DETERMINING SYSTEM REQUIREMENTS FOR HUMAN-MACHINE INTEGRATION IN CYBER SECURITY INCIDENT RESPONSE

by

Megan M. Nyre-Yu

A Dissertation

Submitted to the Faculty of Purdue University

In Partial Fulfillment of the Requirements for the degree of

Doctor of Philosophy

School of Industrial Engineering

West Lafayette, Indiana

December 2019

THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF COMMITTEE APPROVAL

Dr. Barrett S. Caldwell, Chair
School of Industrial Engineering
Dr. C. Robert Kenley
School of Industrial Engineering
Dr. Natalie Lambert
Brian Lamb School of Communication
Dr. Kelly Sprehn
The Charles Stark Draper Laboratory, Inc.

Approved by:
Dr. Abhijit Deshmukh
Head of the Graduate Program

For my family

ACKNOWLEDGMENTS

I would like to acknowledge several people who were instrumental in completing this dissertation. First, I would like to acknowledge my committee members, who went above and beyond in terms of feedback frequency. You each understood my goal and timeline, and worked with me to get it complete with a high level of quality. I deeply respect each of you, and appreciate your particular feedback styles more than you know. Thank you.

To my labmates, Siobhan Heiden, Michelle (Shelly) Jahn Holbrook, and Jordan Hill, I would like to thank each of you for your individual advice and team support. You all together made my GROUPER Lab experience one I will not forget. I am thankful to have found a friend in each of you as well as a colleague.

To Pat and Priyanka Brunese, you both know I would not be in graduate school without the nudge from Pat, and I would not have finished without the encouragement, support, and extroverted brainstorming sessions with Priyanka. You opened your home and your hearts to me, and made me a better person, and academic, for it.

To my family, Janine, David, Lauren, Patrick, Eryn, Grant, Ian and Katharina, (and all the rest!) for the love and support you endlessly provided over the last 5 years during graduate school. You put up with me through all phases, and still came out the other side rooting for me. Thank you for honoring my goals and helping me achieve them.

Finally, to my husband Andy, who I met during graduate school but he never once doubted my capabilities. Thank you for supporting me in every possible way, and never second guessing me, even when I did myself. Thank you for your unwavering love and friendship.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

IT –	Information Technology
IoT –	Internet of Things
CSIRT –	Computer Security Incident Response Team
SOC –	Security Operations Center
SE –	Systems Engineering
KSAs –	Knowledge, skills, and abilities
SOAR –	Security Orchestration Automation and Response
FA –	Function Allocation
HAI –	Human-automation Interaction
HCI –	Human Computer Interaction
LOA –	Levels of Automation
DFA –	Dynamic Function Allocation
CSA –	Cyber Situation Awareness
MTS –	Multi-Team System
CISO –	Chief Information Security Officer
NICE –	National Initiative for Cybersecurity Education
VSM –	Viable Systems Model
CTA –	Cognitive Task Analysis
IR –	Incident Response
CONOPS –	Concept of Operations
T1 / T2 –	Tier 1 / Tier 2
ACTA –	Applied Cognitive Task Analysis
IRB –	Institutional Review Board
CIO –	Chief Information Officer
AAR –	After Action Review
SIEM –	Security Information and Event Management
CDT –	Cognitive Demands Table
IRR –	Inter-Rater Reliability
XAI –	Explainable Artificial Intelligence

ABSTRACT

Author: Nyre-Yu, Megan, M. PhD
Institution: Purdue University
Degree Received: December 2019
Title: Determining System Requirements for Human-Machine Integration in Cyber Security Incident Response
Committee Chair: Barrett S. Caldwell

In 2019, cyber security is considered one of the most significant threats to the global economy and national security. Top U.S. agencies have acknowledged this fact, and provided direction regarding strategic priorities and future initiatives within the domain. However, there is still a lack of basic understanding of factors that impact complexity, scope, and effectiveness of cyber defense efforts. Computer security incident response is the short-term process of detecting, identifying, mitigating, and resolving a potential security threat to a network. These activities are typically conducted in computer security incident response teams (CSIRTs) comprised of human analysts that are organized into hierarchical tiers and work closely with many different computational tools and programs. Despite the fact that CSIRTs often provide the first line of defense to a network, there is currently a substantial global skills shortage of analysts to fill open positions. Research and development efforts from educational and technological perspectives have been independently ineffective at addressing this shortage due to time lags in meeting demand and associated costs. This dissertation explored how to combine the two approaches by considering how human-centered research can inform development of computational solutions toward augmenting human analyst capabilities. The larger goal of combining these approaches is to effectively complement human expertise with technological capability to alleviate pressures from the skills shortage.

Insights and design recommendations for hybrid systems to advance the current state of security automation were developed through three studies. The first study was an ethnographic field study which focused on collecting and analyzing contextual data from three diverse CSIRTs from different sectors; the scope extended beyond individual incident response tasks to include aspects of organization and information sharing within teams. Analysis revealed larger design implications regarding collaboration and coordination in different team environments, as well as considerations about usefulness and adoption of automation. The second study was a cognitive task analysis with

CSIR experts with diverse backgrounds; the interviews focused on expertise requirements for information sharing tasks in CSIRTs. Outputs utilized a dimensional expertise construct to identify and prioritize potential expertise areas for augmentation with automated tools and features. Study 3 included a market analysis of current automation platforms based on the expertise areas identified in Study 2, and used Systems Engineering methodologies to develop concepts and functional architectures for future system (and feature) development.

Findings of all three studies support future directions for hybrid automation development in CSIR by identifying social and organizational factors beyond traditional tool design in security that supports human-systems integration. Additionally, this dissertation delivered functional considerations for automated technology that can augment human capabilities in incident response; these functions support better information sharing between humans and between humans and technological systems. By pursuing human-systems integration in CSIR, research can help alleviate the skills shortage by identifying where automation can dynamically assist with information sharing and expertise development. Future research can expand upon the expertise framework developed for CSIR and extend the application of proposed augmenting functions in other domains.

CHAPTER 1. INTRODUCTION

1.1. Problem Statement

Cyber security is becoming one of the most prevalent concerns across multiple levels of government, sectors of industry, and academia. Recent statements from heads of top U.S. agencies have highlighted cyber defense as a major weakness, and have provided direction regarding strategic priorities and future initiatives to improve it (Coats, 2017; Oltsik, 2018b; Papadopoulous, 2017; Pomerleau, 2016; L. C. Williams, 2017). Integral commercial information technology (IT) providers are also in agreement that cyber defense is a top priority. Industry leaders and experts commonly identify rising cyber crime and the associated costs as reasons to increase focus and spending in cyber security (Clark, Berson, & Lin, 2014; Morgan, 2015; Ponemon Institute, 2017).

Multiple forces are influencing the importance of security from different angles. As mentioned, cyber crime is on the rise. As shown in Figure 1, the Identity Theft Resource Center (2019) reported that the trend of cyber crime has steadily increased in the last decade across multiple industries (e.g. banking, business, education, government, and healthcare). Though the number of reported breaches show some improvement in 2019, they still remain a concern (Darktrace, 2018) as the complexity and sophistication of attacks is on the rise with the evolution of self-propagating threats (Cisco Systems, 2018). The attack surface, or the number of potential ways a system can be infiltrated, is also quickly growing as more people connect to the internet (Boden, 2016; Clark et al., 2014) and the number of connected (and potentially unprotected) devices increases (Darktrace, 2018; Symantec, 2017). In 2017, Darktrace noted a 400% increase in security incidents related to internet of things (IoT); this number seems daunting when combined with the expected increase in IoT devices to reach 20.4 billion by 2020 (Darktrace, 2018).

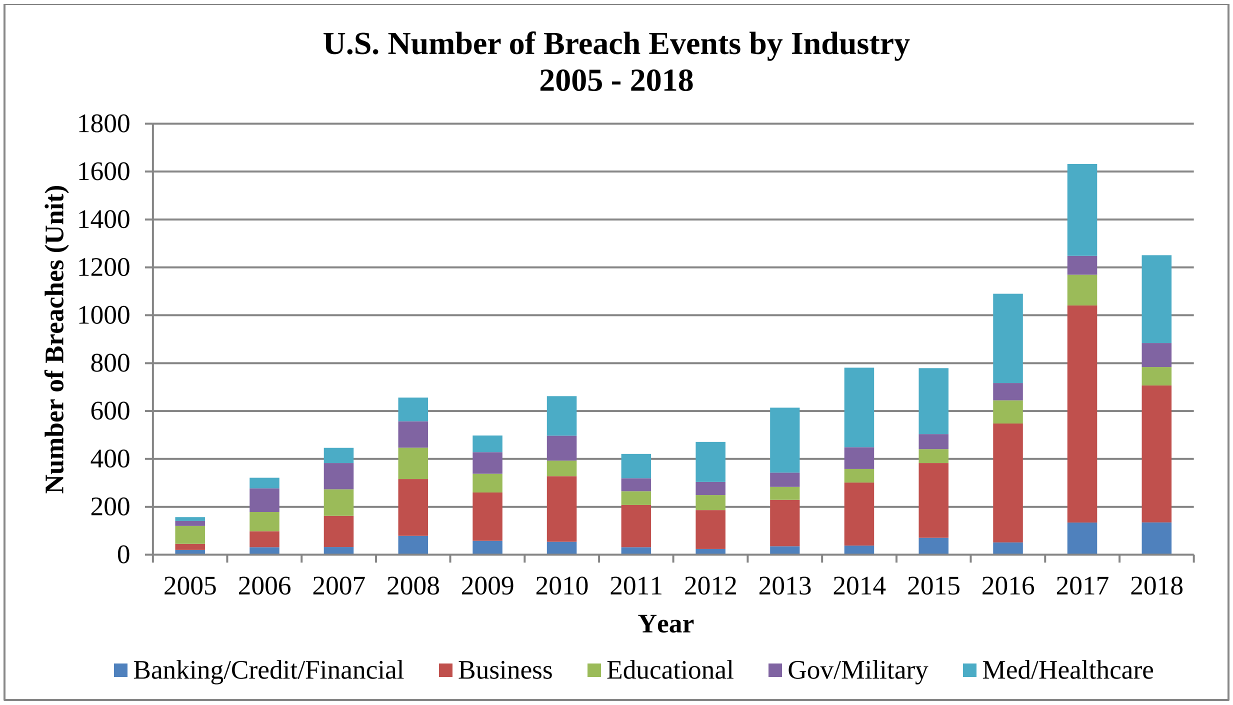

Figure.1. U.S. Reported Breaches (Identity Theft Resource Center, 2019)

The size of threats themselves are also increasing; IBM and Ponemon Institute published an industry benchmark report in 2017 that stated the average likelihood of experiencing a data breach involving more than 10,000 records is at 27.7% globally (Ponemon Institute, 2017).

