessary? should reference explicit guidance	insufficient process	Probability Occasional	Severity Negligible	Low
2.4	> For each mitigation plan created or revised in Part 2.3, implement the plan within the timeframe specified in the plan, unless a revision to the plan or an extension to the timeframe specified in Part 2.3 is approved by the CIP Senior Manager or delegate.	> For each mitigation plan created or revised in Part 2.3, implement the plan within the timeframe specified in the plan, unless a revision to the plan or an extension to the timeframe specified in Part 2.3 is approved by the CIP Senior Manager or delegate. Specify guidelines for implementation relative to risk (critical --> immediate, negligible --> 30 days)	insufficient process	Probability Occasional	Severity Negligible	Low
3.2	> Examples of evidence may include, but are not limited to: Records of response processes for malicious code detection. Records of the performance of these processes when malicious code is detected.	> Examples of evidence may include, but are not limited to: Records of response processes for malicious code detection. Records of the performance of these processes when malicious code is detected. What specifics should technicians record? Playbook creation? time to ID, time to quarantine? what metrics? why?	insufficient process	Probability Occasional	Severity Negligible	Low
3.3	> For those methods identified in Part 3.1 that use signatures or patterns, have a process for the update of the signatures or patterns. The process must address testing and installing the signatures or patterns.	> For those methods identified in Part 3.1 that use signatures or patterns, have a process for the update of the signatures or patterns. The process must address testing and installing the signatures or patterns. What about sharing those IOCs across organizations?	insufficient process	Probability Occasional	Severity Negligible	Low
4	Log events at the BES Cyber System level (per BES Cyber System capability) or at the Cyber Asset level (per Cyber Asset capability)	This session specifies many low level event that require logging, but does not specify a minimum set of criteria to include in those logs. Define the minimum telemetry (datetime, system identifier, event time, event risk classification, event severity, etc.)	insufficient process	Probability Likely	Severity Negligible	Low
5.5.2	>Minimum password complexity that is the lesser of three or more different types of characters (e.g., uppercase alphabetic, lowercase alphabetic, numeric, non- alphanumeric) or the maximum complexity supported by the Cyber Asset.	Passwords should have high entropy, but enforcing the above standards have proven to be ineffective.	insufficient process	Probability Likely	Severity Negligible	Low
5.3	> Identify individuals who have authorized access to shared accounts.	> Identify individuals who have authorized access to shared accounts. Don't share accounts!!!	insufficient process	Probability Likely	Severity Moderate	Low
5.6	> Where technically feasible, for password-only authentication for interactive user access, either technically or procedurally enforce password changes or an obligation to change the password at least once every 15 calendar months.	> Where technically feasible, for password-only authentication for interactive user access, either technically or procedurally enforce password changes or an obligation to change the password at least once every 15 calendar months. Obsolete practice, change on by-need basis	insufficient process	Probability Frequent	Severity Negligible	Low
4	Each Responsible Entity shall implement one or more documented process(es) that collectively include each of the applicable requirement parts in CIP-007-6 Table R4	Maybe I missed it, but I don't see anything on network security monitoring. Security posture should include both host and network oriented security monitoring, and prevention.	insufficient process	Probability Likely	Severity Moderate	Low
2.2	"_At least once every 35 calendar days, evaluate security patches for applicability that have been released since the last evaluation from the source or sources identified in part 2.1_"	This does not provide a provision for frequency of patch evaluation based on severity of risk. I believe if there is a very severe vulnerability discovered, then there should be a system for immediate patch validation without waiting the specified 35 calendar days.	insufficient process	Probability Likely	Severity Moderate	Low

Compliance Audit Results for FedRAMP						
Section	Text of concern	Concern	Issue	Probability	Severity	Impact
AC-17 (2)	The information system implements cryptographic mechanisms to protect the confidentiality and integrity of remote access sessions. Supplemental Guidance: The encryption strength of mechanism is selected based on the security categorization of the information. Related controls: SC-8, SC-12, SC-13.	Does not specify storage location or protection mechanisms for keys; could be in adversarial control; No specifications for minimum cryptographic mechanisms	Risk to Data	Frequent	Critical	Extremely High
AC-18 (1)	The information system protects wireless access to the system using authentication of [Selection (one or more): users; devices] and encryption. Supplemental Guidance: Related controls: SC-8, SC-13.	should disallow weak passwords and encryption algos	Ambiguous Specification	Likely	Critical	High
AC-4	The information system enforces approved authorizations for controlling the flow of information within the system and between interconnected systems based on [Assignment: organization-defined information flow control policies]. Supplemental Guidance: Information flow control regulates where information is allowed to travel within an information system and between information systems (as opposed to who is allowed to access the information) and without explicit regard to subsequent accesses to that information. Flow control restrictions include, for example, keeping export-controlled information from being transmitted in the clear to the Internet, blocking outside traffic that claims to be from within the organization, restricting web requests to the Internet that are not from the internal web proxy server, and limiting information transfers between organizations based on data structures and content. Transferring information between information systems representing different security domains with different security policies introduces risk that such transfers violate one or more domain security policies. In such situations, information owners/stewards provide guidance at designated policy enforcement points between interconnected systems. Organizations consider mandating specific architectural solutions when required to enforce specific security policies. Enforcement includes, for example: (i) prohibiting information transfers between interconnected systems (i.e., allowing access only); (ii) employing hardware mechanisms to enforce one-way information flows; and (iii) implementing trustworthy regarding mechanisms to reassign security attributes and security labels. Organizations commonly employ information flow control policies and enforcement mechanisms to control the flow of information between designated sources and destinations (e.g., networks, individuals, and devices) within information systems and between interconnected systems. Flow control is based on the characteristics of the information and/or the information path. Enforcement occurs, for example, in boundary protection devices (e.g., gateways, routers, guards, encrypted tunnels, firewalls) that employ rule sets or establish configuration settings that restrict information system services, provide a packet-filtering capability based on header information, or message-filtering capability based on message content (e.g., implementing key word searches or using document characteristics). Organizations also consider the trustworthiness of filtering/inspection mechanisms (i.e., hardware, firmware, and software components) that are critical to information flow enforcement. Control enhancements 3 through 22 primarily address cross-domain solution needs which focus on more advanced filtering techniques, in-depth analysis, and stronger flow enforcement mechanisms implemented in cross-domain products, for example, high-assurance guards. Such capabilities are generally not available in commercial off-the-shelf information technology products. Related controls: AC-3, AC-17, AC-19, AC-21, CM-6, CM-7, SA-8, SC-2, SC-5, SC-7, SC-18. References: None.	makes provisions for unencrypted transmission of controlled information via intranet -- does not account for insider threats	Risk to Data	Likely	Critical	High
AC-1	The organization: a. Develops, documents, and disseminates to [Assignment: organization-defined personnel or roles]: 1. An access control policy that addresses purpose, scope, roles, responsibilities, management commitment, coordination among organizational entities, and compliance; and 2. Procedures to facilitate the implementation of the access control policy and associated access controls; and b. Reviews and updates the current: 1. Access control policy [Assignment: organization-defined frequency]; and 2. Access control procedures [Assignment: organization-defined frequency]. Supplemental Guidance: This control addresses the establishment of policy and procedures for the effective implementation of selected security controls and control enhancements in the AC family. Policy and procedures reflect applicable federal laws, Executive Orders, directives, regulations, policies, standards, and guidance. Security program policies and procedures at the organization level may make the need for system-specific policies and procedures unnecessary. The policy can be included as part of the general information security policy for organizations or conversely, can be represented by multiple policies reflecting the complex nature of certain organizations. The procedures can be established for the security program in general and for particular information systems, if needed. The organizational risk management strategy is a key factor in establishing policy and procedures. Related control: PM-9. Control Enhancements: None. References: NIST Special Publications 800-12, 800-100.	Does not require any level of training, certification, or responsible role for creating/managing ACs.	Ambiguous Specification	Unlikely	Critical	Low
AC-2 (1)	The organization employs automated mechanisms to support the management of information system accounts. Supplemental Guidance: The use of automated mechanisms can include, for example: using email or text messaging to automatically notify account managers when users are terminated or transferred; using the information system to monitor account usage; and using telephonic notification to report atypical system account usage.	The supplemental guidance propose insecure mechanisms, including email and text messaging.	Ambiguous Specification	Occasional	Moderate	Low
AC-2 (2)	The information system automatically [Selection: removes; disables] temporary and emergency accounts after [Assignment: organization-defined time period for each type of account]. Supplemental Guidance: This control enhancement requires the removal of both temporary and emergency accounts automatically after a predefined period of time has elapsed, rather than at the convenience of the systems administrator.	Needs maximum time. Organization shouldn't get to pick outrageous time periods.	Ambiguous Specification	Occasional	Moderate	Low
AC-2 (3)	The information system automatically disables inactive accounts after [Assignment: organization-defined time period].	Needs maximum time. Organization shouldn't get to pick outrageous time periods.	Ambiguous Specification	Occasional	Moderate	Low
AC-2 (5)	The organization requires that users log out when [Assignment: organization-defined time-period of expected inactivity or description of when to log out]. Supplemental Guidance: Related control: SC-23.	Needs maximum time. Organization shouldn't get to pick outrageous time periods.	Ambiguous Specification	Occasional	Moderate	Low

AC-6 (9)	<p>The information system audits the execution of privileged functions.</p> <p>Supplemental Guidance: Misuse of privileged functions, either intentionally or unintentionally by authorized users, or by unauthorized external entities that have compromised information system accounts, is a serious and ongoing concern and can have significant adverse impacts on organizations. Auditing the use of privileged functions is one way to detect such misuse, and in doing so, help mitigate the risk from insider threats and the advanced persistent threat (APT). <u>Related control: AU-2</u></p>	Who and how often are audits reviewed?	Ambiguous Specification	Occasional	Moderate	Low
AC-17 (9)	<p>The organization provides the capability to expeditiously disconnect or disable remote access to the information system within [Assignment: organization-defined time period].</p> <p>Supplemental Guidance: This control enhancement requires organizations to have the capability to rapidly disconnect current users remotely accessing the information system and/or disable further remote access. The speed of disconnect or disablement varies based on the criticality of missions/business functions and the need to eliminate immediate or future remote access to organizational information systems.</p>	No maximum time period since organization can define.	Ambiguous Specification	Occasional	Negligible	Low
AC-18	<p>The organization:</p> <p>a. Establishes usage restrictions, configuration/connection requirements, and implementation guidance for wireless access; and</p> <p>b. Authorizes wireless access to the information system prior to allowing such connections.</p> <p>Supplemental Guidance: Wireless technologies include, for example, microwave, packet radio (UHF/VHF), 802.11x, and Bluetooth. Wireless networks use authentication protocols (e.g., EAP/TLS, PEAP), which provide credential protection and mutual authentication. Related controls: AC-2, AC-3, AC-17, AC-19, CA-3, CA-7, CM-8, IA-2, IA-3, IA-8, PL-4, SI-4.</p> <p>References: NIST Special Publications 800-48, 800-94, 800-97.</p>	WCE: This would allow for wifi access within data centers to things like AWS host machines. Thankfully, I don't imagine any cloud providers actually allow datacenter hosts to access wifi, but you could and it would be compliant	Ambiguous Specification	Seldom	Moderate	Low
AC-2 (12)	<p>The organization:</p> <p>(a) Monitors information system accounts for [Assignment: organization-defined atypical use]; and</p> <p>(b) Reports atypical usage of information system accounts to [Assignment: organization-defined personnel or roles].</p> <p>Supplemental Guidance: Atypical usage includes, for example, accessing information systems at certain times of the day and from locations that are not consistent with the normal usage patterns of individuals working in organizations. <u>Related control: CA-7</u></p>	required follow-up action?	Ambiguous Specification	Likely	Moderate	Medium
AC-3	<p>The information system enforces approved authorizations for logical access to information and system resources in accordance with applicable access control policies.</p> <p>Supplemental Guidance: Access control policies (e.g., identity-based policies, role-based policies, attribute-based policies) and access enforcement mechanisms (e.g., access control lists, access control matrices, cryptography) control access between active entities or subjects (i.e., users or processes acting on behalf of users) and passive entities or objects (e.g., devices, files, records, domains) in information systems. In addition to enforcing authorized access at the information system level and recognizing that information systems can host many applications and services in support of organizational missions and business operations, access enforcement mechanisms can also be employed at the application and service level to provide increased information security. Related controls: AC-2, AC-4, AC-5, AC-6, AC-16, AC-17, AC-18, AC-19, AC-20, AC-21, AC-22, AU-9, CM-5, CM-6, CM-11, MA-3, MA-4, MA-5, PE-3.</p> <p>References: None.</p>	Who approves authorizations and how?	Ambiguous Specification	Likely	Moderate	Medium
AC-4	<p>The information system enforces approved authorizations for controlling the flow of information within the system and between interconnected systems based on [Assignment: organization-defined information flow control policies].</p> <p>Supplemental Guidance: Information flow control regulates where information is allowed to travel within an information system and between information systems (as opposed to who is allowed to access the information) and without explicit regard to subsequent accesses to that information. Flow control restrictions include, for example, keeping export-controlled information from being transmitted in the clear to the Internet, blocking outside traffic that claims to be from within the organization, restricting web requests to the Internet that are not from the internal web proxy server, and limiting information transfers between organizations based on data structures and content. Transferring information between information systems representing different security domains with different security policies introduces risk that such transfers violate one or more domain security policies. In such situations, information owners/stewards provide guidance at designated policy enforcement points between interconnected systems. Organizations consider mandating specific architectural solutions when required to enforce specific security policies. Enforcement includes, for example: (i) prohibiting information transfers between interconnected systems (i.e., allowing access only); (ii) employing hardware mechanisms to enforce one-way information flows; and (iii) implementing trustworthy regarding mechanisms to reassign security attributes and security labels.</p> <p>Organizations commonly employ information flow control policies and enforcement mechanisms to control the flow of information between designated sources and destinations (e.g., networks, individuals, and devices) within information systems and between interconnected systems. Flow control is based on the characteristics of the information and/or the information path. Enforcement occurs, for example, in boundary protection devices (e.g., gateways, routers, guards, encrypted tunnels, firewalls) that employ rule sets or establish configuration settings that restrict information system services, provide a packet-filtering capability based on header information, or message-filtering capability based on message content (e.g., implementing key word searches or using document characteristics). Organizations also consider the trustworthiness of filtering/inspection mechanisms (i.e., hardware, firmware, and software components) that are critical to information flow enforcement. Control enhancements 3 through 22 primarily address cross-domain solution needs which focus on more advanced filtering techniques, in-depth analysis, and stronger flow enforcement mechanisms implemented in cross-domain products, for example, high-assurance guards. Such capabilities are generally not available in commercial off-the-shelf information technology products. Related controls: AC-3, AC-17, AC-19, AC-21, CM-6, CM-7, SA-8, SC-2, SC-5, SC-7, SC-18.</p>	WCE: there are some examples given in the supplemental guidance, but this is wide open for the interpretation of the implementor. As long as you do some sort of restriction of information from your FedRAMP system into "interconnected" systems, you're probably compliant here. Often there are pretty weak data classification policies, and they just say "don't copy files with confidential data out of this datastore to employee workstations" without a lot of technical controls to prevent that.	Ambiguous Specification	Likely	Moderate	Medium

AC-6	<p>The organization employs the principle of least privilege, allowing only authorized accesses for users (or processes acting on behalf of users) which are necessary to accomplish assigned tasks in accordance with organizational missions and business functions.</p> <p>Supplemental Guidance: Organizations employ least privilege for specific duties and information systems. The principle of least privilege is also applied to information system processes, ensuring that the processes operate at privilege levels no higher than necessary to accomplish required organizational missions/business functions. Organizations consider the creation of additional processes, roles, and information system accounts as necessary, to achieve least privilege. Organizations also apply least privilege to the development, implementation, and operation of organizational information systems. Related controls: AC-2, AC-3, AC-5, CM-6, CM-7, PL-2.</p> <p>References: None.</p>	does not specify review period	Ambiguous Specification	Likely	Moderate	Medium
AC-7	<p>The information system:</p> <p>a. Enforces a limit of [Assignment: organization-defined number] consecutive invalid logon attempts by a user during a [Assignment: organization-defined time period]; and</p> <p>b. Automatically [Selection: locks the account/node for an [Assignment: organization-defined time period]; locks the account/node until released by an administrator; delays next logon prompt according to [Assignment: organization-defined delay algorithm]] when the maximum number of unsuccessful attempts is exceeded.</p> <p>Supplemental Guidance: This control applies regardless of whether the logon occurs via a local or network connection. Due to the potential for denial of service, automatic lockouts initiated by information systems are usually temporary and automatically release after a predetermined time period established by organizations. If a delay algorithm is selected, organizations may choose to employ different algorithms for different information system components based on the capabilities of those components. Responses to unsuccessful logon attempts may be implemented at both the operating system and the application levels. Related controls: AC-2, AC-9, AC-14, IA-5.</p> <p>References: None.</p>	<p>does not discuss implications across enterprise (cred spraying). MITRE ATT&CK would agree this is insufficient; does not discuss alerting mechanisms or response mechanisms (AU-5 is for audit failures, not attack detection)</p> <p>WCE: the "lock accounts after N unsuccessful attempts" control can be dangerous, because you can fuzz the admin accounts, get them locked, then go do something while they can't log in.</p>	Ambiguous Specification	Likely	Moderate	Medium
AC-2 (9)	<p>The organization only permits the use of shared/group accounts that meet [Assignment: organization-defined conditions for establishing shared/group accounts].</p>	<p>Use aliases or distro lists instead; do not share accounts period</p> <p>WCE: Shared accounts should not be used. Service accounts that can't get an interactive shell and have restrictive permissions, maybe. But should use AWS Security groups (or unix groups, or AD roles) instead of a shared login if something like this is needed.</p>	Risk to Data	Likely	Moderate	Medium
AU-6 (1)	<p>The organization employs automated mechanisms to integrate audit review, analysis, and reporting processes to support organizational processes for investigation and response to suspicious activities.</p> <p>Supplemental Guidance: Organizational processes benefiting from integrated audit review, analysis, and reporting include, for example, incident response, continuous monitoring, contingency planning, and Inspector General audits. Related controls: AU-12, PM-7.</p>	no time period specified (CA-5 only makes provisions for security assessments, not real attacks)	Ambiguous Specification	Occasional	Critical	High
IA-5 (1)	<p>The information system, for password-based authentication:</p> <p>(a) Enforces minimum password complexity of [Assignment: organization-defined requirements for case sensitivity, number of characters, mix of upper-case letters, lower-case letters, numbers, and special characters, including minimum requirements for each type];</p> <p>(b) Enforces at least the following number of changed characters when new passwords are created: [Assignment: organization-defined number];</p> <p>(c) Stores and transmits only encrypted representations of passwords;</p> <p>(d) Enforces password minimum and maximum lifetime restrictions of [Assignment: organization-defined numbers for lifetime minimum, lifetime maximum];</p> <p>(e) Prohibits password reuse for [Assignment: organization-defined number] generations; and</p> <p>(f) Allows the use of a temporary password for system logons with an immediate change to a permanent password.</p> <p>Supplemental Guidance: This control enhancement applies to single-factor authentication of individuals using passwords as individual or group authenticators, and in a similar manner, when passwords are part of multifactor authenticators. This control enhancement does not apply when passwords are used to unlock hardware authenticators (e.g., Personal Identity Verification cards). The implementation of such password mechanisms may not meet all of the requirements in the enhancement. Encrypted representations of passwords include, for example, encrypted versions of passwords and one-way cryptographic hashes of passwords. The number of changed characters refers to the number of changes required with respect to the total number of positions in the current password. Password lifetime restrictions do not apply to temporary passwords. Related control: IA-6.</p>	Not current with NIST 800-63	Obsolete Reference	frequent	moderate	High

IA-5	<p>The organization manages information system authenticators by:</p> <ul style="list-style-type: none"> a. Verifying, as part of the initial authenticator distribution, the identity of the individual, group, role, or device receiving the authenticator; b. Establishing initial authenticator content for authenticators defined by the organization; c. Ensuring that authenticators have sufficient strength of mechanism for their intended use; d. Establishing and implementing administrative procedures for initial authenticator distribution, for lost/compromised or damaged authenticators, and for revoking authenticators; e. Changing default content of authenticators prior to information system installation; f. Establishing minimum and maximum lifetime restrictions and reuse conditions for authenticators; g. Changing/refreshing authenticators [Assignment: organization-defined time period by authenticator type]; h. Protecting authenticator content from unauthorized disclosure and modification; i. Requiring individuals to take, and having devices implement, specific security safeguards to protect authenticators; and j. Changing authenticators for group/role accounts when membership to those accounts changes. <p>Supplemental Guidance: Individual authenticators include, for example, passwords, tokens, biometrics, PKI certificates, and key cards. Initial authenticator content is the actual content (e.g., the initial password) as opposed to requirements about authenticator content (e.g., minimum password length). In many cases, developers ship information system components with factory default authentication credentials to allow for initial installation and configuration. Default authentication credentials are often well known, easily discoverable, and present a significant security risk. The requirement to protect individual authenticators may be implemented via control PL-4 or PS-6 for authenticators in the possession of individuals and by controls AC-3, AC-6, and SC-28 for authenticators stored within organizational information systems (e.g., passwords stored in hashed or encrypted formats, files containing encrypted or hashed passwords accessible with administrator privileges). Information systems support individual authenticator management by organization-defined settings and restrictions for various authenticator characteristics including, for example, minimum password length, password composition, validation time window for time synchronous one-time tokens, and number of allowed rejections during the verification stage of biometric authentication. Specific actions that can be taken to safeguard authenticators include, for example, maintaining possession of individual authenticators, not loaning or sharing individual authenticators with others, and reporting lost, stolen, or compromised authenticators immediately. Authenticator management includes issuing and revoking, when no longer needed, authenticators for temporary access such as that required for remote maintenance. Device authenticators include, for example, certificates and passwords. Related controls: AC-2, AC-3, AC-6, CM-6, IA-2, IA-4, IA-8, PL-4, PS-5, PS-6, SC-12, SC-13, SC-17, SC-28.</p> <p>References: OMB Memoranda 04-04, 11-11; FIPS Publication 201; NIST Special Publications 800-73, 800-63, 800-76, 800-78; FICAM Roadmap and Implementation Guidance</p>	<p>RAS: does not specific control mechanisms for authenticators. Can it be replicated? How are seeds protected? Who has access to seed data? WCE: doesn't directly require enough complexity/length. The password "password" could be used & would meet the requirements</p>	Ambiguous Specification	Unlikely	Critical	Low
IA-5 (4)	<p>The organization employs automated tools to determine if password authenticators are sufficiently strong to satisfy [Assignment: organization-defined requirements].</p> <p>Supplemental Guidance: This control enhancement focuses on the creation of strong passwords and the characteristics of such passwords (e.g., complexity) prior to use, the enforcement of which is carried out by organizational information systems in IA-5 (1). Related controls: CA-2, CA-7, RA-5.</p>	<p>RAS: does not consider previous password breaches associated with username or most common passwords from OWASP Top XX</p>	Ambiguous Specification	Occasional	Moderate	Low
IA-5 (2)	<p>The information system, for PKI-based authentication:</p> <ul style="list-style-type: none"> (a) Validates certifications by constructing and verifying a certification path to an accepted trust anchor including checking certificate status information; (b) Enforces authorized access to the corresponding private key; (c) Maps the authenticated identity to the account of the individual or group; and (d) Implements a local cache of revocation data to support path discovery and validation in case of inability to access revocation information via the network. <p>Supplemental Guidance: Status information for certification paths includes, for example, certificate revocation lists or certificate status protocol responses. For PIV cards, validation of certifications involves the construction and verification of a certification path to the Common Policy Root trust anchor including certificate policy processing. Related control: IA-6.</p>	<p>RAS: does not consider access control to PKI keys</p>	Risk to Data	Unlikely	Critical	Low
IA-2 (1)	<p>The information system implements multifactor authentication for network access to privileged accounts.</p> <p>Supplemental Guidance: Related control: AC-6.</p>	<p>WCE: Allows for less secure methods of multifactor authentication, like SMS, instead of more secure ones, like Yubikeys or use of an authenticator app.</p>	Ambiguous Specification	Frequent	Negligible	Medium
IA-2 (2)	<p>The information system implements multifactor authentication for network access to non-privileged accounts.</p>	<p>WCE: Similar to IA-02 (01); Allows for less secure methods of multifactor authentication, like SMS, instead of more secure ones, like Yubikeys or use of an authenticator app.</p>	Ambiguous Specification	Frequent	Negligible	Medium
IA-2 (3)	<p>The information system implements multifactor authentication for local access to privileged accounts.</p> <p>Supplemental Guidance: Related control: AC-6.</p>	<p>WCE: Similar to IA-02 (01); Allows for less secure methods of multifactor authentication, like SMS, instead of more secure ones, like Yubikeys or use of an authenticator app.</p>	Ambiguous Specification	Frequent	Negligible	Medium

IR-1	<p>The organization:</p> <p>a. Develops, documents, and disseminates to [Assignment: organization-defined personnel or roles]:</p> <ol style="list-style-type: none"> 1. An incident response policy that addresses purpose, scope, roles, responsibilities, management commitment, coordination among organizational entities, and compliance; and 2. Procedures to facilitate the implementation of the incident response policy and associated incident response controls; and <p>b. Reviews and updates the current:</p> <ol style="list-style-type: none"> 1. Incident response policy [Assignment: organization-defined frequency]; and 2. Incident response procedures [Assignment: organization-defined frequency]. <p>Supplemental Guidance: This control addresses the establishment of policy and procedures for the effective implementation of selected security controls and control enhancements in the IR family. Policy and procedures reflect applicable federal laws, Executive Orders, directives, regulations, policies, standards, and guidance. Security program policies and procedures at the organization level may make the need for system-specific policies and procedures unnecessary. The policy can be included as part of the general information security policy for organizations or conversely, can be represented by multiple policies reflecting the complex nature of certain organizations. The procedures can be established for the security program in general and for particular information systems, if needed. The organizational risk management strategy is a key factor in establishing policy and procedures. Related control: PM-9.</p> <p>Control Enhancements: None.</p> <p>References: NIST Special Publications 800-12, 800-61, 800-83, 800-100.</p>	<p>RAS: does not specify how controls/procedures are developed or improved (no feedback loop). should be derived from best practices and improved upon from exercising controls</p>	<p>Ambiguous Specification</p>	<p>Likely</p>	<p>Moderate</p>	<p>Medium</p>
MP-7	<p>The organization [Selection: restricts; prohibits] the use of [Assignment: organization-defined types of information system media] on [Assignment: organization-defined information systems or system components] using [Assignment: organization-defined security safeguards].</p> <p>Supplemental Guidance: Information system media includes both digital and non-digital media. Digital media includes, for example, diskettes, magnetic tapes, external/removable hard disk drives, flash drives, compact disks, and digital video disks. Non-digital media includes, for example, paper and microfilm. This control also applies to mobile devices with information storage capability (e.g., smart phones, tablets, E-readers). In contrast to MP-2, which restricts user access to media, this control restricts the use of certain types of media on information systems, for example, restricting/prohibiting the use of flash drives or external hard disk drives. Organizations can employ technical and nontechnical safeguards (e.g., policies, procedures, rules of behavior) to restrict the use of information system media. Organizations may restrict the use of portable storage devices, for example, by using physical cages on workstations to prohibit access to certain external ports, or disabling/removing the ability to insert, read or write to such devices. Organizations may also limit the use of portable storage devices to only approved devices including, for example, devices provided by the organization, devices provided by other approved organizations, and devices that are not personally owned. Finally, organizations may restrict the use of portable storage devices based on the type of device, for example, prohibiting the use of writeable, portable storage devices, and implementing this restriction by disabling or removing the capability to write to such devices. Related controls: AC-19, PL-4.</p> <p>References: None.</p>	<p>Organization shouldn't get to define the safeguards</p>	<p>Ambiguous Specification</p>	<p>occasional</p>	<p>moderate</p>	<p>Low</p>
MP-5 (4)	<p>The information system implements cryptographic mechanisms to protect the confidentiality and integrity of information stored on digital media during transport outside of controlled areas.</p> <p>Supplemental Guidance: This control enhancement applies to both portable storage devices (e.g., USB memory sticks, compact disks, digital video disks, external/removable hard disk drives) and mobile devices with storage capability (e.g., smart phones, tablets, E-readers). Related control: MP-2.</p> <p>References: FIPS Publication 199; NIST Special Publication 800-60.</p>	<p>RAS: controls access to the physical device and encrypts data, but does not protect keys/passwords used to protect data</p>	<p>Risk to Data</p>	<p>Unlikely</p>	<p>Critical</p>	<p>Low</p>
PE-13 (3)	<p>The organization employs an automatic fire suppression capability for the information system when the facility is not staffed on a continuous basis.</p>	<p>RAS: who pentests the automated protection? a server room flood would be an enormous loss</p>	<p>Ambiguous Specification</p>	<p>Unlikely</p>	<p>Critical</p>	<p>Low</p>