As the security problem grows, organizations are recognizing the need to go beyond the traditional reactive approach (Graves, 2019; Solomon, n.d.), but there is evidence that there may be a lack of consensus in how to change strategies. A recent survey indicated that U.S. CEOs recognize cyber security as the biggest external threat to their companies (*C-Suite Challenge 2019: Survey*, 2019). Yet, another recent report from the United Kingdom suggested that, internally, companies struggle to find balance and agreement between security leaders and policy-making bodies (including boards of directors) on how to address these problems (Optiv, 2019). In essence, the first step has been taken by acknowledging a needed change in strategy, but the path forward is still unclear in how to execute.

One major concern for addressing security issues is the recognized skills gap in cyber security (United States Government Accountability Office, 2017). High rates of burnout and turnover make it difficult to retain qualified workers across multiple levels of the security profession (Bourget, 2017; Richmond & Lindstrom, 2015). Furthermore, demand for analysts increases with the recognition of cyber threats and the implications thereof (Bureau of Labor Statistics, 2016), especially in the commercial sector, which has largely been reactive to cyber threats (KPMG International, 2015). Industry research in this domain estimates that the shortage of these human analysts will range from 1.5 million by 2020 (Frost & Sullivan, 2017) to 3.5 million by 2021 (Morgan, 2017). These reports and other articles (HCL Technologies, 2019; Oltsik, 2019) highlight a critical need for more qualified workers. Though academia has responded with curricula to increase the pipeline (Bishop et al., 2017; L. Hoffman, Burley, & Toregas, 2012), the time it takes to recruit, train, and graduate novice-level candidates is too long to meet immediate demand for skilled, experienced individuals.

Incident response in the context of computer security is the act of addressing the short-term effects of security threats in an information system (NICCS, 2017). Teams that perform incident response, called Computer Security Incident Response Teams (CSIRTs), are often assembled at the enterprise level in both commercial and government sectors, and are comprised of skilled cyber security analysts, management, and specialized operators. CSIRTs typically reside within Security Operations Centers (SOCs) as a specialized group. As one of the major career channels for cyber professionals, CSIRTs are among the groups most affected by the labor shortage.

CSIRTs are an evolution of Computer Emergency Response Teams (CERTs), which date back to the late 1980's. The first CERT was assembled after the US Defense Advanced Research Projects Agency analyzed a widespread virus event and concluded that lack of communication and coordination in response efforts adversely affected efficient and timely response (Ruefle, Dorofee, & Mundie, 2014). The Software Engineering Institute at Carnegie Mellon University was then charged with handling major internet security incidents, and the CERT was born. In the following two decades, other government agencies, institutions, and commercial organizations formed teams of their own to help address the rapidly evolving security threats.

Since their inception, CSIRTs have evolved from ad hoc, reactive bodies into larger, more formal teams with structure and protocols that also incorporate monitoring and security maintenance. They are also distinctly different from IT solution teams in terms of scope of knowledge needed and types of tasks performed. The task scope of these teams includes an array of services, from proactive, to reactive, to maintenance (Killcreece, Kossakowski, Ruefle, & Zajicek, 2003). In order to remain effective, CSIRTs must adapt to a constantly changing threat landscape across the three service areas (Bada, Creese, Goldsmith, Mitchell, & Phillips, 2014; Steinke et al., 2015).

Within the various task areas, CSIRTS perform critical communication processes, such as handoffs, to coordinate within the team and beyond. In a recent manual for improving CSIRT effectiveness, Tetrick et al. (2016) identify ten (10) key areas to focus strategies and future research, which includes *information sharing*. Information sharing is central to collaboration and effective response. Information sharing encompasses basic, yet important, components of communication between sender and receiver (Sasaki et al., 2017). Within cyber operations, this includes navigating team expertise, knowing what information to share, when to share it, with whom, and how to do so in a way that establishes and maintains trust. For additional context, in order to achieve problem solving within the CSIRT, analysts must share or solicit knowledge from other members or outside parties, and quickly apply findings to mitigate the issue. Thus, another important aspect of CSIRT operations is maintaining clear and effective communication within and between constituents of the team, and with their parent or partner organizations.

Despite their importance to incident response, coordination activities remain a challenge in CSIRTs from both a within-team perspective (Steinke et al., 2015) and multi-team system perspective (T. R. Chen et al., 2014; Tøndel, Line, & Jaatun, 2014; Van der Kleij, Kleinhuis, & Young, 2017). Knowledge sharing and retention within the team is critical to team effectiveness, and is still identified as a current need for improvement (Steinke et al., 2015; Van der Kleij et al., 2017). Communication effectiveness within CSIRTs also has a direct impact on overall CSIRT effectiveness, yet continues to be neglected by managers as a focus for improvement (Tetrick et al., 2016). Information sharing outside the organization is even more complex with respect to trust and policy (Tanczer, Brass, & Carr, 2018). Though literature does not highlight specific issues pertaining to inter-organization collaboration in security, the need for effective collaboration between entities is explicitly identified, with implications for what could happen in its absence (Bada et al., 2014; Skierka, Morgus, Hohmann, & Maurer, 2015; Tøndel et al., 2014; Werlinger, Muldner, Hawkey, & Beznosov, 2010).

Recalling the shortage of cyber analyst talent and information sharing problems in CSIRTs, one potential and promising solution is employing aspects of automation to assist with knowledge sharing activities in CSIRTs, specifically in task execution and coordination. Not surprisingly, research in cyber defense has largely taken a computational approach to this solution space, which is consistent with trends in security science research as a whole. Computational approaches tend to focus on algorithms that can assist or replace human analysts performing security tasks (National Academies of Sciences Engineering and Medicine, 2017; Tyworth, Giacobe, & Mancuso, 2012). However, these efforts still require human interaction in incident response processes, and thus require supporting social and behavioral research to ensure effectiveness of design and execution. While computer scientists are making strides in automating cyber tasks using machine learning and natural language processing (Faysel & Haque, 2010), other research still suggests that human-automation teaming will produce better results than pure automation solutions (Lathrop, 2017; Shoshitaishvili et al., 2017). Prominent leaders and organizations in government and industry have also recently spoken about this disconnect between full automation approaches and human-automation integration. They encouraged human-centered sciences to bring expertise to the cyber security domain (National Academies of Sciences Engineering and Medicine, 2017), with the goal of integrating human analysts and autonomous capabilities to complement the strengths and weaknesses of each (Papadopoulous, 2017; Pomerleau, 2016; Williams, 2017).

From a human factors perspective, there are numerous opportunities for automation to assist CSIRT operations, and as the domain matures, these opportunities are likely to become more viable (Lathrop, 2017). Research and development efforts in security technology have started to address aspects of analyst tasks, including visualization (Tyworth, Giacobe, & Mancuso, 2012), asset mapping (Goodall, D’Amico, & Kopylec, 2009), and interface design (Lathrop, 2017). Many of these examples focus on individual functions operators and analysts, but could also be extended to other aspects of CSIRT operations, such as information sharing and collaboration efforts during incident response. Furthermore, there are gaps in addressing boundary information sharing practices between individuals and teams (Steinke et al., 2015), which can extend above the analyst level of the organization.

Automation may be able to assist with these critical information sharing activities and handoff points to increase awareness and security safety at higher levels of the organization. Some experts in cyber security and technology support the belief that the path forward in complex domains involves combining the strengths of humans and computers (Lathrop, 2017; Scharre, 2003; Williams, 2017), effectively creating a hybrid team. In the future, automation may become an actual ‘artificial teammate’, or an independent computer-based member of the team with its own set of roles, functions, and responsibilities.

Designing automation to work collaboratively with CSIRTs as a team member is an enormous undertaking, and could engage many different approaches from human factors at various stages of development. Function allocation, levels of automation, and at-the-screen interaction all represent critical issues and human factors research applications in the cyber security domain. The approach employed in this dissertation evaluates expertise and information sharing needs of security analysts, and compares those needs to current technology development. The gaps between needs and current technology informed a systems approach, or a broadened perspective that includes the physical systems, environment, human operators, and relationships between them. This approach was used to strategically develop concepts for future automation that can augment expertise and information sharing within a given team. Ultimately, this augmentation can improve the range of expertise available to the team during incident response, effectively improving task coordination, knowledge sharing, and information flow.