RA-3	<p>The organization:</p> <p>a. Conducts an assessment of risk, including the likelihood and magnitude of harm, from the unauthorized access, use, disclosure, disruption, modification, or destruction of the information system and the information it processes, stores, or transmits;</p> <p>b. Documents risk assessment results in [Selection: security plan; risk assessment report; [Assignment: organization-defined document]];</p> <p>c. Reviews risk assessment results [Assignment: organization-defined frequency];</p> <p>d. Disseminates risk assessment results to [Assignment: organization-defined personnel or roles]; and</p> <p>e. Updates the risk assessment [Assignment: organization-defined frequency] or whenever there are significant changes to the information system or environment of operation (including the identification of new threats and vulnerabilities), or other conditions that may impact the security state of the system.</p> <p>Supplemental Guidance: Clearly defined authorization boundaries are a prerequisite for effective risk assessments. Risk assessments take into account threats, vulnerabilities, likelihood, and impact to organizational operations and assets, individuals, other organizations, and the Nation based on the operation and use of information systems. Risk assessments also take into account risk from external parties (e.g., service providers, contractors operating information systems on behalf of the organization, individuals accessing organizational information systems, outsourcing entities). In accordance with OMB policy and related E-authentication initiatives, authentication of public users accessing federal information systems may also be required to protect nonpublic or privacy-related information. As such, organizational assessments of risk also address public access to federal information systems.</p> <p>Risk assessments (either formal or informal) can be conducted at all three tiers in the risk management hierarchy (i.e., organization level, mission/business process level, or information system level) and at any phase in the system development life cycle. Risk assessments can also be conducted at various steps in the Risk Management Framework, including categorization, security control selection, security control implementation, security control assessment, information system authorization, and security control monitoring. RA-3 is noteworthy in that the control must be partially implemented prior to the implementation of other controls in order to complete the first two steps in the Risk Management Framework. Risk assessments can play an important role in security control selection processes, particularly during the application of tailoring guidance, which includes security control supplementation. Related controls: RA-2, PM-9.</p> <p>Control Enhancements: None.</p> <p>References: OMB Memorandum 04-04; NIST Special Publication 800-30, 800-39; Web.idmanagement.gov.</p>	No requirement to remediate or otherwise inform prioritization based on results of risk assessment.	Ambiguous Specification	occasional	moderate	Low
RA-5 (1)	<p>The organization employs vulnerability scanning tools that include the capability to readily update the information system vulnerabilities to be scanned.</p> <p>Supplemental Guidance: The vulnerabilities to be scanned need to be readily updated as new vulnerabilities are discovered, announced, and scanning methods developed. This updating process helps to ensure that potential vulnerabilities in the information system are identified and addressed as quickly as possible. Related controls: SI-3, SI-7.</p>	RAS: should require continual security assessment of scanning products	Risk to Data	Unlikely	Critical	Low
CA-8	<p>The organization conducts penetration testing [Assignment: organization-defined frequency] on [Assignment: organization-defined information systems or system components].</p> <p>Supplemental Guidance: Penetration testing is a specialized type of assessment conducted on information systems or individual system components to identify vulnerabilities that could be exploited by adversaries. Such testing can be used to either validate vulnerabilities or determine the degree of resistance organizational information systems have to adversaries within a set of specified constraints (e.g., time, resources, and/or skills). Penetration testing attempts to duplicate the actions of adversaries in carrying out hostile cyber attacks against organizations and provides a more in-depth analysis of security-related weaknesses/deficiencies. Organizations can also use the results of vulnerability analyses to support penetration testing activities. Penetration testing can be conducted on the hardware, software, or firmware components of an information system and can exercise both physical and technical security controls. A standard method for penetration testing includes, for example: (i) pretest analysis based on full knowledge of the target system; (ii) pretest identification of potential vulnerabilities based on pretest analysis; and (iii) testing designed to determine exploitability of identified vulnerabilities. All parties agree to the rules of engagement before the commencement of penetration testing scenarios. Organizations correlate the penetration testing rules of engagement with the tools, techniques, and procedures that are anticipated to be employed by adversaries carrying out attacks. Organizational risk assessments guide decisions on the level of independence required for personnel conducting penetration testing. Related control: SA-12.</p> <p>References: None.</p>	WCE: Requires a pen test, but the ODV is defined here, I believe by the organization, so it can be infrequent. There also aren't requirements to give the pen test firm access to source code and configurations, or what type of accounts to provision them (i.e. grant them a privileged user account and test see what a malicious insider could do with leaving audit trails or other traces).	Ambiguous Specification	Occasional	Moderate	Low
SC-7 (5)	<p>The information system at managed interfaces denies network communications traffic by default and allows network communications traffic by exception (i.e., deny all, permit by exception).</p> <p>Supplemental Guidance: This control enhancement applies to both inbound and outbound network communications traffic. A deny-all, permit-by-exception network communications traffic policy ensures that only those connections which are essential and approved are allowed.</p>	RAS: whitelisted comms should be continually re-assessed WCE: bit of a meta SI comment... CSPs define where integrity needs to be monitored (SI-7), where input validation needs to be done (SI-10), and what memory protection needs to be used (SI-16). These are all up to the provider to define; they can define weak protections here	Ambiguous Specification	Occasional	Moderate	Low

SC-8	<p>The information system protects the [Selection (one or more): confidentiality; integrity] of transmitted information.</p> <p>Supplemental Guidance: This control applies to both internal and external networks and all types of information system components from which information can be transmitted (e.g., servers, mobile devices, notebook computers, printers, copiers, scanners, facsimile machines). Communication paths outside the physical protection of a controlled boundary are exposed to the possibility of interception and modification. Protecting the confidentiality and/or integrity of organizational information can be accomplished by physical means (e.g., by employing physical distribution systems) or by logical means (e.g., employing encryption techniques). Organizations relying on commercial providers offering transmission services as commodity services rather than as fully dedicated services (i.e., services which can be highly specialized to individual customer needs), may find it difficult to obtain the necessary assurances regarding the implementation of needed security controls for transmission confidentiality/integrity. In such situations, organizations determine what types of confidentiality/integrity services are available in standard, commercial telecommunication service packages. If it is infeasible or impractical to obtain the necessary security controls and assurances of control effectiveness through appropriate contracting vehicles, organizations implement appropriate compensating security controls or explicitly accept the additional risk. Related controls: AC-17, PE-4.</p> <p>References: FIPS Publications 140-2, 197; NIST Special Publications 800-52, 800-77, 800-81, 800-113; CNSS Policy 15; NSTISSI No. 7003.</p>	Both confidentiality and integrity should be mandatory, not an option.	Ambiguous Specification	Occasional	Moderate	Low
SC-12	<p>The organization establishes and manages cryptographic keys for required cryptography employed within the information system in accordance with [Assignment: organization-defined requirements for key generation, distribution, storage, access, and destruction].</p> <p>Supplemental Guidance: Cryptographic key management and establishment can be performed using manual procedures or automated mechanisms with supporting manual procedures. Organizations define key management requirements in accordance with applicable federal laws, Executive Orders, directives, regulations, policies, standards, and guidance, specifying appropriate options, levels, and parameters. Organizations manage trust stores to ensure that only approved trust anchors are in such trust stores. This includes certificates with visibility external to organizational information systems and certificates related to the internal operations of systems. Related controls: SC-13, SC-17.</p> <p>References: NIST Special Publications 800-56, 800-57.</p>	If organization has to establish cryptographic keys then they probably cannot but SSL certs (since that entity establishes the key).	Ambiguous Specification	Occasional	Moderate	Low
SC-28	<p>The information system protects the [Selection (one or more): confidentiality; integrity] of [Assignment: organization-defined information at rest].</p> <p>Supplemental Guidance: This control addresses the confidentiality and integrity of information at rest and covers user information and system information. Information at rest refers to the state of information when it is located on storage devices as specific components of information systems. System-related information requiring protection includes, for example, configurations or rule sets for firewalls, gateways, intrusion detection/prevention systems, filtering routers, and authenticator content. Organizations may employ different mechanisms to achieve confidentiality and integrity protections, including the use of cryptographic mechanisms and file share scanning. Integrity protection can be achieved, for example, by implementing Write-Once-Read-Many (WORM) technologies. Organizations may also employ other security controls including, for example, secure off-line storage in lieu of online storage when adequate protection of information at rest cannot otherwise be achieved and/or continuous monitoring to identify malicious code at rest. Related controls: AC-3, AC-6, CA-7, CM-3, CM-5, CM-6, PE-3, SC-8, SC-13, SI-3, SI-7.</p> <p>References: NIST Special Publications 800-56, 800-57, 800-111.</p>	Require both confidentiality and integrity, not an option.	Ambiguous Specification	Occasional	Moderate	Low
SC-8 (1)	<p>The information system implements cryptographic mechanisms to [Selection (one or more): prevent unauthorized disclosure of information; detect changes to information] during transmission unless otherwise protected by [Assignment: organization-defined alternative physical safeguards].</p> <p>Supplemental Guidance: Encrypting information for transmission protects information from unauthorized disclosure and modification. Cryptographic mechanisms implemented to protect information integrity include, for example, cryptographic hash functions which have common application in digital signatures, checksums, and message authentication codes. Alternative physical security safeguards include, for example, protected distribution systems. Related control: SC-13.</p>	RAS: no protection for protection mechanisms; JD: should not allow OR, needs to be both	Risk to Data	Unlikely	Critical	Low
SC-12 (2)	The organization produces, controls, and distributes symmetric cryptographic keys using [Selection: NIST FIPS-compliant; NSA-approved] key management technology and processes.	Necessary but insufficient. "Controls" is a weak requirement. Allows even intentionally giving to adversary.	Risk to Data	Unlikely	Critical	Low
SC-12 (3)	The organization produces, controls, and distributes asymmetric cryptographic keys using [Selection: NSA-approved key management technology and processes; approved PKI Class 3 certificates or prepositioned keying material; approved PKI Class 3 or Class 4 certificates and hardware security tokens that protect the user's private key].	Does not limit control over who can access keys (giving access to foreign adversary)	Risk to Data	Unlikely	Critical	Low

SI-2	<p>The organization:</p> <ul style="list-style-type: none"> a. Identifies, reports, and corrects information system flaws; b. Tests software and firmware updates related to flaw remediation for effectiveness and potential side effects before installation; c. Installs security-relevant software and firmware updates within [Assignment: organization-defined time period] of the release of the updates; and d. Incorporates flaw remediation into the organizational configuration management process. <p>Supplemental Guidance: Organizations identify information systems affected by announced software flaws including potential vulnerabilities resulting from those flaws, and report this information to designated organizational personnel with information security responsibilities. Security-relevant software updates include, for example, patches, service packs, hot fixes, and anti-virus signatures. Organizations also address flaws discovered during security assessments, continuous monitoring, incident response activities, and system error handling. Organizations take advantage of available resources such as the Common Weakness Enumeration (CWE) or Common Vulnerabilities and Exposures (CVE) databases in remediating flaws discovered in organizational information systems. By incorporating flaw remediation into ongoing configuration management processes, required/anticipated remediation actions can be tracked and verified. Flaw remediation actions that can be tracked and verified include, for example, determining whether organizations follow US-CERT guidance and Information Assurance Vulnerability Alerts. Organization-defined time periods for updating security-relevant software and firmware may vary based on a variety of factors including, for example, the security category of the information system or the criticality of the update (i.e., severity of the vulnerability related to the discovered flaw). Some types of flaw remediation may require more testing than other types. Organizations determine the degree and type of testing needed for the specific type of flaw remediation activity under consideration and also the types of changes that are to be configuration-managed. In some situations, organizations may determine that the testing of software and/or firmware updates is not necessary or practical.</p> <p>for example, when implementing simple anti-virus signature updates. Organizations may also consider in testing decisions, whether security-relevant software or firmware updates are obtained from authorized sources with appropriate digital signatures. Related controls: CA-2, CA-7, CM-3, CM-5, CM-8, MA-2, IR-4, RA-5, SA-10, SA-11, SI-11.</p>	No maximum time period since organization can define.	Ambiguous Specification	Occasional	Moderate	Low
SI-3	<p>The organization:</p> <ul style="list-style-type: none"> a. Employs malicious code protection mechanisms at information system entry and exit points to detect and eradicate malicious code; b. Updates malicious code protection mechanisms whenever new releases are available in accordance with organizational configuration management policy and procedures; c. Configures malicious code protection mechanisms to: <ul style="list-style-type: none"> 1. Perform periodic scans of the information system [Assignment: organization-defined frequency] and real-time scans of files from external sources at [Selection (one or more); endpoint; network entry/exit points] as the files are downloaded, opened, or executed in accordance with organizational security policy; and 2. [Selection (one or more): block malicious code; quarantine malicious code; send alert to administrator; [Assignment: organization-defined action]] in response to malicious code detection; and d. Addresses the receipt of false positives during malicious code detection and eradication and the resulting potential impact on the availability of the information system. <p>Supplemental Guidance: Information system entry and exit points include, for example, firewalls, electronic mail servers, web servers, proxy servers, remote-access servers, workstations, notebook computers, and mobile devices. Malicious code includes, for example, viruses, worms, Trojan horses, and spyware. Malicious code can also be encoded in various formats (e.g., UUENCODE, Unicode), contained within compressed or hidden files, or hidden in files using steganography. Malicious code can be transported by different means including, for example, web accesses, electronic mail, electronic mail attachments, and portable storage devices. Malicious code insertions occur through the exploitation of information system vulnerabilities. Malicious code protection mechanisms include, for example, anti-virus signature definitions and reputation-based technologies. A variety of technologies and methods exist to limit or eliminate the effects of malicious code. Pervasive configuration management and comprehensive software integrity controls may be effective in preventing execution of unauthorized code. In addition to commercial off-the-shelf software, malicious code may also be present in custom-built software. This could include, for example, logic bombs, back doors, and other types of cyber attacks that could affect organizational missions/business functions. Traditional malicious code protection mechanisms cannot always detect such code. In these situations, organizations rely instead on other safeguards including, for example, secure coding practices, configuration management and control, trusted procurement processes, and monitoring practices to help ensure that software does not perform functions other than the functions intended. Organizations may determine that in response to the detection of malicious code, different actions may be warranted. For example, organizations can define actions in response to malicious code detection during periodic scans, actions in response to detection of malicious downloads, and/or actions in response to detection of maliciousness when attempting to open or execute files. Related controls: CM-3, MP-2, SA-4, SA-8, SA-12, SA-13, SC-7, SC-26, SC-44, SI-2, SI-4, SI-7.</p> <p>References: NIST Special Publication 800-83.</p>	RAS: no protection from protection mechanisms WCE: this control often causes issues with developer workstations, as most commercially available solutions will trip over the new software code that devs build and execute on their workstations. There aren't clear directions here for how to use malware protection and actively develop new software on the same environment.	Risk to Data	Unlikely	Critical	Low

SA-4	<p>The organization includes the following requirements, descriptions, and criteria, explicitly or by reference, in the acquisition contract for the information system, system component, or information system service in accordance with applicable federal laws, Executive Orders, directives, policies, regulations, standards, guidelines, and organizational mission/business needs:</p> <p>a. Security functional requirements; b. Security strength requirements; c. Security assurance requirements; d. Security-related documentation requirements; e. Requirements for protecting security-related documentation; f. Description of the information system development environment and environment in which the system is intended to operate; and g. Acceptance criteria.</p> <p>Supplemental Guidance: Information system components are discrete, identifiable information technology assets (e.g., hardware, software, or firmware) that represent the building blocks of an information system. Information system components include commercial information technology products. Security functional requirements include security capabilities, security functions, and security mechanisms. Security strength requirements associated with such capabilities, functions, and mechanisms include degree of correctness, completeness, resistance to direct attack, and resistance to tampering or bypass. Security assurance requirements include: (i) development processes, procedures, practices, and methodologies; and (ii) evidence from development and assessment activities providing grounds for confidence that the required security functionality has been implemented and the required security strength has been achieved. Security documentation requirements address all phases of the system development life cycle.</p> <p>Security functionality, assurance, and documentation requirements are expressed in terms of security controls and control enhancements that have been selected through the tailoring process. The security control tailoring process includes, for example, the specification of parameter values through the use of assignment and selection statements and the specification of platform dependencies and implementation information. Security documentation provides user and administrator guidance regarding the implementation and operation of security controls. The level of detail required in security documentation is based on the security category or classification level of the information system and the degree to which organizations depend on the stated security capability, functions, or mechanisms to meet overall risk response expectations (as defined in the organizational risk management strategy). Security requirements can also include organizationally mandated configuration settings specifying allowed functions, ports, protocols, and services. Acceptance criteria for information systems, information system components, and information system services are defined in the same manner as such criteria for any organizational acquisition or procurement. The Federal Acquisition Regulation (FAR) Section 7.103 contains information security requirements from FISMA. Related controls: CM-6, PL-2, PS-7, SA-3, SA-5, SA-8, SA-11, SA-12.</p>	<p>RAS: for multi-national corps, this needs to explicitly consider access/control from foreign govns</p>	Risk to Data	Unlikely	Critical	Low
SA-10 (1)	<p>The organization requires the developer of the information system, system component, or information system service to enable integrity verification of software and firmware components.</p> <p>Supplemental Guidance: This control enhancement allows organizations to detect unauthorized changes to software and firmware components through the use of tools, techniques, and/or mechanisms provided by developers. Integrity checking mechanisms can also address counterfeiting of software and firmware components. Organizations verify the integrity of software and firmware components, for example, through secure one-way hashes provided by developers. Delivered software and firmware components also include any updates to such components. Related control: SI-7.</p>	<p>RAS: should also consider backdoored updates where adversary co-opts legit update to inject own code; checksum would not detect. should continually monitor communications of security devices</p>	Risk to Data	Unlikely	Critical	Low
SA-4 (2)	<p>The organization requires the developer of the information system, system component, or information system service to provide design and implementation information for the security controls to be employed that includes: [Selection (one or more): security-relevant external system interfaces; high-level design; low-level design; source code or hardware schematics; [Assignment: organization-defined design/implementation information]] at [Assignment: organization-defined level of detail].</p> <p>Supplemental Guidance: Organizations may require different levels of detail in design and implementation documentation for security controls employed in organizational information systems, system components, or information system services based on mission/business requirements, requirements for trustworthiness/resiliency, and requirements for analysis and testing. Information systems can be partitioned into multiple subsystems. Each subsystem within the system can contain one or more modules. The high-level design for the system is expressed in terms of multiple subsystems and the interfaces between subsystems providing security-relevant functionality. The low-level design for the system is expressed in terms of modules with particular emphasis on software and firmware (but not excluding hardware) and the interfaces between modules providing security-relevant functionality. Source code and hardware schematics are typically referred to as the implementation representation of the information system. Related control: SA-5.</p>	<p>RAS: what standardized things are you looking for here and who is trained to look for it? this should not be left up to organizations. tons of previous work about how difficult it is to detect backdoors even with source code</p>	Risk to Data	Seldom	Critical	Medium
All	FedRAMP allows for variable security controls: high, moderate, low	Security controls for US Gov should not be variable	Ambiguous Specification	Frequent	Negligible	Medium
All	FIPS Publications 140-2	140-3 replaced 140-2	Obsolete Reference	Frequent	Negligible	Medium

B.3 Additional data from Chapter 6

B.4 Playbook examples

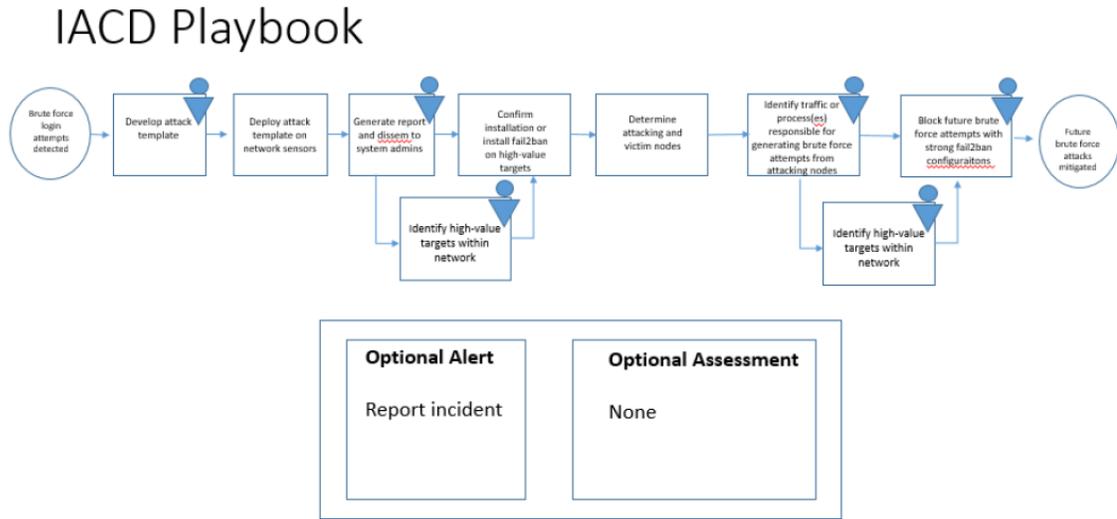


Figure B.4: Participant P4’s IACD playbook for brute force login attempts.

B.5 Additional data from Chapter 7

B.5.1 Quantitative analysis

Factor	Type	Description	Baseline
Category	Fixed	Measure groups	Passive Defense
ID	Random	Participant ID	–

Table B.2: Factors used in creating the Cumulative Link Mixed Model.

Our Cumulative Link Mixed Model took into account fixed and random effects (Table B.2). The fixed effect was the set of all complementary measure categories

Scenario: Detecting and responding to the login of a fake account (honeypot).

<https://attack.mitre.org/techniques/T1078/>

1. Preparation:

Preparation helps speed up response time. In this step, a list of critical assets and critical endpoints associated with a threat is compiled. This list is ranked by level of importance and monitored.

- a. Passwords for “honeypot” fake user account
- b. Logging Services
- c. Alerting mechanisms
 - i. Dashboard
 - ii. Email

2. Detection and Analysis:

Now that the threat incident has been identified, information must be gathered on the threat and an analysis is done.

Where is the entry point of this breach?

- a. Identify the location that the “honeypot” account login originates from by reviewing logs.
- b. Dashboard alerts cannot be removed until after the breach is resolved.

What is the breadth of this breach?

- a. Correlate the time and origin of the “honeypot” account login with other attempted logins.
- b. Determine what access each compromised account had and if failed login attempts also occurred.
 - i. If there were no failed login attempts, this could be indicative that the password used to login was not guessed, but acquired through some other means.
- c. Identify if any data was exfiltrated from the network.

Analyze the threat to the best of your ability. Think about these questions and add anything else you can think of.

-What are the consequences of not resolving this incident?

Figure B.5: A sample of Participant P11’s NIST playbook for credential misuse attempt.

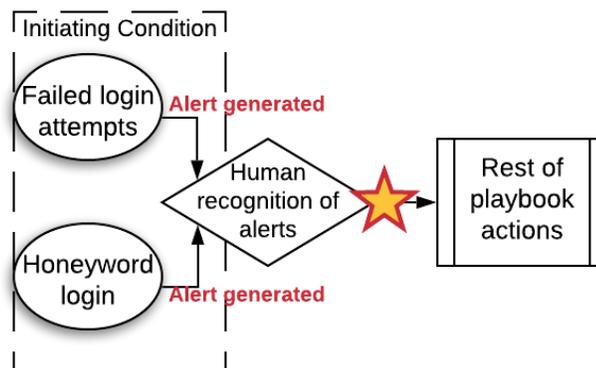


Figure B.6: A visualization of the common issue novice defenders had using playbooks. Despite alerts being present, the defender failed to recognize the significance of the alerts and failed to enact the follow-on steps described in the playbooks.

and the random effect was the participant set, forming the mixed effect input. The random effect includes duplicate identification numbers since different participants could choose the same measure. We compared each proactive control to the baseline that we selected.

Contrast	Estimate	p-value
Passive Def - Continually Evolving Def	2.170	0.0001*
Passive Def - Training and Exercises	1.807	0.0046*
Passive Def - Human-focused Reviews	1.679	0.0068*
Continually Evolving Def - Training and Exercises	-0.363	0.8321
Continually Evolving Def - Human-focused Reviews	-0.491	0.6279
Training and Exercises - Human-focused Reviews	-0.128	0.9934

*Statistically significant

Table B.3: Contrasts and estimates between combinations of proactive control groups

In Table B.3, we provide contrast values between all combinations of control groups that shows the satisfaction level difference estimates between the control groups and the associated p-value that indicates significance.

B.5.2 Reported measures

Table B.4 lists all the complementary measures reported by our participants, organized according to the four high-level categories.

B.5.3 List of reported compliance standards

Table B.5 details the reported compliance standards used by study participants.

B.5.4 Demographics

Table [B.6](#) details the collected demographics of each participant. Clientele size (C/S) indicates the number of supported customers, whereas the number of reported supported organizations (S/O) indicates how many *external* organizations that the company supports.

B.5.5 Codebook

Codebook

Category	Code	Explanation	Example
unaddressed_threat	Untrained_auditors	Concern about how auditors conduct their assessments	"Compliance is a patchwork and it is often interpreted and enforced "tribally" or locally – the person deciding whether you meet the requirements is an outsider who often has no security background. The configuration they approve may be useless for preventing the threat that the control is intended to stop. Many compliance controls are put in place to satisfy the inspectors, with no traceability to mission/business requirements."
unaddressed_threat	Compliance_is_baseline_min_sec	Compliance provides protection guidelines only up to an acceptable threshold	" Frameworks are a baseline to ensure you're thinking about controls in many domains at a minimal/moderate. Very often even the baseline controls are not even implemented well to begin with... Its a starting point not a destination."
unaddressed_threat	Compliance_is_insufficient	Compliance without complementary measures is not enough for protection	"Compliance doesn't fully address any threat."
unaddressed_threat	Rely_on_ext_experts_for_understanding	Relying on external party recommendations for complete protection	"Contract companies evaluate our systems yearly. We add on security based on evaluations."
unaddressed_threat	Compliance_does_not_protect_against_sophisticated_attacks	Compliance provides only basic level protection	"I can't get into specifics, but in general we believe compliance mechanisms only protect against 80% of threats (i.e. the low hanging fruit)."
unaddressed_threat	Rely_on_self-reports_for_employee_non-compliance	Threat and non-compliance mitigation based on employee self-reports	We encourage self-reported cases of employee negligence or failure to adopt and employ standard operating procedures within our company with amnesty.
unaddressed_threat	0day_emerging_threats	Previously unknown threats	"We allow full scope pentests against our network to ID all emerging threats."
unaddressed_threat	Insider_threats	Threats that originate from people within the organization	"Foreign state actors. Insiders."
unaddressed_threat	Phishing	Maliciously obtaining sensitive information in the disguise of a trustworthy party	"Insider Threats, Phishing, Denial of Service attacks."
unaddressed_threat	Nation_state	Highly sophisticated and state-sponsored adversaries	"Nation state actors (APTs)"
unaddressed_threat	DoS	Rendering a resource unavailable for legitimate users	"Insider Threats, Phishing, Denial of Service attacks."
unaddressed_threat	Operational_Security_Threats	Comprehensive analysis of assets and infrastructure from an adversary's perspective	"Operational security threats not covered by compliance frameworks."
what_worked_well	PNTA_No_answer	The participant did not respond.	N/A
what_worked_well	Time_to_implement_or_use	Ability to implement or use a complementary measure in a time-efficient manner	"We struggle to quickly integrate paid vendor intel into our analysis systems."
what_worked_well	Valuable_outcomes	whether the complementary measure provided benefits or not	"Participation was high for this program. Over 100 vulnerabilities were identified."
what_worked_well	Better_security_understanding	Following best practices and an educated approach to security	"The ability to focus resources has allowed us to take a deeper and more methodical approach to our security implementation."
what_worked_well	Reduces_attack_vector	Mitigation of a certain attack types	"Able to get bugs identified and submitted, seems like patches were applied as well."
what_worked_well	Reduce_Cost	Reducing overall security-based spending of the organization	"Able to forecast budgets and resources for future threats."
what_worked_well	Time_to_replace	Time spent to replace an already deployed complementary measure	"Authenticator apps can cause issues when employees break or lose their phones. Often it takes a long time to disable MFA, restore access, and set up again."
what_worked_well	Proper_implementation_/design	Implementing a complementary measure by following its guidelines or in a well-established and planned manner	"It's not done as periodically as the policy says it should be."
what_worked_well	Embracing_Modern_Techniques	Replacing obsolete techniques with modern and up-to-date approaches	"We do it improperly and based off of old policies and intelligence."
what_worked_well	Convincing_Departmental_Entities	Ability to talk departmental entities to embrace or acquire a complementary measure	"Legal liability and permissions are difficult to overcome."
what_worked_well	Scalability	The ability to handle a growing amount of work	"Initial roll outs of two-factor authentication can be time consuming with lots of user misunderstandings and/or errors."