Human incident response teams are central to addressing the cyber security problem, but also are currently operating at a team-level expertise deficit. Task-related knowledge is essential to team performance (Paris, Salas, & Cannon-Bowers, 2000). Human factors and computational social science research have created methods for team knowledge assessment through individual assessment of knowledge (Cooke, 2004; Su, Huan, & Contractor, 2010). However, having individual knowledge or expertise alone does not ensure success. Sharing of that knowledge, or information sharing, is equally critical as it is the delivery process of the knowledge to where it is needed. Information sharing has been known to be a weak point in decision making teams (McGrath, 1984), and within CSIRTs, has been identified as one of the top ten opportunities for improving effectiveness (Tetrick et al., 2016). These facts provide a compelling case for addressing the gaps in expertise and information sharing by pursing options in the automation solution space.

To pursue development of human-automation teaming, it is essential to understand the human team's needs in terms of information and expertise, and how those needs might be addressed or supported with automation. Thus, in order to augment expertise in human teams performing knowledge-dependent functions, one must first identify knowledge needs in CSIRTs, and how information sharing occurs within the team and across team boundaries. In simple terms, this process will highlight gaps in information sharing, which automation developers and designers can address by strategically augmenting with autonomous capabilities. In so doing, researchers may find wide variation in how human analyst teams utilize expertise and execute information sharing (Ruefle et al., 2014). Thus, CSIRTs in diverse settings may have differing needs with respect to artificial teammates.

In summary, the investigation of information sharing and expertise in diverse CSIRTs aids in understanding the opportunities for autonomous systems to work within different team by augmenting expertise in information sharing functions. Due to the complexity of the environment and processes, systems approaches may be appropriate to frame and guide investigation. This dissertation employs several constructs and methodologies from systems engineering (SE) accordingly. This research concludes with proposing concepts for enhancing CSIRT operation through automation support for effective communication and collaboration. By remediating the expertise pipeline and assisting with critical coordination, research can then improve specific aspects of task coordination and information flow within incident response.

1.2. Objectives

The overall goal of this research was to determine expertise-based functional requirements for partially autonomous artificial teammates in CSIRTs. This involved identifying functions of computer security incident response that an artificial teammate should be able to perform or assist with, and functions to support interaction with human teammates. The scope included information sharing in incident response tasks and processes, as well as information sharing across team boundaries, such as handoffs and communication protocols between a given CSIRT and its parent organization.

The results of this research are aimed at cyber security researchers and technology developers who are currently designing and developing partially or fully autonomous solutions for CSIRTs. By incorporating research contributions focused on human behavior and expertise, this community will help to address the strategic aim of the National Academies of Sciences, Engineering, and Medicine (2017). The approach employed in this dissertation will be valuable in understanding the human aspect of cyber security, particularly analysts and team leaders in incident response, and how developing technologies could augment aspects of human expertise to improve effectiveness of teams at different stages of this process. Ultimately, the products of this research help direct new developments for better human-machine integration across other domains beyond cyber defense.

In addition to human factors theories and methods, systems engineering literature provides both frameworks and methodologies that are useful for determining organizational information flows and deriving functional requirements for solutions. Jackson (2000) provides a useful categorization of these approaches, which includes functionalist, interpretive, emancipatory, and postmodern systems approaches. This categorization helps determine which approaches are appropriate for a given system. Within the context of this dissertation, the ‘system’ to be investigated is the CSIRT, which is a tiered, complex organization with a common goal, and comprised of many different specialized teams of human analysts and managers that coordinate between them. One goal of this research is examining CSIRTs using multiple systems approaches to better understand system components and stakeholder needs for analysis and improvement. This dissertation research focuses on the human component, which is in relation to the tasks being performed and the larger context of the system.

The intersection of humans and technology presents ample opportunity to take systems approaches in designing appropriate solutions to improve overall system performance. The cyber community lacks expertise in scientific domains focused on understanding the human component of the system, leaving a large gap between the problem and proposed solutions produced by researchers in computer and security sciences (National Academies of Sciences Engineering and Medicine, 2017). The research presented in this dissertation aims to address this gap by providing the missing human context in computer security incident response, helping to highlight potential solution spaces and provide benefit to the entire system with high ecological validity. By studying the performance, expertise, and dynamics of human CSIRTs, this research provides a set of system requirements for cyber researchers to develop and integrate solutions into partially autonomous agents for computer security incident response.

1.3. Document Overview

The rest of this dissertation will address relevant literature, methods, results of three (3) studies, and presentation of findings and implications. Chapter 2 highlights relevant literature in cyber security research, human factors topics related to cyber security, and elements of cyber incident response. Chapter 2 also outlines the research conceptual framework used in this dissertation, which employed frameworks from expertise and systems engineering to shape the lens through which the problem was analyzed. Chapter 3 presents the research questions and methodology for conducting this research, including details about the methods chosen. Chapter 4, 5, and 6 present approaches and findings from the three respective studies, and include answers to the research questions. Chapter 7 provides discussion on research implications. Finally Chapter 8 completes the dissertation with general summaries, conclusions, and broader implications, and proposes future directions for research.

CHAPTER 2. LITERATURE REVIEW

2.1. Introduction

This chapter provides background in specific areas that are relevant to the proposed research topic. Section 2.1 presents background on the state of cyber security in 2019, followed by surrounding context of the size of the problem in Section 2.2. Section 2.3 highlights past and current research areas within security science including technological approaches. Relevant topics in human factors are reviewed in Section 2.4, emphasizing key theories and concepts that relate to human-automation in cyber security presented in Chapter 1. The last sections focus on the intersection of cyber security literature and human factors research (Section 2.5), followed by further scoping (Section 2.6) and review of background research on incident response teams in cyber security. Section 2.7 frames the problem statement using systems engineering perspectives, queuing the transition to Chapter 3, which addresses the research conceptual framework.

2.2. Challenges in Cyber Security

Cyber security touches almost all aspects of today’s interconnected world. The need to address issues in this domain is recognized at national (The Council of Economic Advisors, 2018; The White House, 2018) and international (EY, 2018) levels, with increasing investment over the last several years (EY, 2018; Gartner, 2018; HCL Technologies, 2019). Healthcare, finance, government, and other business data are all at risk of security breaches (FireEye, 2019; Identity Theft Resource Center, 2019), and the fight to protect these critical areas is ongoing. The risk to these sectors increases as the attack surface grows. Attack surfaces are affected by a variety of factors, including but not limited to the number of devices connected to a network, the security of those devices, user awareness of threats, and handling practices of data contained in a network. Threats to cyber security grow and evolve daily, making it difficult for researchers and analysts to make significant progress across all areas of cyber defense (FireEye, 2019), especially regarding internet of things botnets, cloud operations, and malware (Cisco Systems, 2018; Darktrace, 2018; Symantec Corporation, 2019). Hackers are becoming more sophisticated in their attack schemes and more effective at finding vulnerabilities in aging IT infrastructures of many companies (Cisco

Systems, 2017), and leveraging automation to conduct hacks (Brundage et al., 2018; Sophos Labs, 2019).

Humans still provide the majority of analysis and response in computer security (Williams, 2017). Yet, the industry is currently in desperate need for qualified candidates to fill a growing number of analyst positions (Bureau of Labor Statistics, 2016; Cobb, 2016; Frost & Sullivan, 2017; *Growing the Security Analyst: Hiring, Training, and Retention*, 2014). One solution to addressing this deficit is to increase the pipeline of humans capable of doing these jobs. Many publications that address educating and recruiting candidates for these roles outline some general skillsets and characteristics, approaching the problem from traditional models of expertise (Bishop et al., 2017; Hoffman et al., 2012) and knowledge, skills, and abilities (KSAs) (Assante & Tobey, 2011). Academia and educators have tried to address the growing demand for cyber security analysts with structured approaches to education (Bishop et al., 2017; Ruefle et al., 2014; West-Brown et al., 2003).

Cited articles above indicate that in the time period from 2003 to 2017, there has been much focus on the educational side of computer security in determining curricula that support technical areas of expertise. Additionally, these efforts also aim to identify and cultivate certain characteristics, such as curiosity and integrity, as well as critical thinking, problem solving, and verbal communication. Other related approaches have proposed to address this problem involve educators working directly with industry and policymakers to develop a more holistic workforce in cyber security (Hoffman et al., 2012). Despite development in educational and recruiting strategies, there continues to be a growing skills gap in 2019, represented by the number unfilled positions in the field (HCL Technologies, 2019; Oltsik, 2019). Moreover, retention of hired analysts is a problem (*Growing the Security Analyst: Hiring, Training, and Retention*, 2014). Limited career growth, high demand, and rising salaries are among reasons why qualified analysts have relatively short tenure at a given firm (HCL Technologies, 2019).

Beyond the technical skills, companies continue to cite business acumen and communication skills as necessary but lacking in cyber analysts, including university graduates trained in technical areas of cyber security (HCL Technologies, 2019). Communication skills in particular are necessary for effective collaboration and information sharing. These activities are central to incident response operations in human CSIRTs, but continue to illustrate a classic social dynamics problem that persists beyond technological development (Tetrick et al., 2016). Some industry-based research has produced specific improvement strategies for managers and human resources, focusing on improving teams through different hiring practices, simulation training, and standard work tools (*Growing the Security Analyst: Hiring, Training, and Retention*, 2014). In summary, one of the biggest concerns in the cyber security domain is in the need for development of both technical and non-technical skills, as well as expansion of both human capital and technological resources on a long-term basis. However, current strategies to meet an increasing demand for CSIRTs have not been sufficient.

2.3. Related Security Science Research

Research in cyber security, also known as security science, has been active since the late 1970's (National Academies of Sciences Engineering and Medicine, 2017). For the sake of narrowing the literature focus, this very large body of research can be broken down in to hardware-based research (which focuses on infrastructure and hardware engineering) and software-based research (which explores the logic and language of the programs operating in a network). The more relevant of these two branches to this dissertation is the software track.

2.3.1. Software-Based Research

Traditionally, software-based research takes one or more of the following technological approaches: cryptography, programming and semantics, and security modeling (National Academies of Sciences Engineering and Medicine, 2017). Cryptography is the mathematical derivation of logic structures and technologies behind encryption. Programming and semantics is a very popular approach currently, and includes development of models, algorithms, and languages for system security. Some applications include insider threat detection (Bowen, Devarajan, & Stolfo, 2011; Buford, Lewis, & Jakobson, 2008), incident detection systems (Faysel & Haque, 2010; Kumar, 2005), and network security assessments using game theory (Roy et al., 2010). Lastly, security modeling helps researchers understand the implications of policies as they are enforced across networks, as well as better understanding of threats and system behavior (Goldstein, 2016; Memon, 2014).

A recent publication from the National Academies of Sciences Engineering and Medicine (2017) points out that many approaches to cyber security research have a strong technological focus, and often leaves out the human component of the problem. Tyworth et al. (2012) also recognized this issue with a short review of cyber security literature that maintains the position that the answer is in technology, not humans. Singh & Nene (2013) reinforce this belief, claiming that humans (as operators and end-users) are viewed as a limitation, and the usefulness lies in abstracting human information processing into models to support algorithmic approaches. In summary, this area of research largely represents an approach that excludes humans as a significant and central piece of the cybersecurity research solution.

2.3.2. Development of Automated Solutions: The SOAR Platform

Considering the labor shortage and need to provide security coverage, a new type of software has been developed within the last several years to help automate low-level tasks and standardize incident response within a given organization. Security Orchestration, Automation and Response (SOAR) technologies are software platforms designed, built, and marketed towards incident response organizations to increase capacity and efficiency. The main issues these platforms aim to address are software integration and analyst time on an incident. CSIRT analysts must monitor, use, and pivot between a variety of programs and appliances. SOAR technology aims to reduce pivoting by integrating signals from different tools into a single interface. Moreover, the industry has recognized the shortage of analysts, which has resulted in an aim to reduce hands-on time of analysts on menial tasks by automating low-level activities and allowing more time for applying expertise in more difficult tasks.

There are at least a dozen different platforms available on the market as of early 2019, and more software companies are aiming to add SOAR capabilities to their existing products and services to compete. According to Gartner, the expected growth of SOAR solutions in practice is 15% by 2020, up from 1% in 2018 (Chuvakin & Barros, 2018). This market validation (Bhargava, 2018) indicates that more firms are recognizing the potential benefits of automation and orchestration in their security organizations, as well as the need to address labor shortages, data deluge, and disparate tools.

Gartner, Inc. conducted a detailed analysis of SOAR capabilities (Neiva, Lawson, Bussa, & Sadowski, 2017) that identified requirements of what platforms should be able to do to meet industry needs. Some of the requirements presented in the Gartner report clearly overlap with the dimensions of expertise framework (subject matter expertise, interface/tool expertise, expert identification expertise, situational context expertise), while others do not (communication expertise, information flow path expertise). This report is commonly referenced in SOAR platform websites, and features closely align with the report recommendations. Indeed, it seems to be the "gold standard" for SOAR requirements. However, much of the recommendations aim to address organizational needs in computer security.

Since 2017, developments in SOAR have helped to alleviate some of the current pressures in the market. However, experts emphasize that continued focus on human-centric aspects of this technology will be critical to actually bridging gaps in the field (Oltsik, 2018a, 2018b; Staples & Sullivan, 2018). Employing other approaches to determining *user* needs may also be useful in developing system requirements, and may help address some underlying causes of the issues felt in the field.

These areas of security science research continue to progress and improve the overall state of security. However, the improvement remains incremental and defensive compared to the advancements of attackers (Curry, 2019; Staples & Sullivan, 2018). Thus, other approaches, which include social and behavioral sciences, have been called for by prominent research directives when scoping the future of research in the cyber domain (National Academies of Sciences Engineering and Medicine, 2017). This dissertation will further investigate how technological solutions (SOAR platforms) address analyst expertise needs, and what gaps persist between those needs and current SOAR features.

2.4. Applicable Human Factors Research

Topics in the human factors domain of human supervisory control are applicable to the cyber security problem and the potential automation solution identified in the introduction. Generally speaking, two key areas include function allocation (FA) and human-automation interaction (HAI), which both include subtopics that relate to the described problem space. Current publications mainly focus on exploring FA and HAI in broader application areas, but do include some applications in cyber security.

The following sections describe two major areas of human factors literature that directly relate to human-automation teaming. FA research, described in Section 2.4.1, has explored assignment of tasks between humans and machines dating back to the 1950s when automated functions were introduced into more complex task domains. As technology further developed, specific research in interaction between humans and automation emerged. The result was the formation of a sub-domain in human factors that is most often referred to as human computer interaction (HCI). These two areas support major findings and discussions regarding decisions and guiding principles for designing systems that facilitate human-automation teaming.

2.4.1. Function Allocation and Human Supervisory Control

Human-machine function allocation is a long-standing topic in human factors, dating back to 1951, when Paul Fitts produced a guideline for deciding task assignment between humans and machines (Fitts, 1951). This approach, later referred to as Humans-are-better-at, Machines-are-better-at (HABA-MABA) (Price, 1985) has persisted over the decades to produce additional versions of the Fitts List. However longstanding, the Fitts Lists were determined to be not as useful due to disconnects between the assumptions and generalness of the HABA-MABA lists and the concepts and specificity of engineering domains (Price, 1985). These lists also do not reflect considerations of operational context of the tasks described (Hancock & Scallen, 1996).

Despite some of its shortcomings, the Fitts List initiated a new domain between human factors and other research domains in FA and human supervisory control. Traditional approaches in this domain base research and design on models of human information processing and decision making

(R Parasuraman, Sheridan, & Wickens, 2000; Sheridan & Ferrell, 1974). To mitigate ambiguity of the benchmark Fitts report, much human factors work has been done in this area to develop systematic approaches to FA (Price, 1985), function allocation requirements (Feigh & Pritchett, 2014), measurement and metrics (Pritchett, Kim, & Feigh, 2014), models (Hollnagel & Bye, 2000), and more. Levels of automation (LOA) are often a topic of interest within FA and HAI. LOA concepts focus on what level of automation independence / autonomy is appropriate for which functions based on the function being performed and the context within which it exists (R Parasuraman et al., 2000).

Despite issues in executing and delivering FA solutions based on the classic, static model (Hancock & Scallen, 1996), it is generally agreed that it is critically important to consider dynamic elements of function allocation contexts and criteria early in automation design, before decisions about displays and interfaces are made (Pritchett et al., 2014). In response to increasing complexity of human-machine interactions, and the recognition of failures of static FA over 40 years of design (Hancock & Scallen, 1996), two new concepts emerged: dynamic function allocation (DFA) and adaptive function allocation. Bridging issues between HAI and FA, these concepts promote real-time function allocation based on operator workload and environmental context. Literature in this area has focused on three general categories (Scallen & Hancock, 2001): 1) identifying critical issues (Feigh, Dorneich, & Hayes, 2012), 2) empirical investigations (Hilburn, Molloy, Wong, & Parasuraman, 1993), and 3) systematic empirical evaluations of adaptive automation (Miller & Parasuraman, 2007). Put simply, DFA (and adaptive function allocation) can be used to frame function allocation in today’s complex systems (Challenger, Clegg, & Shepherd, 2013). Discussions on function allocation and automation development in CSIRTs will be further explored in Chapter 7 of this manuscript.

2.4.2. Human-Automation Interaction and Teaming

Human-Automation Interaction is a subset of HCI research that focuses on the design and effectiveness of interactions between humans and automation, from single-function machines to complex autonomous agents. HAI research ranges from at-the-screen tasks to higher-order problems with cognition and information processing of humans in cyber-physical systems.

Literature in HCI is vast, and divided amongst several larger domains, including human factors, design, and now computer science. This subset has produced some literature in cyber security applications, which are reviewed in 2.5.2.

Concepts surrounding HAI often cross those of function allocation, and expand into research questions regarding how automation can assist or augment the human performing physical and cognitive tasks (de Visser & Parasuraman, 2011; Demir, McNeese, & Cooke, 2016; Mercado et al., 2016; Raja Parasuraman, Barnes, & Cosenzo, 2007). With the growth of digital technology, work environments have become more complex and automation has rapidly advanced; automation technology is able to learn from incoming data, compute at rapid speeds, and predict behaviors of humans and processes. These trends have sparked new concepts surrounding collaboration between humans and automation, effectively creating a team between the two entities.

Drawing on the concepts of DFA and LOA from 2.4.1, research in human-automation teaming has grown in the last 20 years. Instead of the mutually exclusive approach to function allocation between humans and computers, a new generation of research has emerged, focusing on how to create a team out of humans and one or more *intelligent agents* (Cuevas, Fiore, Caldwell, & Strater, 2007; Sycara & Lewis, 2004). Here, an intelligent agent refers to an automated computer or machine component that amplifies the performance of the human in concert with the human, thus allowing the machine component to perform tasks for which it is better suited than the human. An example of this concept is a search task through large databases for specific points (Bradshaw, 2015). These human-agent teams have been referred to as *centaurs,* a term coined by Garry Kasparov to describe successful human-machine teaming in chess (Holtel, 2015). The centaur system (involving human and computer) was able to defeat a (fully autonomous) computer in a famous chess competition, opening avenues for scientific research and development in hybrid teams based on empirical evidence of their efficacy (Scharre, 2003). These hybrid teams present compelling research opportunities for many domains, including within cyber security (Truvé, 2017).