Table 1: Codebook

what_worked_well	Availability	The measure of accessibility	"EDR is one of the best ways to prevent multiple types of attacks from impacting our organization and does not bog down the endpoints resources."
what_worked_well	Good_Management	Skilled, educated, and competent management	"Improper understanding by senior leadership limits its potential. Leadership approaches threat hunting as 'vacuum up everyones data and search it as big data' vs hunting on specific organizations networks."
what_worked_well	Having_Skilled_Employees	Having employees who are skilled and successful at what they do in their daily work routine	"Not convinced we brought in the right hunters."
what_worked_well	Validation	Accuracy checking	"Nice to reassess security gaps frequently"
what_worked_well	Reliability	The degree of trustworthiness	"Being able to get intel, comb the environment and incorporate the new intel in searches has worked well."
what_worked_well	Convincing_Users_or_Employees	Convincing the users and employees to embrace a complementary measure	"Developing the program was great... informing everyone of its existence has been a struggle."
what_worked_well	Sentiment	Overall feeling or opinion regarding a complementary measure	"Being able to get intel, comb the environment and incorporate the new intel in searches has worked well."
why_implement	N/A	The participant responded with "N/A."	Not applicable.
why_implement	Support_customers	To better protect clients or providing the ability to protect themselves better	"As part of our clients services offering, we perform these internally and for clients."
why_implement	Informs_how_to_implement_security	Information on how to properly implement a complementary measure	"Password-based protection alone is insufficient, we needed another layer."
why_implement	Best_practice	Well-known, established, and proven methods	"This is a best practice something you have and something you know."
why_implement	Support_other_defense	Indirectly supporting the security of other parts of the system while providing security to a certain part of the system while	"It is imperative that additional layers of security are implemented where necessary. Multi-factor authentication provides an additional level of security that assists in reducing the occurrence of gaining access to critical systems."
why_implement	Better_than_compliance_rec	The complementary measure provides better security compared to the compliance recommendations	"simple signature based AV is dead. EDR tooling gives a much richer vision of process execution that his valuable for both detection and forensics."
why_implement	Expert_recommendation	Getting advice from external parties	"Third party vendor recommended it to help monitor key systems."
why_implement	Security_incident_or_frequency	Protection for previously happened or highly frequently happening incidents	"We have the most attacked network in the world."
why_implement	Skill_Shortage	The shortage of knowledgeable individuals who are also working in the security field	"Humans do not scale and are in short supply, and security data is growing exponentially. Automation has to be employed."
why_implement	Reduce_costs/Cost-Benefit	Same effectiveness and quality, but cheaper solutions	"We needed to economize and focus our efforts."
why_implement	Support_Employees	Solutions that help employees to be better at tasks that they do in their daily routine	"Threat intelligence can not only educate our organization on new or continuing threats, but can help us understand how they might impact us."
why_implement	Situational_Awareness	Improving the adequate knowledge of the potential threats, the environment, and the organization's mission	"To discuss process/methodologies for response to various potential events."
why_implement	Effort_Prioritization	Prioritizing efforts based on several factors, such as budget limitations	"We needed to economize and focus our efforts."
why_implement	Organization_Req	Organizational requirements or rules that must be followed	"Required as individuals come in and out of organizations."
why_implement	Compliance_is_Insufficient	Compliance without complementary measures is not enough for protection	"AV is just flatly insufficient. attackers often use "living off the land" tools, this helps to detect and prevent normal tools used in bad ways."
why_implement	Vendor_or_App_Reputation	Purchasing and deploying a complementary measure based on previous facts regarding an application or vendor	"We leverage the campus FireEye system. It has good ratings based upon what I have read, it was approved by security, and it is free."
compatible	PTNA	The participant did not respond.	
compatible	N/A	The participant responded with "N/A".	Not applicable.
compatible	Do_not_check	Not checking if the complementary measure is compatible with enforced compliance programs	"We don't bother and I cannot understand why."
compatible	Not_covered_in_compliance	the complementary measure is not covered in the enforced compliance programs	"above and beyond"

Table 2: Codebook - cont.

compatible	Helps_streamline	Making compliance more efficient	"TI can help us understand if the control we have in place will be effective against threats."
compatible	"Bonus_points"	Implementing a complementary control for a better (extra) image	"Most compliance regimes allow for an organization to "get credit" for additional controls in the front matter by describing them in a narrative."
compatible	Out-of-the-box_compliance	Uncommon, but effective ways to make complementary controls compatible with enforced compliance programs	"Selecting a good product that meets compliance requirements in addition to security requirements."
compatible	Supports_some_compliance_but_not_all	A complementary measure satisfying a part of an enforced compliance program	"Most of our compliance regimes expect MFA. We take credit for it wherever it applies."
compatible	Legal_Review	Making sure that a complementary measure is compatible with enforced compliance programs by conducting legal reviews	"Lots of legal review."
compatible	Dictated_by_compliance	Already a requirement of enforced compliance programs	"We defer to the compliance controls, even if the threat model has changed and would better serve us."
compatible	Proper_Design_/Implementation	Making sure that a complementary measure is implemented by following its guidelines or in a well-established and planned manner	"Strict scoping, bounded timelines and user registration."
compatible	Incremental_Improvement	Slowly, but surely improvement over time	"Incremental improvement over time."
compatible	Testing	Testing if a complementary measure is compatible with the enforced compliance programs or not	"PCI evaluations show it was compatible."
compatible	Knowledgeable_Decision Makers	Having skilled and knowledgeable decision-makers that understand how enforced compliance programs work and how to be compatible with them	"By having personnel that understand the compliance standards and the strategy of the organization."
compatible	Identifying_Gaps	Making sure that a complementary measure is compatible with enforced compliance programs by identifying gaps	"We probe whether we have the controls or not, and if we do not, we work to remediate them."
compatible	Keeping_up_With_Tech	Being up-to-date with new developments and technology	"Integration into systems throughout the lifecycle with initial focus on end point authorization and access controls."
compatible	Continuous_Education	Contentious knowledge and awareness improvement	"Continuous education against common threats."
compatible	Documentation	Recording events and incidents	"Document how, where, when an issue was found. Show proof of mitigation."
compatible	Managed_by_third_party	Requirements of an enforced compliance program managed by third-party experts	"This isn't really required, but is useful to have as a way for external parties to report security issues so that they aren't lost in Customer Service queues."
key_factors	No_Response	The participant did not respond.	N/A
key_factors	Directed_security_(internal)	The quality of	"Operations are king. Once it is determined what product is wanted security is brought in to assess, which is slightly backwards."
key_factors	Expert_recommendation	Importance of recommendations from experts	"Third party companies evaluate our systems and make recommendations."
key_factors	Budget	Spending power	"Attempted self assessments and then prioritize bang/buck."
key_factors	Current_threats	Current threats to the organization	"Yes. Threat intelligence, regulations, and quantitative risk assessments server to help prioritize."
key_factors	Business_goals	Importance of the organization's mission	"Business needs, client requests."
key_factors	Inter-departmental_collaboration	The degree of communication and collaboration between different parties in the same organization	"Allowing our red team to do their jobs. they find the gaps, and we fix based on threat levels."
key_factors	Reputable_Vendors	Previous facts regarding a vendor	" We also have to spend a good bit of time finding vendors with truly useful technologies and not just well marketed snake oil."
key_factors	Established_programs	The degree of previous success regarding a program	"Vendor sales pitches and entrenched constituencies."
key_factors	Time_Saving	Being able to spend less time	"Look for low cost and items can easily implement."
key_factors	Automation	Automating repetitive manual labor	"Possibility of automation."
key_factors	Regulatory_Requests	Requests that come from compliance regulators	"Typically this is related to regulatory requests."

Table 3: Codebook - cont.

key_factors	Applicability	Integration difficulty or compatibility	"The applicability to our systems and threat models."
key_factors	Potential_Benefits	Valuable outcomes	"Value/ROI, efficacy, lifecycle and management overhead."
key_factors	User_Feedback	User experiences and feedback on a complementary measure	"Depends on if someone is asking for new tool or explains why they can't see something malicious."
key_factors	Capability_/Features	Feature-rich solutions	"We also have to spend a good bit of time finding vendors with truly useful technologies."
key_factors	Necessity	Current needs for better protecting the organization	"We invest in solutions that meet our highest priority needs."
key_factors	Implementation_Difficulty	The degree of difficulty in implementing a complementary measure	"Look for low cost and items can easily implement."
key_factors	Compliance_Support	Supporting an existing or a part of a compliance	"Cost, threats, whether part of compliance or not."

Table 4: Codebook - cont.

Measure	Prevalence
Training and exercises	
Hands-on training	19
On-the-job mentoring	16
Tabletop	15
Live security	13
On-the-job peering	11
Internal phishing exercises*	1
Human-focused reviews	
Periodic access reviews	21
Change control	18
Threat modeling	16
Integrity review	14
Incident response playbooks*	1
Passive defenses	
Multi-factor authentication	32
Zero clients	9
Sandboxing	9
End-to-end encryption*	1
Microsegmentation*	1
Continually evolving defenses	
Endpoint detection and response	26
Threat hunting	20
Threat intelligence	19
Vulnerability disclosure / bug bounty	14
Physical access barriers	13
Machine learning	10
Dogfooding*	1

Table B.4: All complementary measures reported by our participants, organized into four high-level categories. Measures provided by participants in the ‘other’ field are marked with [^]*

Reported Standard	Count
HIPAA (Health Insurance Portability and Accountability Act)	17
PCI DSS (Payment Card Industry Data Security Standard)	14
FISMA (Federal Information Security Management Act)	12
NIST Cybersecurity Framework	33
IRS Publication 1075	1
NERC CIP (Critical Infrastructure Protection)	2
FedRAMP	10
FERPA (The Family Educational Rights and Privacy Act of 1974)	8
COPPA (Children’s Online Privacy Protection Rule)	4
ISO (International Organization for Standardization)	12
CIS Controls (Center for Internet Security Controls)	8
SOX (Sarbanes-Oxley Act)	2
GDPR (General Data Protection Regulation)	9
CCPA (California Consumer Privacy Act)	4
DoD Instruction 8510.01*	4
Genome data protection guidelines*	1
NY Department of Financial Services Regulation*	1
Trusted Computer System Evaluation Criteria “Orange Book”*	1
U.S. Nuclear Regulatory Commission Standards*	1
University IT standards*	2
Other financial regulations*	1

Table B.5: Reported compliance standards used by study participants. Standards provided by participants in the ‘other’ field are marked with [^]*

ID	Sec	Job	Deg	Size	C/S	S/O	Exp
P01	Consumer services	Compliance/Governance SME	Graduate Deg	151-500	1-500	1-3	15
P02	Consumer services	Management	Graduate Deg	0-50	10k-100k	1-3	2
P03	Consumer services	Security Analyst	Bachelors Deg	151-500	100k+	1-3	10
P04	Education	Compliance/Governance SME	Graduate Deg	1000+	10k-100k	1-3	20
P05	Education	Management	Graduate Deg	1000+	1-500	1-3	4
P06	Education	Management	Graduate Deg	1000+	501-5000	1-3	25
P07	Education	Management	Graduate Deg	1000+	10k-100k	1-3	21
P08	Education	Management	Graduate Deg	0-50	501-5000	1-3	22
P09	Education	Management	Graduate Deg	501-1000	5001-10k	1-3	20
P10	Education	Security Analyst	Bachelors Deg	1000+	501-5000	1-3	5
P11	Financial services	Management	Graduate Deg	1000+	10k-100k	1-3	27
P12	Government	Developer	Bachelors Deg	151-500	501-5000	4-10	10
P13	Government	Developer	Bachelors Deg	1000+	100k+	51+	10
P14	Government	Developer	Graduate Deg	1000+	1-500	1-3	11
P15	Government	Management	Graduate Deg	1000+	10k-100k	51+	14
P16	Government	Management	Bachelors Deg	1000+	100k+	51+	24
P17	Government	Management	Graduate Deg	501-1000	501-5000	1-3	15
P18	Government	Management	Graduate Deg	1000+	100k+	51+	10
P19	Government	Management	Graduate Deg	1000+	100k+	51+	16
P20	Government	Management	Graduate Deg	151-500	10k-100k	11-50	20
P21	Government	Management	Graduate Deg	51-150	1-500	11-50	15
P22	Government	Management	Graduate Deg	1000+	100k+	1-3	21
P23	Government	Management	Graduate Deg	51-150	501-5000	1-3	22
P24	Government	Management	Graduate Deg	1000+	501-5000	1-3	9
P25	Government	Security Analyst	Graduate Deg	1000+	100k+	51+	19
P26	Healthcare	Management	Graduate Deg	501-1000	501-5000	1-3	20
P27	Healthcare	Management	Graduate Deg	1000+	100k+	51+	26
P28	Healthcare	Management	Bachelors Deg	1000+	100k+	4-10	15
P29	Healthcare	Security Analyst	Bachelors Deg	1000+	100k+	1-3	5
P30	Info Tech	Compliance/Governance SME	Bachelors Deg	151-500	10k-100k	51+	8
P31	Info Tech	Compliance/Governance SME	Associates Deg	151-500	1-500	11-50	12
P32	Info Tech	Compliance/Governance SME	Graduate Deg	151-500	5001-10k	1-3	3
P33	Info Tech	Compliance/Governance SME	PNTA	1000+	100k+	51+	25
P34	Info Tech	Developer	Graduate Deg	151-500	10k-100k	1-3	10
P35	Info Tech	Management	Graduate Deg	0-50	100k+	51+	25
P36	Info Tech	Management	Bachelors Deg	0-50	501-5000	4-10	12
P37	Info Tech	Security Analyst	Graduate Deg	151-500	10k-100k	11-50	25
P38	Info Tech	Security Analyst	Bachelors Deg	51-150	501-5000	1-3	10
P39	Info Tech	Security Engineer	Graduate Deg	1000+	100k+	51+	10
P40	Info Tech	Security Engineer	Graduate Deg	0-50	1-500	11-50	20

Table B.6: Participant Demographics. The columns show: participant identifiers, business sector, job role, educational degree, organization’s size of hired employees, clientele size, number of supported organizations, and years of IT and compliance experience.