There are a variety of research and development publications in this area addressing humans and automation on the same team. A framework in augmented cognition, or the idea of assisting the human with cognitive tasks, has been proposed for teams in complex environments, and draws on team performance and HAI research to determine some level of requirements for augmented cognition systems (Cuevas et al., 2007). Embodied cognition, an area which explores mechanisms of social interaction and decision-making, is also starting to be incorporated into human-automation settings (Dautenhahn, Ogden, & Quick, 2002). Researchers have also been investigating human-automation teaming by exploring how to use humans as sensors for computers in pattern recognition (Rutkin, 2015). These studies provide groundwork for replication and further investigation in other complex environments, including the cyber domain.

Though the scope of research opportunities in human-automation teaming is vast from both human and machine domain perspectives, there are limitations from development of silos within this landscape, as well as application and integration of results. Examples from recent workshops by the National Academies of Science Engineering and Medicine ("Session 7: Humans and Machines Working Together with Big Data," 2017; "Session 8: Use of Machine Learning for Privacy Ethics," 2017) indicate that, while specific aspects of human-automation teaming are being addressed through research, additional approaches continue to exist independently within computer science disciplines. These examples implicate recent within-discipline studies that may scope problems to exclude social and behavioral science perspectives, thus limiting application generalizability and ecological validity. In order to help bridge this gap, this dissertation will explore how to effectively collect human-centered data and integrate findings in to technological development methodologies.

2.5. Current Human Factors Research in Cyber Security

Despite the vast opportunities to apply human factors methods and principles in the cyber security domain (Borghetti, Funke, Pastel, & Gutzwiller, 2017), the body of research is relatively small at present. This section provides an overview of current human factors research in the cyber security application domain. Cyber situation awareness, described in Section 2.5.1, is the largest of the topics within this subdomain. Section 2.5.2 provides an overview of the next largest topic: human-automation interaction within cyber security. A potpourri of other human factors topics, defined in Section 2.5.3, have also been explored in cyber, but do not amount to a significant corpus as a whole.

2.5.1. Cyber Situation Awareness (CSA)

Much of the work being done by human factors researchers in cyber security is focused on improving situation awareness of operators. Situation awareness of cyber networks is becoming more challenging as networks grow and become more complex (Endsley & Connors, 2014). As related problems become a prominent issue in cyber defense, CSA has grown to become a relatively large body of literature within the human factors cyber security domain. Some notable review papers have been published in the last decade to cover work in this area, which highlight the variety of research topics that ultimately connect back to CSA (Franke & Brynielsson, 2014; Tadda & Salerno, 2010a). One such review by Franke and Brynielsson (2014) also identified 11 clusters of article topics within CSA, including visualization, human-computer interaction, tool and algorithm development, and information exchange, to name a few. Recent examples of CSA work focus on context-driven CSA research (Tyworth, Giacobe, & Mancuso, 2012), visualization and user-centered design concepts to improve CSA (Mancuso, Staheli, Leahy, & Kalke, 2016), user alerting in internet-of-things (Kammüller, 2018), and improving cognitive models computer gaming (Domínguez, Goodwin, Roberts, & Amant, 2017).

Literature that studies security commonly cites the need to understand situational context as driving requirement for analysts (Ahrend, Jirotka, & Jones, 2016; Bishop et al., 2017; Ruefle et al., 2014; Steinke et al., 2015; Tøndel et al., 2014; Yufik, 2014) thus human factors research in this area is timely and appropriate. Recent developments in specific applications of CSA in practice focus on consolidating tools to a single screen (Engelbrecht, 2018; Oltsik, 2018b; Sophos, 2019) and shared awareness across collaborating analysts (Gutzwiller, Fugate, Sawyer, & Hancock, 2015b; Tyworth, Giacobe, & Mancuso, 2012; Vieane et al., 2016). Though this dissertation does not explore situation awareness as a central topic, the research design does explore contextual settings to support with evidence this critical aspect of incident response.

2.5.2. Human-Automation Interaction in Cyber

As mentioned in previous sections, opportunities to apply HAI concepts in complex domains like cyber defense are abundant. Consequently, HAI is the next most significant topic in the subdomain of human factors in cyber security after CSA. However, much of this research stops short of human-automation *teaming*.

Cain & Schuster (2014) discuss human-automation interaction in terms of types of CSA and high-level propositions for how information should be communicated between humans and autonomous agents working together in cyber applications. Other prevalent HAI literature in cyber concerns the Smart Grid, or the updated power grid infrastructure of North America. Specifically, researchers considered a framework that included levels of automation, adaptive autonomy, and performance shaping factors for Smart Grid operations (Boroomand et al., 2010). Another study, not involving the Smart Grid, explored trust and reliability aspects of HAI within the cyber domain, though a major limitation of the study was the lack of experience of the pool of participants (Brown, Christensen, & Schuster, 2016). Within cyber operations, studies have explored HAI in vulnerability analysis functions specifically (Shoshitaishvili et al., 2017). However, this study shifts the paradigm from tools assisting humans to the reverse, focusing on how automated systems can leverage humans better based on human expertise. In general, existing research does not address next steps of HAI, which is the interaction of a human and an artificial autonomous teammate (Lathrop, 2017). This dissertation will further explore diverse contextual settings to identify opportunities for human-automation teaming, as well as compare those findings to current automation development in cyber security.

2.5.3. Other Human Factors Topics in Cyber

As mentioned in (Gutzwiller et al., 2015b; Vieane et al., 2016), there are many opportunities for human factors work to contributed to research in cyber security. Human factors literature stretches beyond the topics covered in the previous sections to include a wide range of subjects within cyber security. Two sub-groups (Table 1) within this research potpourri are cognitive processes and design. The first sub-group focuses on understanding how analysts and end users process information and make decisions, which includes a myriad of different situations and software-specific studies. Visualization is a popular topic in human factors as a whole, with some representation in the relatively young cyber domain. Research in this area focuses on how to help analysts and users increase awareness and understanding of threats and vulnerabilities (Healey, Hao, & Hutchinson, 2014). Additionally, general usability of cyber security software tools has been a subject of interest since the early 2000s, and continues to draw focus from researchers in human factors as well as industry practitioners (Schultz, 2012).

Table.1. Potpourri Topics on Cyber Security in Human Factors

Topic Area	**Relevant Work**
Usability; design	(Healey et al., 2014; Schultz, 2012)
Information processing; decision making	(Beitzel, Dykstra, Toliver, & Youzwak, 2018; Gonzalez, Ben-Asher, Oltramari, & Lebiere, 2014; Massey, Seker, & Nicholson, 2018; Mihajlov & Jerman-Blazic, 2018; Muggler, Eshwarappa, & Cankaya, 2018; Proctor, 2015; Yen, Erbacher, Zhong, & Liu, 2014; Zinke, Anke, Meyer, & Schmidt, 2018)

2.6. Human Teams in Cyber Security Incident Response

Incident response is typically performed by teams called CSIRTs, which can range in size from less than 5 people to more than 100 (Ruefle et al., 2014). These teams of ‘cyber first responders’ also exist at varying levels of operational scope, from national teams to smaller teams within companies or organizations (Bada et al., 2014; Killcreece et al., 2003; Ruefle et al., 2014). The goal of incident response teams in any context is to address incidents (or emergencies) as quickly as possible while minimizing cost and damage to the relevant subject in the application domain; this could be a human patient, the general public, or the control system, respectively. For CSIRTs, the relevant subjects include hardware, software, activity of and information stored in the system or network.

CSIRTs are just one component of the cyber defense system, but provide critical support in mitigating malicious events once identified. The various roles within these teams require different levels of different types of expertise, even beyond traditional emphases on subject matter domain knowledge (Garrett, Caldwell, Harris, & Gonzalez, 2009), which all impact the team’s ability to perform (Tetrick et al., 2016). Research in this area, reviewed in the subsections below, provides some alternative approaches to improving the state of cyber security, mainly by drawing on organizational psychology and human factors literature to improve team effectiveness and performance.

2.6.1. Parallels in Incident Response

Some approaches have drawn the parallel between incident response teams in other domains and those in cyber (Steinke et al., 2015). In fact, some of the same language used in other incident response literature is used to describe cyber incident responder roles (Newhouse, Keith, Scribner, & Witte, 2017). In theory, cyber incident response has common characteristics with other models in event response, such as disaster response.

The Disaster Management Cycle describes four stages: response, recovery, mitigation, and preparedness (FEMA, n.d.). When comparing this cycle to cyber defense, it is evident that the equivalent process in cyber often stops at response, and rarely reaches recovery. The cause of this truncation stems from poor executive-level decisions regarding cyber risk mitigation strategies and spending (KPMG International, 2015), analyst workload (Borghetti et al., 2017), and a mindset of “patch and pray” (Grose, 2007; Ravindranath, 2015), a common approach to fixing only known problems, and only when they manifest themselves. Essentially, many firms do not have the resources to support more sophisticated or developed models of response.

Disaster response models inspire possible future extensions of research as the cyber domain matures. Efforts must first focus on getting the field beyond infancy (Ruefle et al., 2014), which could begin with exploring smaller, more specific aspects of incident response. The lessons learned from other incident response environments may provide important contextual considerations for improving performance of CSIRTs. For instance, comparisons to emergency medical teams, national power plant operations teams, and military response teams have highlighted some opportunities for increasing team effectiveness in CSIRTs (Steinke et al., 2015). One finding in Steinke et al (2015) identified that, as in other incident response contexts, handoffs in CSIRT operations are critical, with a recognized need for potential investigation and improvement. Thus, handoff functions in CSIRT operations offer a vetted opportunity for research. Handoffs will be explored within this dissertation as information sharing functions, with the goal of identifying specific types in different teams, and how these might be improved.

2.6.2. Functions of a CSIRT

The title of the team handling incidents may vary across organizations, such as incident response, incident handling, or incident management teams (Killcreece et al., 2003). These teams typically exist within large organizations or governments which can support the cost of a team of highly qualified security professionals (Horne, 2014). Not all organizations have their own incident response team, but rather have some sort of lower level capabilities with threat detection and incident handling. Should incidents escalate, these smaller organizations may rely on external entities, such as managed security service providers (MSSPs), to investigate and resolve issues.

CSIRTs are responsible for an array of functions that range from reactive services to security management (Killcreece et al., 2003). Ruefle et al. (2014) present graphical representations of CSIRT functions (Figure 2) and the incident handling lifecycle (Figure 3). Complimenting the functions presented in Figure 2, these teams are typically comprised of a range of roles, positions, and specialties (Lathrop, 2017; Newhouse et al., 2017). Not all members will have every set of skills required to perform various tasks (Ruefle et al., 2014). Additionally, the team itself may not possess or have consistent access to knowledge needed to perform certain functions in information sharing.

Oftentimes, CSIRTs are considered multi-team systems (MTS), comprised of different groups working on specialized tasks (Tetrick et al., 2016). Sometimes, like the in the case of an MSSP working with a small customer, the MTS is distributed over multiple organizations. Accordingly, CSIRT operations tend to be complex in both structure and operations as constituent teams coordinate between each other to perform various functions and processes. Each specialized team typically has a structure that includes operators performing the specialized response tasks, and a lead or manager directing work tasks, performing critical decision making, and communicating upward and outside the scope of the team. Coordination of the entire MTS is typically under the responsibility of the Chief Information Security Officer (CISO) or equivalent position (Hale, 2017).

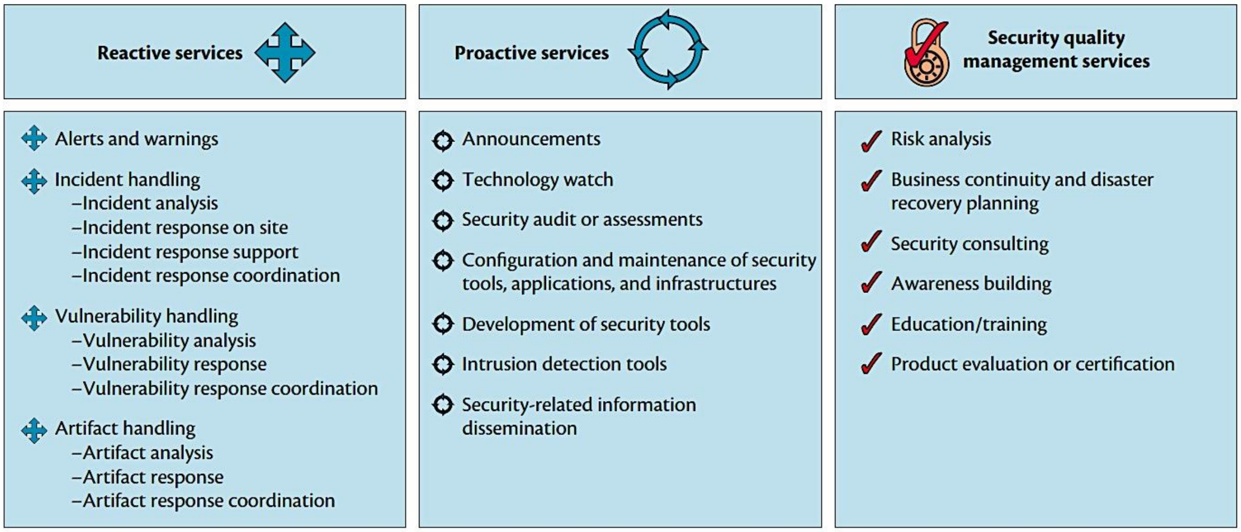

Reactive services

- Alerts and warnings
- Incident handling
  - –Incident analysis
  - –Incident response on site
  - –Incident response support
  - –Incident response coordination
- Vulnerability handling
  - –Vulnerability analysis
  - –Vulnerability response
  - –Vulnerability response coordination
- Artifact handling
  - –Artifact analysis
  - –Artifact response
  - –Artifact response coordination

Proactive services

- Announcements
- Technology watch
- Security audit or assessments
- Configuration and maintenance of security tools, applications, and infrastructures
- Development of security tools
- Intrusion detection tools
- Security-related information dissemination

Security quality management services

- ✓ Risk analysis
- ✓ Business continuity and disaster recovery planning
- ✓ Security consulting
- ✓ Awareness building
- ✓ Education/training
- ✓ Product evaluation or certification

Figure.2. Computer Security Incident Response Team functions (Ruefle et al., 2014)

As defined in (Killcreece et al., 2003), incident handling includes (a) receipt of an incident from a variety of internal and external sources; (b) triaging the incident in terms of category and severity to direct actions towards the proper resources and at an appropriate priority level; (c) analysis and investigation of the incident, searching for potential causes, vulnerabilities, and extent of the incident; (d) mitigation of incident damages; (e) patching or resolving vulnerabilities identified in the process; and (f) trending and correlating across incident reporting. Figure 3 depicts general steps, indicating that incidents vary in type, severity, and scope. Furthermore, the operational focus of CSIRTs may vary across organizations, depending on their scope and services (Ruefle et al., 2014). After the response protocol, some sort of lesson learned or related maintenance task is typically incorporated back into the repertoire of activities to be performed by the CSIRT on an ongoing basis. Consequently, these ongoing maintenance activities add to the extensive list of responsibilities of the team. The addition of these activities does not result in additional operators or time to complete, and increasing the workload on analyst in existing teams.

Information sharing is central to incident response effectiveness. Teams need to be able to share knowledge and expertise within the team and between other constituents in order to coordinate response and remediation efforts. Despite the criticality of these points highlighted in CSIRT literature and other domains of incident response, there is a dearth of literature concerning communication and coordination tasks during escalation and handoff. Some research has highlighted information sharing as a problem, and provided some deployable hiring and training techniques for improving the social dynamics of CSIRTs (Tetrick et al., 2016). However, it is unclear how much these approaches will improve team effectiveness, and if they are feasible to implement given the currently stressed state of the system. This gap will be further researched in this dissertation, specifically regarding the domains of expertise required for information sharing within CSIRT operations.

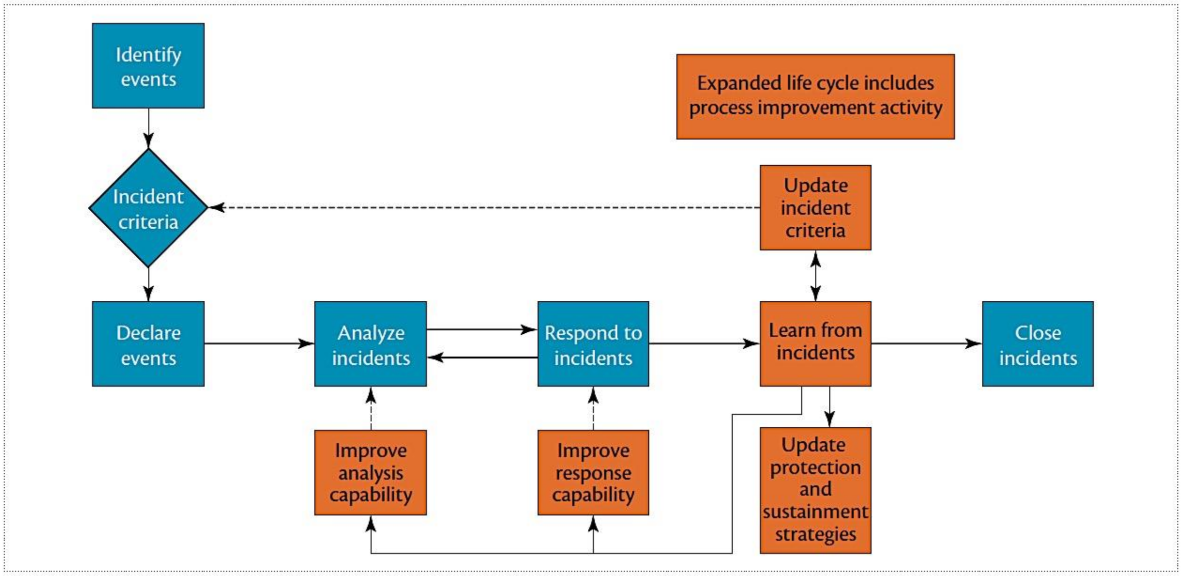
Identify events
Incident criteria
Declare events
Analyze incidents
Respond to incidents
Learn from incidents
Close incidents
Expanded life cycle includes process improvement activity
Update incident criteria
Improve analysis capability
Improve response capability
Update protection and sustainment strategies

Figure.3. Incident Handling Lifecycle (Ruefle et al., 2014)

2.6.3. Human Analyst Qualifications

Providing a variety of services to internal and external organizations, members of CSIRTs tend to have varying degrees of expertise across cyber defense subjects, including incident handling, vulnerability analysis, and artifact handling, to name a few (Ruefle et al., 2014). In order to remain effective, CSIRTs must adapt to a constantly changing threat landscape with little to no warning if, when, and where events will occur (Bada et al., 2014; Steinke et al., 2015). Communication within the team is vital for effective response (T. R. Chen et al., 2014; Tetrick et al., 2016; Werlinger et al., 2010). Communication beyond the scope of the CSIRT is also critical to keep the parent organization abreast of plans, findings, and activities of the team; external communication with industry and government are also becoming increasingly important. Communication within the organization is the responsibility of the team lead, who must also direct tasks and help make decisions regarding incidents and their respective responses.