Bibliography

- [1] Hervé Abdi. 2007. The Kendall rank correlation coefficient. *Encyclopedia of Measurement and Statistics* (2007), 508–510.
- [2] Lillian Ablon and Timothy Bogart. 2017. *Zero Days, Thousands of Nights*. (2017).
- [3] Sherly Abraham and InduShobha Chengalur-Smith. 2010. An overview of social engineering malware: Trends, tactics, and implications. *Technology in Society* 32, 3 (2010), 183–196.
- [4] AdaptiveMobile Security. 2019. Simjacker Technical Paper. (2019). https://simjacker.com/downloads/technicalpapers/AdaptiveMobile_Security_Simjacker_Technical_Paper_v1.01.pdf
- [5] Sushant Agarwal, Simon Steyskal, Franjo Antunovic, and Sabrina Kirrane. 2018. Legislative compliance assessment: framework, model and GDPR instantiation. In *Annual Privacy Forum*. Springer, 131–149.
- [6] Icek Ajzen. 2011. *The theory of planned behaviour: Reactions and reflections*. (2011).
- [7] Hirotugu Akaike. 1974. A new look at the statistical model identification. *IEEE transactions on automatic control* 19, 6 (1974), 716–723.
- [8] Noura Alomar, Primal Wijesekera, Edward Qiu, and Serge Egelman. 2020. “You’ve Got Your Nice List of Bugs, Now What?” Vulnerability Discovery and Management Processes in the Wild. In *Sixteenth Symposium on Usable Privacy and Security ({SOUPS} 2020)*. 319–339.
- [9] Amazon. Summary of the Amazon S3 Service Disruption in the Northern Virginia (US-EAST-1) Region. (????). <https://aws.amazon.com/message/41926/>
- [10] Amazon. 2018. IRS Publication 1075. (2018). <https://aws.amazon.com/compliance/irs-1075/>

- [11] Amazon Web Services. 2018. Compliance and Top Security Threats in the Cloud – Are You Protected? (2018). <https://www.youtube.com/watch?v=Rc55aY0DnMI&feature=youtu.be&t=18m10s>
- [12] Evan E Anderson and Joobin Choobineh. 2008. Enterprise information security strategies. *Computers & security* 27, 1-2 (2008), 22–29.
- [13] Ross Anderson, Chris Barton, Rainer Böhme, Richard Clayton, Michel JG Van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. 2013. Measuring the cost of cybercrime. In *The economics of information security and privacy*. Springer, 265–300.
- [14] Dorine Andrews, Blair Nonnecke, and Jennifer Preece. 2003. Electronic survey methodology: A case study in reaching hard-to-involve Internet users. *International journal of human-computer interaction* 16, 2 (2003), 185–210.
- [15] Andy Applebaum, Shawn Johnson, Michael Limiero, and Michael Smith. 2018. Playbook Oriented Cyber Response. In *2018 National Cyber Summit (NCS)*. IEEE, 8–15.
- [16] Ashish Arora, Rahul Telang, and Hao Xu. 2008. Optimal policy for software vulnerability disclosure. *Management Science* 54, 4 (2008), 642–656.
- [17] Hala Assal and Sonia Chiasson. 2018. Security in the software development lifecycle. (2018).
- [18] John W Atkinson. 1957. Motivational determinants of risk-taking behavior. *Psychological review* 64, 6p1 (1957), 359.
- [19] Albert Bandura. 1993. Perceived self-efficacy in cognitive development and functioning. *Educational psychologist* 28, 2 (1993), 117–148.
- [20] Albert Bandura. 2006. Guide for constructing self-efficacy scales. *Self-efficacy beliefs of adolescents* 5, 307-337 (2006).
- [21] Albert Bandura and Richard H Walters. 1977. *Social learning theory*. Prentice-Hall Englewood Cliffs, NJ.
- [22] Talya Niehaus Bauer and Berrin Erdogan. 1996. Organizational socialization. *APA Handbook of I/O Psychology* 3 (1996), 51–64.
- [23] Adam Beautement, M Angela Sasse, and Mike Wonham. 2009. The compliance budget: managing security behaviour in organisations. In *Proceedings of the 2008 New Security Paradigms Workshop*. ACM, 47–58.
- [24] Cory Bennett and Ariel Tseitlin. 2012. Chaos monkey released into the wild. *Netflix Tech Blog* 30 (2012).

- [25] Betsy Beyer, Chris Jones, Jennifer Petoff, and Niall Richard Murphy. 2016. *Site Reliability Engineering: How Google Runs Production Systems*. ” O’Reilly Media, Inc.”.
- [26] Josh Blum, Simon Booth, Oded Gal, Maxwell Krohn, Karan Lyons, Antonio Marcedone, Mike Maxim, Merry Ember Mou, Jack O’Connor, and Miles Steele. 2020. Zoom End-to-End Encryption Whitepaper. (2020). <https://github.com/zoom/zoom-e2e-whitepaper>
- [27] Michael Blyth. 2009. *Business continuity management: building an effective incident management plan*. John Wiley & Sons.
- [28] Jeff Bollinger, Brandon Enright, and Matthew Valites. 2015. *Crafting the InfoSec Playbook: Security Monitoring and Incident Response Master Plan*. ” O’Reilly Media, Inc.”.
- [29] David Botta, Rodrigo Werlinger, André Gagné, Konstantin Beznosov, Lee Iverson, Sidney Fels, and Brian Fisher. 2007. Towards understanding IT security professionals and their tools. In *Proceedings of the 3rd symposium on Usable privacy and security*. 100–111.
- [30] Cristian Bravo-Lillo, Lorrie Cranor, Saranga Komanduri, Stuart Schechter, and Manya Sleeper. 2014. Harder to ignore? Revisiting pop-up fatigue and approaches to prevent it. In *10th Symposium On Usable Privacy and Security ({SOUPS} 2014)*. 105–111.
- [31] Travis Breaux and Annie Antón. 2008. Analyzing regulatory rules for privacy and security requirements. *IEEE transactions on software engineering* 34, 1 (2008), 5–20.
- [32] Travis D Breaux and Annie I Antón. 2005. Mining rule semantics to understand legislative compliance. In *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*. 51–54.
- [33] Travis D Breaux and David G Gordon. 2013. Regulatory requirements traceability and analysis using semi-formal specifications. In *International working conference on requirements engineering: Foundation for software quality*. Springer, 141–157.
- [34] Travis D Breaux, Matthew W Vail, and Annie I Anton. 2006. Towards regulatory compliance: Extracting rights and obligations to align requirements with regulations. In *14th IEEE International Requirements Engineering Conference (RE’06)*. IEEE, 49–58.
- [35] John Brooke. 2013. SUS: a retrospective. *Journal of usability studies* 8, 2 (2013), 29–40.
- [36] John Brooke and others. 1996. SUS-A quick and dirty usability scale. *Usability evaluation in industry* 189, 194 (1996), 4–7.

- [37] Rune Todnem By. 2005. Organisational change management: A critical review. *Journal of change management* 5, 4 (2005), 369–380.
- [38] Kelly Caine. 2016. Local standards for sample size at CHI. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 981–992.
- [39] Alvaro A Cárdenas, Saurabh Amin, and Shankar Sastry. 2008. Research Challenges for the Security of Control Systems.. In *HotSec*.
- [40] Adam Carlson. 2013. 3 Reasons Anti-Virus Software Alone Is No Longer Enough. (2013). <https://www.lawtechnologytoday.org/2013/03/3-reasons-anti-virus-software-alone-is-no-longer-enough/>
- [41] Belinda Carne, Marcus Kennedy, and Tim Gray. 2012. Crisis resource management in emergency medicine. *Emergency Medicine Australasia* 24, 1 (2012), 7–13.
- [42] Rune Haubo Bojesen Christensen and Per B Brockhoff. 2013. Analysis of sensory ratings data with cumulative link models. *Journal de la Societe Francaise de Statistique & Revue de Statistique Appliquee* 154, 3 (2013), 58–79.
- [43] Anton Chuvakin. 2013. Named: Endpoint Threat Detection & Response. (2013). <https://blogs.gartner.com/anton-chuvakin/2013/07/26/named-endpoint-threat-detection-response/>
- [44] Paul Cichonski, Tom Millar, Tim Grance, and Karen Scarfone. 2012. Computer security incident handling guide. *NIST Special Publication* 800, 61 (2012), 1–147.
- [45] Robert Clark. 2018. Compliance != Security (Except When It Might Be). In *Enigma 2018 (Enigma 2018)*. USENIX Association, Santa Clara, CA. <https://www.usenix.org/node/208142>
- [46] Chris Cleary. DEF CON 19: Operational Use of Offensive Cyber. (????). <https://www.youtube.com/watch?v=1EDCiUyJa2U>
- [47] Jane Cleland-Huang. 2014. How Well Do You Know Your Personae Non Gratae? *IEEE software* 31, 4 (2014), 28–31.
- [48] Gregory M Coates, Kenneth M Hopkinson, Scott R Graham, and Stuart H Kurkowski. 2010. A trust system architecture for SCADA network security. *IEEE Transactions on Power Delivery* 25, 1 (2010), 158–169.
- [49] Carl Colwill. 2009. Human factors in information security: The insider threat—Who can you trust these days? *Information security technical report* 14, 4 (2009), 186–196.
- [50] Gregory Conti and David Raymond. 2017. *On Cyber: Towards an Operational Art for Cyber Conflict*. Kopidion Press.

- [51] Gregory W Corder and Dale I Foreman. 2009. *Nonparametric statistics for non-statisticians: a step-by-step approach*. John Wiley & Sons.
- [52] Anthony H. Cordesman. 1994. Iraq's Military Forces: 1988-1993. (1994). https://csis-prod.s3.amazonaws.com/s3fs-public/legacy_files/files/media/csis/pubs/iraq88-93.pdf
- [53] Lorrie F Cranor. 2008. A framework for reasoning about the human in the loop. (2008).
- [54] CrowdStrike. 2019. 2019 CrowdStrike Global Threat Report: Adversary Trade-craft and The Importance of Speed. (2019). <https://go.crowdstrike.com/rs/281-0BQ-266/images/Report2019GlobalThreatReport.pdf>
- [55] Fred D Davis. 1989. Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS quarterly* (1989), 319–340.
- [56] Richard Dempsey and Jonathan M Chavous. 2013. Commander's intent and concept of operations. *Military Review* 93, 6 (2013), 58–66.
- [57] Tamara Denning, Batya Friedman, and Tadayoshi Kohno. The Security Cards: A Security Threat Brainstorming Toolkit. (????). <http://securitycards.cs.washington.edu/>
- [58] DEPARTMENT OF HOMELAND SECURITY. 2019. DHS Bomb Threat Checklist. (2019). <https://www.cisa.gov/sites/default/files/publications/dhs-bomb-threat-checklist-2014-508.pdf>
- [59] Jacques Dopagne. 2011. The European air traffic management response to volcanic ash crises: Towards institutionalised aviation crisis management. *Journal of business continuity & emergency planning* 5, 2 (2011), 103–117.
- [60] Josiah Dykstra and Celeste Lyn Paul. 2018. Cyber Operations Stress Survey (COSS): Studying fatigue, frustration, and cognitive workload in cybersecurity operations. In *11th {USENIX} Workshop on Cyber Security Experimentation and Test ({CSET} 18)*.
- [61] Josiah ABS Dykstra and Stephen R Orr. 2016. Acting in the unknown: the cynefin framework for managing cybersecurity risk in dynamic decision making. In *Proceedings of the 8th International Conference on Cyber Conflict (CyCon US '16)*. IEEE, 1–6.
- [62] Allen L Edwards. 1957. The social desirability variable in personality assessment and research. (1957).
- [63] Jack E Edwards, Thomas Edwards, Marie D Thomas, Paul Rosenfeld, and Stephanie Booth-Kewley. 1997. *How to conduct organizational surveys: A step-by-step guide*. Sage.

- [64] Dale C Eikmeier. 2004. Center of gravity analysis. *Military Review* 84, 4 (2004), 2–5.
- [65] Steve Elky. 2006. An Introduction to Information System Risk Management. *SANS Institute InfoSec Reading Room* (2006). <https://www.sans.org/reading-room/whitepapers/auditing/introduction-information-system-risk-management-1204>
- [66] Karen Evans and Franklin Reeder. 2010. *A human capital crisis in cybersecurity: Technical proficiency matters*. CSIS.
- [67] Federal Emergency Management Agency. 2020. FEMA Playbooks. (2020). <https://www.fema.gov/media-library/search/playbook#>
- [68] Christina Rosa Filipowski. 2017. *A Qualitative Case Study of Airline Pilot Leadership Behaviors and Practices During Crisis Situations*. Ph.D. Dissertation. Grand Canyon University.
- [69] J Floyd and JR Fowler. 2009. Survey research methods. *Survey Research Methods (4th ed.)*. SAGE Publications, Inc. Thousand Oaks, CA: SAGE Publications, Inc (2009).
- [70] Food and Drug Administration. 2020. FDA Guidance Documents. (October 2020). <https://www.fda.gov/regulatory-information/search-fda-guidance-documents> Accessed: 2020-10-10.
- [71] Donelson R Forsyth. 2008. Self-serving bias. In *International Encyclopedia of the Social Sciences*, William A Darity (Ed.). Vol. 7. Macmillan Reference USA, Detroit.
- [72] Steven M Furnell, Nathan Clarke, Rodrigo Werlinger, Kasia Muldner, Kirstie Hawkey, and Konstantin Beznosov. 2010. Preparation, detection, and analysis: the diagnostic work of IT security incident response. *Information Management & Computer Security* (2010).
- [73] Mirta Galesic and Michael Bosnjak. 2009. Effects of questionnaire length on participation and indicators of response quality in a web survey. *Public opinion quarterly* 73, 2 (2009), 349–360.
- [74] Daniel Geer and John Harthorne. 2002. Penetration testing: A duet. In *Computer Security Applications Conference, 2002. Proceedings. 18th Annual*. IEEE, 185–195.
- [75] General Services Administration. FedRAMP Moderate Security Controls. (????). https://www.fedramp.gov/assets/resources/documents/FedRAMP_Moderate_Security_Controls.xlsx

- [76] Arnab Ghosh, Prashant Kumar Gajar, and Shashikant Rai. 2013. Bring your own device (BYOD): Security risks and mitigating strategies. *Journal of Global Research in Computer Science* 4, 4 (2013), 62–70.
- [77] Paolo Giorgini, Fabio Massacci, John Mylopoulos, and Nicola Zannone. 2005. Modeling security requirements through ownership, permission and delegation. In *Proceedings. 13th IEEE International Conference on Requirements Engineering (RE '05)*. IEEE, 167–176.
- [78] Cristin Flynn Goodwin and J Paul Nicholas. 2013. Developing a National Strategy for Cybersecurity. *Foundation for Security Growth and Innovation* (2013).
- [79] William E Gortney. 2016. *Department of Defense Dictionary of Military and Associated Terms*. Technical Report. Joint Chiefs of Staff, Washington, United States.
- [80] Julia Graham and David Kaye. 2015. *A Risk Management Approach to Business Continuity: Aligning Business Continuity and Corporate Governance*. Rothstein Publishing.
- [81] Paul Grassi, James Fenton, Elaine Newton, Ray Perlner, Andrew Regenscheid, William Burr, and Justin Richer. 2017. NIST Special Publication 800-63B Digital Identity Guidelines Authentication and Lifecycle Management. (2017). <https://pages.nist.gov/800-63-3/sp800-63b.html>
- [82] Varun Grover. 1999. From business reengineering to business process change management: a longitudinal study of trends and practices. *IEEE Transactions on Engineering Management* 46, 1 (1999), 36–46.
- [83] Robert M Groves, Floyd J Fowler, Mick P Couper, James M Lepkowski, Eleanor Singer, Roger Tourangeau, and others. 2009. Survey Methodology. (2009).
- [84] Greg Guest, Arwen Bunce, and Laura Johnson. 2006. How many interviews are enough? An experiment with data saturation and variability. *Field methods* 18, 1 (2006), 59–82.
- [85] Charles Haley, Robin Laney, Jonathan Moffett, and Bashar Nuseibeh. 2008. Security requirements engineering: A framework for representation and analysis. *IEEE Transactions on Software Engineering* 34, 1 (2008), 133–153.
- [86] GM Hardy. 2012. Beyond Continuous Monitoring: Threat Modeling for Real-time Response. *SANS Institute* (2012).
- [87] Wajih Ul Hassan, Shengjian Guo, Ding Li, Zhengzhang Chen, Kangkook Jee, Zhichun Li, and Adam Bates. 2019. Nodoze: Combatting threat alert fatigue with automated provenance triage. In *Network and Distributed Systems Security Symposium*.