Generally, the search for CSIRT candidates focuses mainly on certifications, such as the Security+ or CISSP certifications, and traditional domain topic areas, such as network protocols, hardware, software, and so on. According to West-Brown et al. (2003), “many people incorrectly consider the most important attribute in CSIRT staff to be their technical expertise” (p.168). Researchers who recognize the importance of other key personal attributes have pointed out the significance of characteristics such as adaptability, learning, teamwork, and flexibility (Bada et al., 2014). A recent handbook on CSIRT social maturity goes so far as to add information sharing practices and communication techniques into simulation interviews to ensure candidates are vetted for qualities related to effective social dynamics (Tetrick et al., 2016).

One existing framework reference in literature is the National Initiative for Cybersecurity Education (NICE) framework, a product of previous research linking work roles to specializations and KSAs in computer security (Newhouse et al., 2017). This purpose of this framework is to provide a lexicon and structure for reference, and provides a comprehensive overview of mapping between tasks, roles, knowledge, skills, and abilities as they relate to specializations within the entire cyber security domain. While useful as a reference, the NICE Framework does not provide context regarding specific information about CSIRTs, such as team structure, prioritization of critical elements for incident response, or how to apply the NICE framework to CSIRTs in general. Tetrick et al. (2016) helped to address the missing context with a handbook that specifically identifies key knowledge, skills, abilities and other attributes in CSIRTs with the overarching theme of improving the organizational effectiveness in the human operations. Still, the changing roles and skills required is a constantly evolving model that adapts to emerging threats and trends; cyber professionals need to engage in continuous learning to stay competitive and effective (L. Hoffman et al., 2012; Oltsik, 2017). Static frameworks of KSAs and role definition may not be robust enough to base long-term research, especially as organizations organically adapt to the changing landscape.

Several studies explore tasks performed in cyber defense: one in particular uses the NICE framework as a starting point (T. R. Chen et al., 2014). This study, from organizational psychology, used a team-based approach to studying CSIRTs, performing individual task analysis, link analysis, cognitive task analysis, and team task analysis to determine KSAs that hiring managers should look for in successful analyst candidates. The goal was to close the gap between the needs of a CSIRT and the hiring criteria used by managers, specifically between expertise and KSAs. Other studies use types of task analysis to determine cognitive demands on operators as they perform analysis, and included some aspects of domain knowledge to construct attribute tables (D'Amico, O'Brien, Whitley, Tesone, & Roth, 2005).

2.6.4. CSIRT Effectiveness

The security industry as a whole has recognized a need to formulate and track metrics at the enterprise level in order to adequately communicate risk to board-level executives (Asher-Dotan, 2015). In 2017, one survey indicated that 82% of respondents had defined such metrics (Scale Venture Partners, 2017). The critical, and perhaps obvious, metric in computer security incident response is time to quarantine an incident. Two emerging metrics are now commonly used: mean-time-to-detect and mean-time-to-respond. These measurable timescales help differentiate between how long a threat was undetected versus how fast the organization could react (Petersen & Lentz, 2015). However, enterprise metrics can oversimplify performance in these complex environments. By only measuring time, metrics focus only on meeting organizational goals around security, and do not go deeper into operational performance factors at the team level.

At the CSIRT level, (Ruefle et al., 2014) state there is no agreed-upon measure of effectiveness for this team. Steinke et al (2015) and Tetrick et al (2016) discuss methods for addressing team effectiveness and the absence of research on CSIRT team performance by using team-based methods. Methods borrowed from emergency response and military response were used to suggest training techniques, briefings, and handoff checklists, which aim at improving communication within a team, building trust in stressful situations, and decreasing errors at critical handoff points (Steinke et al., 2015). Other human factors literature could also be useful in developing frameworks for team effectiveness in incident response teams, including CSIRTs (Caldwell, 2015). These opportunities for research and development have not yet been pursued, yet offer rich potential not only for traditional human factors research, but also for automation development. This dissertation considers this development of measures of effectiveness during observational and interview studies with CSIRTs, especially in relation to information sharing.

As previously mentioned, information sharing was one of the key areas identified in (Tetrick et al., 2016) as central to CSIRT success. Information sharing in CSIRTs includes the sharing of knowledge or expertise. *Transactive memory* is a well-studied theory in teamwork related to knowledge sharing. This theory explains how individual members of a group store and recall information, including where knowledge is stored, used, and how it can be accessed (Palazzolo, 2005). Transactive memory has been identified as an important aspect in incident handling (Tetrick et al., 2016) and relates to the expertise framework presented in Chapter 3. Mancuso (2012) applied a variety of methods to study transactive memory on distributed cyber security teams. Other recent expertise studies investigated the performance differences between novices and experts in cyber (Eldardiry & Caldwell, 2015; Silva, Emmanuel, McClain, Matzen, & Forsythe, 2015). The Silva study specifically addressed strategies and methods for studying these teams in situ; the Eldardiry & Caldwell study emphasized information display and presentation tools to support knowledge sharing. Team-based methods have also been applied to understand skills and tasks within a team, and how that affects overall team effectiveness (Steinke et al., 2015).

Other approaches have been explored to study effectiveness in CSIRTs. Anthropological approaches identified tool development as a priority, as tools helped analysts perform tasks faster (Sundaramurthy, McHugh, Ou, Rajagopalan, & Wesch, 2014). This approach also found that studying CSIRTs is challenging due to team members not trusting researchers, as well as the closed culture of information security. Another study used the Delphi method, a qualitative method that procuring information from experts, to formulate a framework for factors affecting CSIRT performance (Y. Lee & Lee, 2004). While this framework may need to be updated, it presents a useful perspective for viewing CSIRTs in terms of specific variables that are critical to performance. Finally, a more recent study proposed human-automation teaming as a potential path in improving CSIRT performance (Lathrop, 2017), and provides solid ground work for expanding upon specific elements of HAI, such as visual presentation of information.

2.7. The CSIRT as a System

Perspectives from systems engineering (SE) can be useful when studying CSIRTs in a broader context, and allows for incorporating more aspects of process and organization. This section describes how the systems engineering research tradition is valuable to CSIRT investigation.

This dissertation adopts from (Meadows, 2008) the definition of a system as *"an interconnected set of elements that is coherently organized in a way that achieves something"* (p.11). A simple system is generally comprised of inputs, a transformation process, outputs, feedback loops, and exists within a specific environmental context, which provides one-way inputs that affect the system in some way. More often than not, definitions of a system also include the idea that *the whole is greater than the sum of the parts,* indicating that the behaviors and outcomes of a system cannot generally be prescribed or calculated based on the known inputs and intended transformation process. The key to understanding a system is in identifying the elements, relationships between them, and the resulting behaviors that manifest as a result. This dissertation uses this definition and asserts that a CSIRT can be described as a system, and will act as the subject of analysis using systems approaches.

SE literature commonly recognizes the human components of systems; this perspective may help overcome the institutional and psychological barriers in security science (National Academies of Sciences Engineering and Medicine, 2017) that currently inhibit application to social and behavioral sciences in the cyber domain. As the field currently relies on human expertise and effective information sharing to maintain stable operations, the human operators are major constituents of the system. However, as threats grow in number and complexity, humans also are considered a hindrance in increasing the speed of cyber response, as they have relatively limited capabilities in fast computation of large datasets (Singh & Nene, 2013). Conversely, the value of humans in the cyber context is that they can perform fast, accurate pattern recognition with which computer algorithms cannot currently compete (Rutkin, 2015). These different perspectives in literature highlight some of the tensions between the value and drawbacks of humans as the central component to incident response.

One SE framework, further discussed in Chapter 3, is the neurocybernetics framework. This framework uses system organization and information flows to aid in understanding the operation and stability of multi-tiered organizations. Not only does this framework accommodate the human component of larger security organizations, it also helps identify processes and policies that might inhibit their success. This framework can also help separate technology-centric and technology-agnostic issues, which is especially useful given the current separation of research traditions, and tensions between humans and technology in security science. This dissertation explores CSIRTs from the cybernetics perspective to investigate how structure, information flow, and policy impact operations at the lowest level.

In summary, SE perspectives can help overcome current tensions in cyber security research by capturing the complex and dynamic nature of CSIRT operations. Systems approaches employed in this dissertation highlight the value of the human component and help develop solutions around this important constituent instead of independent of it.

2.8. Summary

The increasing realization of the cyber security problem has led to increased demand for CSIRT candidates, which traditional channels of talent are struggling to meet. Furthermore, hiring practices are starting to incorporate non-traditional attributes into candidate qualification requirements; this indicates that the needs of the field are also evolving. In essence, the definition of the "right expertise" is evolving, and there are not enough people who have 'the right expertise' to fill the gap.

With current pressures on the workforce, businesses and researchers are turning to automation for reprieve. Automation development for the future of cyber security must consider the history of CSIRTs, current trends, and expected changes, such that the technology can be just as adaptive as the CSIRTs themselves. Opportunities are currently being explored to introduce automation into computer incident response, with the goal of filling the expertise gap, specifically in task execution. Other critical candidate areas for automation, such as information sharing and task coordination, are currently uncharted. These potential areas for development rely on human factors approaches to provide expertise in understanding human operators, the operational context, and the implications of those factors on system design. This dissertation focuses on operational, organizational, and technological needs at the CSIRT team level in order to holistically identify requirements for future automation development of information sharing functions.