- [88] Andrew F Hayes and Klaus Krippendorff. 2007. Answering the call for a standard reliability measure for coding data. *Communication methods and measures* 1, 1 (2007), 77–89. <http://dx.doi.org/10.1080/19312450709336664>
- [89] Bob Hayes and Kathleen Kotwica. 2013. *Business Continuity: Playbook*. Elsevier.
- [90] Donald Hedeker. 2008. Multilevel models for ordinal and nominal variables. In *Handbook of multilevel analysis*. Springer, 237–274.
- [91] Rich Heidorn. 2019. NERC Seeks \$10M Fine for Duke Energy Security Lapses. (Feb 2019). <https://www.rtoinsider.com/nerc-fine-duke-energy-cip-110308/>
- [92] Joan C Henderson. 2008. Managing crises: UK civil aviation, BAA airports and the August 2006 terrorist threat. *Tourism and Hospitality Research* 8, 2 (2008), 125–136.
- [93] Kim Herzig. 2016. Improving software security with stack traces from bug reports. (2016). <https://docs.microsoft.com/en-us/azure/devops/learn/devops-at-microsoft/improving-software-security-stack-traces-bug-reports>
- [94] Andrew Hiles. 2010. *The definitive handbook of business continuity management*. John Wiley & Sons.
- [95] Naoki Hiroshima. 2014. How I lost my \$50,000 Twitter Username. (2014). <https://arstechnica.com/information-technology/2014/01/how-i-lost-my-50000-twitter-username/>
- [96] Allyson L Holbrook, Melanie C Green, and Jon A Krosnick. 2003. Telephone versus face-to-face interviewing of national probability samples with long questionnaires: Comparisons of respondent satisficing and social desirability response bias. *Public opinion quarterly* 67, 1 (2003), 79–125.
- [97] Jesper F Hopstaken, Dimitri Van Der Linden, Arnold B Bakker, and Michiel AJ Kompier. 2015. A multifaceted investigation of the link between mental fatigue and task disengagement. *Psychophysiology* 52, 3 (2015), 305–315.
- [98] Qing Hu, Tamara Dinev, Paul Hart, and Donna Cooke. 2012. Managing employee compliance with information security policies: The critical role of top management and organizational culture. *Decision Sciences* 43, 4 (2012), 615–660.
- [99] C Derrick Huang and Ravi S Behara. 2013. Economics of information security investment in the case of concurrent heterogeneous attacks with budget constraints. *International Journal of Production Economics* 141, 1 (2013), 255–268.

- [100] Larry Hugick and Jonathan Best. 2008. Questionnaire Length. *Encyclopedia of Survey Research Methods* (2008).
- [101] Jez Humble. 2017. Continuous Delivery Sounds Great, but Will It Work Here? *Queue* 15, 6 (2017), 70.
- [102] Jeffrey Hunker and Christian W Probst. 2011. Insiders and Insider Threats-An Overview of Definitions and Mitigation Techniques. *J. Wirel. Mob. Networks Ubiquitous Comput. Dependable Appl.* 2, 1 (2011), 4–27.
- [103] Eric M Hutchins, Michael J Cloppert, and Rohan M Amin. 2011. Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. *Leading Issues in Information Warfare & Security Research* 1, 1 (2011), 80.
- [104] IACD. 2019a. About IACD. (September 2019). <https://www.iacdautomate.org/aboutiacd> Accessed: 2019-09-01.
- [105] IACD. 2019b. IACD Playbook and Workflow Examples. (September 2019). <https://www.iacdautomate.org/playbook-and-workflow-examples> Accessed: 2019-09-01.
- [106] Cybersecurity Insiders. 2018. Crowd Research Partners. *Insider threat 2017* (2018).
- [107] Internal Revenue Service. 1998. Tax Information Security Guidelines For Federal, State and Local Agencies. (1998). <http://www.unclefed.com/ForTaxProfs/irs-drop/1998/pub1075.pdf>
- [108] Internal Revenue Service. 2016. Publication 1075: Tax Information Security Guidelines For Federal, State and Local Agencies. (2016). <https://www.irs.gov/pub/irs-pdf/p1075.pdf>
- [109] International Organization for Standardization. 2016. How to Write Standards: Tips for standards writers. (2016). <https://www.iso.org/publication/PUB100335.html>
- [110] Fehér Dávid János and Nguyen Huu Phuoc Dai. 2018. Security concerns towards security operations centers. In *2018 IEEE 12th International Symposium on Applied Computational Intelligence and Informatics (SACI)*. IEEE, 000273–000278.
- [111] M Eric Johnson and Scott Dynes. 2007. Inadvertent Disclosure-Information Leaks in the Extended Enterprise.. In *WEIS*.
- [112] Ari Juels and Ronald L Rivest. 2013. Honeywords: Making password-cracking detectable. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 145–160.

- [113] Klaus Julisch. 2009. Security compliance: the next frontier in security research. In *Proceedings of the 2008 New Security Paradigms Workshop*. ACM, 71–74.
- [114] Peter Karpati, Andreas L Opdahl, and Guttorm Sindre. 2011. Experimental comparison of misuse case maps with misuse cases and system architecture diagrams for eliciting security vulnerabilities and mitigations. In *Proceedings of the 6th International Conference on Availability, Reliability and Security (ARES '11)*. IEEE, 507–514.
- [115] Nancy Katz, David Lazer, Holly Arrow, and Noshir Contractor. 2004. Network theory and small groups. *Small group research* 35, 3 (2004), 307–332.
- [116] Jason Kick. 2014. *Cyber exercise playbook*. Technical Report. MITRE CORP BEDFORD MA.
- [117] Faris Bugra Kokulu, Ananta Soneji, Tiffany Bao, Yan Shoshitaishvili, Ziming Zhao, Adam Doupé, and Gail-Joon Ahn. 2019. Matched and Mismatched SOCs: A Qualitative Study on Security Operations Center Issues. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. 1955–1970.
- [118] Alice Y Kolb and David A Kolb. 2005. Learning styles and learning spaces: Enhancing experiential learning in higher education. *Academy of management learning & education* 4, 2 (2005), 193–212.
- [119] Brian Krebs. 2008. Cyber incident blamed for nuclear power plant shutdown. *Washington Post*, June 5 (2008), 2008.
- [120] Klaus Krippendorff. 2004. Reliability in Content Analysis. *Human Communication Research* 30, 3 (2004), 411–433. DOI:<http://dx.doi.org/10.1111/j.1468-2958.2004.tb00738.x>
- [121] Katsiaryna Labunets, Fabio Massacci, Federica Paci, and others. 2013. An experimental comparison of two risk-based security methods. In *Proceedings of the 7th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM '13)*. IEEE, 163–172.
- [122] Jintae Lee and Younghwa Lee. 2002. A holistic model of computer abuse within organizations. *Information management & computer security* (2002).
- [123] Paul Legris, John Ingham, and Pierre Collerette. 2003. Why do people use information technology? A critical review of the technology acceptance model. *Information & management* 40, 3 (2003), 191–204.
- [124] Vector Guo Li, Matthew Dunn, Paul Pearce, Damon McCoy, Geoffrey M Voelker, and Stefan Savage. 2019. Reading the tea leaves: A comparative analysis of threat intelligence. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*. 851–867.

- [125] Dimitri Van Der Linden, Ger PJ Keijsers, Paul Eling, and Rachel Van Schaijk. 2005. Work stress and attentional difficulties: An initial study on burnout and cognitive failures. *Work & Stress* 19, 1 (2005), 23–36.
- [126] Matthew Lombard, Jennifer Snyder-Duch, and Cheryl Campanella Bracken. 2002. Content analysis in mass communication: Assessment and reporting of intercoder reliability. *Human communication research* 28, 4 (2002), 587–604.
- [127] Johnny Long, Bill Gardner, and Justin Brown. 2011. *Google hacking for penetration testers*. Vol. 2. Elsevier.
- [128] Mass Soldal Lund, Bjørnar Solhaug, and Ketil Stølen. 2010. *Model-driven risk analysis: the CORAS approach*. Springer Science & Business Media.
- [129] Steve Mansfield-Devine. 2015. The Ashley Madison affair. *Network Security* 2015, 9 (2015), 8–16.
- [130] Bill Marczak and John Scott-Railton. 2020. Move Fast and Roll Your Own Crypto: A Quick Look at the Confidentiality of Zoom Meetings. (2020). <https://citizenlab.ca/2020/04/move-fast-roll-your-own-crypto-a-quick-look-at-the-confidentiality-of-zoom-meetings/>
- [131] Fabio Massacci and Federica Paci. 2012. How to select a security requirements method? a comparative study with students and practitioners. *Secure IT Systems* (2012), 89–104.
- [132] Christopher J May, Joshua Hammerstein, Jeffrey Mattson, and Kristopher Rush. 2006. Defense in Depth: Foundations for Secure and Resilient IT Enterprises.
- [133] Gary McGraw and Brian Chess. 2009. The Building Security in Maturity Model ({BSIMM}). (2009).
- [134] Daniel Mellado, Eduardo Fernández-Medina, and Mario Piattini. 2006. Applying a security requirements engineering process. *Computer Security—ESORICS 2006* (2006), 192–206.
- [135] Ola Aleksandra Michalec, Dirk van der Linden, Sveta Milyaeva, and Awais Rashid. 2020. Industry Responses to the European Directive on Security of Network and Information Systems (NIS): Understanding policy implementation practices across critical infrastructures. In *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*. USENIX Association, 301–317. <https://www.usenix.org/conference/soups2020/presentation/michalec>
- [136] Microsoft Corporation. 2005. *The STRIDE Threat Model*. Technical Report. Microsoft Corporation. [https://msdn.microsoft.com/en-us/library/ee823878\(v=cs.20\).aspx](https://msdn.microsoft.com/en-us/library/ee823878(v=cs.20).aspx)

- [137] Microsoft Corporation. 2016. *Microsoft Threat Modeling Tool 2016*. Technical Report. Microsoft Corporation. <https://www.microsoft.com/en-us/download/details.aspx?id=49168>
- [138] Devesh Mishra. 2018. Cybersecurity Playbook-An Executive Response. *Available at SSRN 3240285* (2018).
- [139] MITRE. 2019a. ATT&ACK. (2019). <https://attack.mitre.org>
- [140] MITRE. 2019b. Brute Force. (2019). <https://attack.mitre.org/techniques/T1110/>
- [141] MITRE. 2019c. Spearphishing Link. (2019). <https://attack.mitre.org/techniques/T1192/>
- [142] MITRE. 2019d. Valid Accounts. (2019). <https://attack.mitre.org/techniques/T1078/>
- [143] MITRE Corporation. 2018. Personal communication. (2018).
- [144] Naomi Miyake. 1986. Constructive interaction and the iterative process of understanding. *Cognitive science* 10, 2 (1986), 151–177.
- [145] Daniel L Moody. 2003. The method evaluation model: a theoretical model for validating information systems design methods. *Proceedings of the 11th European Conference on Information Systems* (2003), 1327–1336.
- [146] Haralambos Mouratidis, Paolo Giorgini, and Gordon Manson. 2003. Integrating security and systems engineering: Towards the modelling of secure information systems. In *Proceedings of the 15th International Conference on Advanced Information Systems Engineering (CAISE '03)*. Springer, 63–78.
- [147] Michael Muckin and Scott C Fitch. 2014. A Threat-Driven Approach to Cyber Security. *Lockheed Martin Corporation* (2014).
- [148] Paula Murrain-Hill, C Norman Coleman, John L Hick, Irwin Redlener, David M Weinstock, John F Koerner, Delaine Black, Melissa Sanders, Judith L Bader, Joseph Forsha, and others. 2011. Medical response to a nuclear detonation: creating a playbook for state and local planners and responders. *Disaster medicine and public health preparedness* 5, S1 (2011), S89–S97.
- [149] Shea Nangle. 2019. Private Communication. (Feb 2019).
- [150] National Institute of Standards and Technology. 2014. NIST Cybersecurity Framework. (2014). <https://www.us-cert.gov/ccubedvp/cybersecurity-framework>
- [151] National Institute of Standards and Technology. 2017. NIST Special Publication 800-53. (2017). <https://nvd.nist.gov/800-53>

- [152] National Institute of Standards and Technology. 2018. Personal communication. (2018).
- [153] National Security Agency. 2019. NSA Cybersecurity Advisory: Malicious Cyber Actors Leveraging VPN Vulnerabilities for Attack; Check VPN Products for Upgrade. (2019). <https://www.nsa.gov/News-Features/News-Stories/Article-View/Article/1982939/nsa-cybersecurity-advisory-malicious-cyber-actors-leveraging-vpn-vulnerabilities/>
- [154] National Security Agency Information Assurance Directorate. 2015. NSA Methodology for Adversary Obstruction. (2015). <http://www.cdse.edu/documents/cdse/nsa-methodology-for-adversary-obstruction.pdf>
- [155] Rahul Neware, Urmila Shrawankar, Pranay Mangulkar, and Sushil Khune. 2020. Review on Multi-Factor Authentication (MFA) Sources and Operation Challenges. *International Journal of Smart Security Technologies (IJSST)* 7, 2 (2020), 62–76.
- [156] Michael Nieves, Kelley Dempsey, and Victoria Yan Pillitteri. 2017. NIST Special Publication 800-12: An Introduction to Information Security. (2017).
- [157] Jakob Nielsen. 2001. Usability metrics. (July 2001). <https://www.nngroup.com/articles/usability-metrics/> Accessed: 2017-09-01.
- [158] North American Electric Reliability Corporation. 2012. NERC Sanction Guidelines. (2012). https://www.nerc.com/FilingsOrders/us/RuleOfProcedureDL/Appendix_4B_SanctionGuidelines_20121220.pdf
- [159] North American Electric Reliability Corporation. 2014. CIP-007-6 — Cyber Security – Systems Security Management. (2014).
- [160] North American Electric Reliability Corporation. 2018. Critical Infrastructure Protection Committee. (2018). <http://www.nerc.com/comm/CIPC/Related%20Files%20DL/CIPC%20Roster%20as%20of%20February%202018.pdf>
- [161] NYC DoITT. 2017a. CityNet. (2017). <https://www1.nyc.gov/site/doitt/agencies/citynet.page>
- [162] NYC DoITT. 2017b. Cybersecurity Requirements for Vendors & Contractors. (2017). <https://www1.nyc.gov/site/doitt/business/it-security-requirements-vendors-contractors.page>
- [163] N Ochoa. 2010. Pass-the-hash toolkit for windows implementation & use. *Retrieved January 1* (2010).
- [164] National Institute of Standards and Technology. 2019. SP 800-53 Rev. 5 (DRAFT) Security and Privacy Controls for Information Systems and Organizations. *Special Publications* (2019). <https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/draft>

- [165] Cyril Onwubiko, Karim Ouazzane, and others. 2019. SOTER: a playbook for cyber security incident management. *IEEE Transaction of Engineering and Management* (2019), 1–22.
- [166] Andreas L Opdahl and Guttorm Sindre. 2009. Experimental comparison of attack trees and misuse cases for security threat identification. *Information and Software Technology* 51, 5 (2009), 916–932.
- [167] Martin T Orne. 1962. On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American psychologist* 17, 11 (1962), 776.
- [168] Danny Palmer. 2020. Hackers are scanning for vulnerable VPNs in order to launch attacks against remote workers. (2020). <https://www.zdnet.com/article/hackers-are-scanning-for-vulnerable-vpns-in-order-to-launch-attacks-against-remote-workers/>
- [169] PCI Security Standards Council. 2016. Payment Card Industry Data Security Standard: Requirements and Security Assessment Procedures v3.2. (2016). https://www.pcisecuritystandards.org/pci_security/
- [170] Andrew Peterson. 2013. *Cracking Security Misconceptions*. O’Reilly Media, Inc.
- [171] Colin Potts. 1993. Software-engineering research revisited. *IEEE software* 10, 5 (1993), 19–28.
- [172] Petri Puhakainen and Mikko Siponen. 2010. Improving employees’ compliance through information systems security training: an action research study. *Mis Quarterly* (2010), 757–778.
- [173] R. Duian O. Alrawi E. Asdar V. Zhu Y. Kwon B. Saltaformaggio R. Kasturi, Y. Sun. 2020. TARDIS: Rolling Back The Clock On CMS-Targeting Cyber Attacks. In *Proceedings of the IEEE Symposium on Security and Privacy 2020*. IEEE.
- [174] Sophia Rabe-Hesketh, Anders Skrondal, and Andrew Pickles. 2002. Reliable estimation of generalized linear mixed models using adaptive quadrature. *The Stata Journal* 2, 1 (2002), 1–21.
- [175] Martin Reznek, Rebecca Smith-Coggins, Steven Howard, Kanthi Kiran, Phillip Harter, Yasser Sowb, David Gaba, and Thomas Krummel. 2003. Emergency Medicine Crisis Resource Management (EMCRM): Pilot study of a simulation-based crisis management course for emergency medicine. *Academic Emergency Medicine* 10, 4 (2003), 386–389.
- [176] Carl Sabottke, Octavian Suci, and Tudor Dumitras. 2015. Vulnerability Disclosure in the Age of Social Media: Exploiting Twitter for Predicting

- Real-World Exploits.. In *Proceedings of the 24th USENIX Security Symposium (USENIX Security '15)*. 1041–1056.
- [177] Chris Salter, O. Sami Saydjari, Bruce Schneier, and Jim Wallner. 1998. Toward a Secure System Engineering Methodolgy. In *Proceedings of the 1998 Workshop on New Security Paradigms (NSPW '98)*. ACM, New York, NY, USA, 2–10. DOI:<http://dx.doi.org/10.1145/310889.310900>
- [178] Edgar H Schein. 2010. *Organizational culture and leadership*. Vol. 2. John Wiley & Sons.
- [179] Bruce Schneier. 1999. Attack trees. *Dr. Dobbs'??s journal* 24, 12 (1999), 21–29.
- [180] Mathew J. Schwartz. 2019. Ransomware Victims Who Pay Cough Up \$6,733. (Feb 2019). <https://www.bankinfosecurity.com/ransomware-victims-who-pay-cough-up-6733-on-average-a-11994>
- [181] Kevin D. Scott. 2017. Joint Planning. *Joint Publication 5-0* (2017).
- [182] Scott Shackelford. 2017. Exploring the “Shared Responsibility” of Cyber Peace: Should Cybersecurity Be a Human Right? *Kelley School of Business Research paper* (2017), 17–55. <https://ssrn.com/abstract=3005062>
- [183] Scott J Shackelford. 2016. Protecting intellectual property and privacy in the digital age: The use of national cybersecurity strategies to mitigate cyber risk. *Chapman Law Review* 19 (2016), 445.
- [184] Adam Shostack. 2014. *Threat modeling: Designing for security*. John Wiley & Sons.
- [185] Guttorm Sindre and Andreas L Opdahl. 2005. Eliciting security requirements with misuse cases. *Requirements engineering* 10, 1 (2005), 34–44.
- [186] Matthew Smith, Martin Strohmeier, Jonathan Harman, Vincent Lenders, and Ivan Martinovic. 2020. A view from the cockpit: exploring pilot reactions to attacks on avionic systems. (2020).
- [187] Wes Sonnenreich, Jason Albanese, Bruce Stout, and others. 2006. Return on security investment (ROSI)-a practical quantitative model. *Journal of Research and practice in Information Technology* 38, 1 (2006), 45.
- [188] Jason Stamp, John Dillinger, William Young, and Jennifer DePoy. 2003. Common vulnerabilities in critical infrastructure control systems. *SAND2003-1772C. Sandia National Laboratories* (2003).
- [189] Rock Stevens. 2017a. Calcifying Crisis Readiness. *USENIX LISA* (2017).

- [190] Rock Stevens. 2017b. Identifying self-inflicted vulnerabilities: The operational implications of technology within US combat systems. In *2017 International Conference on Cyber Conflict (CyCon US)*. IEEE, 112–118.
- [191] Rock Stevens, Colin Ahern, Daniel Votipka, Elissa Redmiles, Patrick Sweeney, and Michelle Mazurek. 2018. The Battle for New York: A Case Study of Applied Digital Threat Modeling at the Enterprise Level. In *27th USENIX Security Symposium*. USENIX Association.
- [192] Rock Stevens, Josiah Dykstra, Wendy Knox Everette, James Chapman, Garrett Bladow, Alexander Farmer, Kevin Halliday, and Michelle L Mazurek. 2020a. Compliance Cautions: Investigating Security Issues Associated with US Digital-Security Standards. *Network and Distributed System Security Symposium* (2020). <https://www.ndss-symposium.org/wp-content/uploads/2020/02/24003-paper.pdf>
- [193] Rock Stevens, Josiah Dykstra, Wendy Knox-Everette, and Michelle L Mazurek. 2020b. How to Hack Compliance: Using Lessons Learned to Repeatably Audit Compliance Programs for Digital Security Concerns. *Learning from Authoritative Security Experiment Results (LASER)* (2020).
- [194] Rock Stevens, Josiah Dykstra, Wendy Knox-Everette, and Michelle L Mazurek. 2020c. It Lurks Within: A Look at the Unexpected Security Implications of Compliance Programs. *IEEE Security & Privacy (In Draft)* (2020).
- [195] Rock Stevens, Daniel Votipka, Elissa M Redmiles, Colin Ahern, and Michelle L Mazurek. 2019. Applied Digital Threat Modeling: It Works. *IEEE Security & Privacy* 17, 4 (2019), 35–42.
- [196] Joe Strange. 1996. *Centers of Gravity & Critical Vulnerabilities: Building on the Clausewitzian Foundation So That We Can All Speak the Same Language*. Technical Report. MARINE CORPS WAR COLLEGE, QUANTICO, VA.
- [197] Joe Strange and Richard Iron. 2005. *Understanding centers of gravity and critical vulnerabilities*. Department of War Studies, Swedish National Defence College.
- [198] Joe Strange, Richard Iron, and UK Army. 2004. Part 2: The CG-CC-CR-CV Construct: A Useful Tool to Understand and Analyze the Relationship Between Centers of Gravity and their Critical Vulnerabilities. *Understanding Centers of Gravity and Critical Vulnerabilities* (2004).
- [199] Anselm Strauss, Juliet Corbin, and others. 1990. *Basics of qualitative research*. Vol. 15. Newbury Park, CA: Sage.
- [200] Blake E Strom, Andy Applebaum, Douglas P Miller, Kathryn C Nickels, Adam G Pennington, and Cody B Thomas. 2018. MITRE ATT&CK: Design and philosophy. *Technical report* (2018).

- [201] Subashini Subashini and Veeraruna Kavitha. 2011. A survey on security issues in service delivery models of cloud computing. *Journal of network and computer applications* 34, 1 (2011), 1–11.
- [202] Sathya Chandran Sundaramurthy, Alexandru G Bardas, Jacob Case, Xinming Ou, Michael Wesch, John McHugh, and S Raj Rajagopalan. 2015. A human capital model for mitigating security analyst burnout. In *Eleventh Symposium On Usable Privacy and Security ({SOUPS} 2015)*. 347–359.
- [203] Sathya Chandran Sundaramurthy, John McHugh, Xinming Ou, Michael Wesch, Alexandru G Bardas, and S Raj Rajagopalan. 2016. Turning contradictions into innovations or: How we learned to stop whining and improve security operations. In *Twelfth Symposium on Usable Privacy and Security ({SOUPS} 2016)*. 237–251.
- [204] Peter P Swire. 2004. A model for when disclosure helps security: What is different about computer and network security. *J. on Telecomm. & High Tech. L.* 3 (2004), 163.
- [205] Leona Tam, Myron Glassman, and Mark Vandenwauver. 2010. The psychology of password management: a tradeoff between security and convenience. *Behaviour & Information Technology* 29, 3 (2010), 233–244.
- [206] Andrew S Tanenbaum and David J Wetherall. 2011. *Computer networks*. Pearson.
- [207] The MITRE Corporation. 2020. MITRE ATT&CK Mitigations. (2020). <https://attack.mitre.org/mitigations/enterprise/>
- [208] Marianthi Theoharidou, Spyros Kokolakis, Maria Karyda, and Evangelos Kiountouzis. 2005. The insider threat to information systems and the effectiveness of ISO17799. *Computers & Security* 24, 6 (2005), 472–484.
- [209] Tyler W Thomas, Madiha Tabassum, Bill Chu, and Heather Lipford. 2018. Security During Application Development: an Application Security Expert Perspective. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 262.
- [210] Jennifer S Tiffany. 2006. Respondent-driven sampling in participatory research contexts: Participant-driven recruitment. *Journal of Urban Health* 83, 1 (2006), 113–124.
- [211] Steven Tom, Dale Christiansen, and Dan Berrett. 2008. Recommended practice for patch management of control systems. *DHS control system security program (CSSP) Recommended Practice* (2008).
- [212] Roger Tourangeau and Ting Yan. 2007. Sensitive questions in surveys. *Psychological bulletin* 133, 5 (2007), 859.

- [213] United States Computer Emergency Readiness Team. 2018. Personal communication. (2018).
- [214] United States Federal Reserve. 2017. The Federal Reserve Payments Study: 2017 Annual Supplement. (2017). <https://www.federalreserve.gov/paymentsystems/2017-December-The-Federal-Reserve-Payments-Study.htm>
- [215] U.S. Department of Health and Human Services. 2018. IRS Safeguards and Publication 1075 Update. (2018). <https://www.acf.hhs.gov/css/resource/irs-safeguards-and-publication-1075-update>
- [216] U.S. Department of the Army. 1998. Field Manual 100-14 Risk Management. (1998).
- [217] U.S. National Security Agency. 2018. NSA’s Top Ten Cybersecurity Mitigation Strategies. (2018). <https://apps.nsa.gov/iaarchive/library/ia-guidance/security-tips/nsas-top-ten-cybersecurity-mitigation-strategies.cfm>
- [218] Mojtaba Vaismoradi, Hannele Turunen, and Terese Bondas. 2013. Content analysis and thematic analysis: Implications for conducting a qualitative descriptive study. *Nursing & health sciences* 15, 3 (2013), 398–405.
- [219] Milan Vego. 1999. On Operational Art. *Newport, RI: Naval War College* 267 (1999).
- [220] Nikos Virvilis, Dimitris Gritzalis, and Theodoros Apostolopoulos. 2013. Trusted Computing vs. Advanced Persistent Threats: Can a defender win this game?. In *Ubiquitous intelligence and computing, 2013 IEEE 10th international conference on and 10th international conference on autonomic and trusted computing (uic/atc)*. IEEE, 396–403.
- [221] Gretchen R Vogelgesang and Paul B Lester. 2009. How leaders can get results by laying it on the line. (2009).
- [222] Carl Von Clausewitz and James John Graham. 1873. *On war*. Vol. 1. London, N. Trübner & Company.
- [223] Daniel Votipka, Rock Stevens, Elissa Redmiles, Jeremy Hu, and Michelle Mazurek. 2018. Hackers vs. testers: A comparison of software vulnerability discovery processes. In *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 374–391.
- [224] Andy Wapling and Chloe Sellwood. 2016. *Health Emergency Preparedness and Response*. CABI.

- [225] Gregory B White, Glenn Dietrich, and Tim Goles. 2004. Cyber security exercises: testing an organization's ability to prevent, detect, and respond to cyber security events. In *37th Annual Hawaii International Conference on System Sciences, 2004. Proceedings of the*. IEEE, 10–pp.
- [226] Frank Wilcoxon. 1945. Individual comparisons by ranking methods. *Biometrics bulletin* 1, 6 (1945), 80–83.
- [227] Rick Wilson, Patrick Lynett, Kevin Miller, Amanda Admire, Aykut Ayca, Edward Curtis, Lori Dengler, Michael Hornick, Troy Nicolini, and Drew Peterson. 2016. Maritime tsunami response playbooks: background information and guidance for response and hazard mitigation use. *California Geological Survey Special Report* 241 (2016), 48.
- [228] Tarun Yadav and Arvind Mallari Rao. 2015. Technical aspects of cyber kill chain. In *International Symposium on Security in Computing and Communication*. Springer, 438–452.
- [229] Dale E Zand. 1972. Trust and managerial problem solving. *Administrative science quarterly* (1972), 229–239.
- [230] Xueqing Zhang. 2005. Critical success factors for public–private partnerships in infrastructure development. *Journal of construction engineering and management* 131, 1 (2005), 3–14.