CSIRTs are teams of first responders. They fill a critical role in network defense, and since their inception have evolved into multi-team systems embedded in complex environments. As technology advances, threats emerge, and qualifications change, it is critical to recognize dynamic needs of CSIRTs moving forward. Many research approaches have explored technological approaches, but do not always consider aspects of human factors for defenders. Government officials and industry executives are now becoming heralds of the idea that humans are at the heart of the system and should be part of the solution. Research and operational examples (such as chess) have suggested that the solution likely involves humans and computers working together as a team. While some human factors research in cyber addresses aspects of cyber operations, it largely lacks approaches for human-automation integration, especially from a top-down perspective that could provide useful results for development. This dissertation provides groundwork for a hybrid approach to closing this research-based expertise gap by employing diverse frameworks and methodologies aimed at translating contextual information to actionable design considerations.

CHAPTER 3. RESEARCH DESIGN

The following chapter presents the research conceptual framework, research questions, and the specific methodology used to explore the problem described in Chapter 1. Section 3.1 presents relevant interdisciplinary frameworks and methodologies employed in this dissertation. Section 3.2 describes the research questions that are explored through three separate but progressive studies that build upon each other. Finally, Section 3.3 goes deeper into each method in relation to the research questions and study approaches.

3.1. Research Conceptual Framework

The intersection of human factors and cyber security research, specifically regarding information sharing in CSIRTs, presents compelling questions, some of which this dissertation will explore. The lens through which this research views the problem draws on frameworks and methodologies from expertise and systems engineering literature. The research conceptual framework is organized in accordance the widely adopted Framework Methodology Application model presented in (Checkland, 1985; Checkland & Scholes, 1990). The following sections present the frameworks and methodologies used in this dissertation, shown in Table 2.

Table.2. Frameworks, Methodologies, and Application

Framework	**Methodology**	**Application**
Organizational Ethnography	Ethnographic Field Research	Computer Security Incident Response Teams (CSIRTs)
Six Dimensions of Expertise	Cognitive Task Analysis	
Neurocybernetics	Viable System Model	
Systems Architecture	Workflow Diagramming	
	Needs Analysis	
	Functional Architecting	

3.1.1. Frameworks

3.1.1.1. Organizational Ethnography

Organizational ethnography marries organizational science and anthropology to study groups of people in their contextual settings (Gaggiotti, Kostera, & Krzyworzeka, 2017). Some literature argues that ethnography, like other qualitative approaches, may not explicitly define theoretical foundations (Green, 2014), as the goal of the research is to build understanding through observation and interaction without assuming a priori a theoretical framework. Though not traditionally defined as a theoretical or conceptual framework, some scholars argue that ethnography is more than a method (Gaggiotti et al., 2017; Van Maanen, 2011), and can also act as an analytic perspective for guiding more holistic approaches to research. As this dissertation makes some assumptions about organizational influences in CSIRTs, organizational ethnography is considered a conceptual framework applied in security operations. The methodology of this dissertation includes some methods from organizational ethnography when viewed from the conventional perspective. This will be further discussed in Section 3.1.2.1.

Organizational ethnography is a framework that guides the researcher to objectively approach a group in their natural setting to better understand sociological and organizational forces. In the context of this dissertation, the organizational ethnography framework was applied to understand how these factors impact operational processes at the CSIRT level of an organization.

3.1.1.2. Dimensions of Expertise and Related Frameworks

Expertise has a range of definitions and research approaches. The definition of expertise adopted here is the ‘extent and organization of knowledge and special reasoning processes to development and intelligence’ (Garrett, Caldwell, Harris, & Gonzalez, 2009; Hoffman, Feltovich, & Ford, 1997 p.454). Examples of traditional research in expertise range across domains, which include chess, music, sports, medicine, and academic areas, such as physics (Ericsson & Smith, 1991). However, foundational works in expert research recognize other types of expertise beyond these traditional definitions of domain expertise (Scardamalia & Bereiter, 1991).

Garrett et al. (2009) propose a multi-dimensional approach to expertise in group contexts, defining six (6) dimensions in which a person or entity can have expertise. Within the cyber security domain, this approach allows for better understanding of the multi-faceted human analyst who needs skills in multiple areas, not just a traditional knowledge domain, in order to do her job effectively. The conceptual framework of this dissertation applies the dimensional expertise lens as it relates to information sharing activities. Definitions of these dimensions include:

- *Subject matter expertise*, or traditional domain expertise. An individual may have varying levels of expertise across multiple domains. Diverse teams include individuals with different domain proficiencies, as well as individuals with low subject matter expertise altogether (novices). Within the context of information sharing, this dimension describes the "what" that is being shared.
- *Situational context expertise*, which supports the individual's situation awareness. This includes knowing how to apply information based on the situation in which one is working. This dimension represents the "when" and "why" of information sharing activities.
- *Interface tool expertise*, or familiarity and proficiency with the system interfaces an individual uses in his or her work. Again, this can have varying degrees of expertise across multiple tool interfaces. While not directly applied in this research, interface tool expertise could contribute to the "what" or "how". However, this dimension is less of a focus for the scope of this dissertation proposal.
- *Expert identification expertise*, which involves knowing to whom to go for knowledge that an individual herself does not have. In diverse incident response teams, this dimension is especially important for quickly mitigating a problem (Steinke et al., 2015). Closely related to another expertise framework called transactive memory theory, this dimension represents the "to or from whom" aspect of information sharing.
- *Communication expertise*, commonly termed 'communication skills'. This involves knowing how to articulate information to other humans. Communication expertise is directly related to teamwork, organizational culture, and information alignment. This dimension describes "how" information sharing occurs.
- *Information flow path expertise*, or knowing which modality of communication is appropriate, effective, and efficient. Like communication expertise and situation context, information flow path expertise relates to "how", "when", and "why" information is shared.

The six dimensions of expertise framework presents solid groundwork for a holistic approach to understanding information sharing expertise in CSIRTs, but lacks its own methodology for identification and quantification of these dimensions in context. However, due to the high applicability of dimensional expertise at the analyst and team lead levels, this framework advocates further investigation and development, specifically for operator tasks and team lead responsibilities. Literature from other domains support this framework by providing specific methods for studying or quantifying dimensional expertise (Steinke et al., 2015; Yuan, Fulk, Monge, & Contractor, 2010), which are discussed in Section 3.1.2.

3.1.1.3. Applying Expertise Approaches to Cyber Security

Previous studies have worked on mapping expertise to tasks or job requirements in different domain settings. Within cyber security specifically, Chen et al. (2014) performed different task analyses and linkage analysis to understand expertise gaps in CSIRTs. The NICE framework also details a comprehensive inventory of knowledge required for different positions in cyber security (Newhouse et al., 2017). Another study produced knowledge maps from surveys to understand how well equipped CSIRTs were to perform daily tasks (Steinke et al., 2015). While not nestled within cyber defense, there have also been attempts to create a visual matrix representation of how expertise maps to other dimensions of a space, such as the Accreditation Board for Engineering and Technology outcomes in engineering education (London, Caldwell, & Patsavas, 2013). These examples illustrate how correlating functions with expertise can provide an understanding of knowledge gaps in groups. One goal of this dissertation is to further explore expertise needs of CSIRT analysts using the six dimensions of expertise framework with the goal of expanding the current list regarding context and technology.

Considering the current rate of technological development in cyber operations, augmenting human expertise in cyber operations with automation is another approach worth exploring. Work by Hoffman et al. (2008) has provided some preliminary groundwork in human-automation system design using expertise approaches with some level of success in other domains. Frameworks and methodologies from human factors were applied to determine expertise and work context, then divided at different levels of function. This methodology proved successful in addressing a gap between human expertise in context and system design for the respective work tasks. This dissertation employed similar methods as (R. R. Hoffman & Deal, 2008) within the CSIRT context to bridge SE concepts with expertise frameworks.

As mentioned in Chapter 2, most incident response functions in cyber defense are performed on teams (CSIRTs). These work groups are assembled based on certain applicable KSAs, then asked to perform decision-making tasks regarding an appropriate response or course of action for a particular incident. While there is a long list of factors that can affect team effectiveness (Tannenbaum, Beard, & Salas, 1992), two of these factors are of particular interest to this research problem: knowledge and information sharing. Expertise or knowledge (often included as part of KSAs) is commonly listed as an important factor in team effectiveness (Cooke, 2004; Paris et al., 2000). Human factors tools have been developed to elicit the knowledge needed on a team in order for the group to perform well (Burke, 2005), which can then be used to direct training or hiring initiatives. However, knowledge by itself will not help the group in performing their tasks. Coordination of expertise within the team, a subset of information sharing, is equally important in order to deliver the knowledge needed to the individual who needs it (Caldwell, 2008; Klinger & Hahn, 2004; Mesmer-Magnus & DeChurch, 2009). Group dynamics literature notes that decision-making tasks, such as those seen in incident response, often rely on knowledge distributed throughout the group, and that this knowledge is often inefficiently or ineffectively applied (McGrath, 1984).

With the acknowledged shortage of labor and the widespread availability of technology, these information sharing functions provide a ripe opportunity for automation development. Expertise related to information sharing at different levels of cyber operations could create a human-automation team centered on improving information coordination within and between components of the CSIRT. The approach presented in this dissertation asserts that, in order to develop human-automation teaming effectively, developers must have better design requirements, specifically regarding relevant information sharing functions. To derive these improved design requirements, researchers must first understand the current landscape of incident response from the view of the individuals that perform the job: the humans. More specifically, it is useful to know *what tasks* these analysts perform, *how* they perform them, *what they need to know* in order to do so, and *how* *they know how to respond* in specific situations (Caldwell, 2002; Garrett & Caldwell, 2011). These questions will be explored within the scope of this dissertation through field research to help create a richer understanding of CSIRT contexts and information sharing needs across different teams.

Additionally, researchers might be interested in the same information, but with a focus on inter-organizational communication and coordination. Organizational psychology and macroergonomics research have indicated the potential effects of these larger scale factors on team performance (T. R. Chen et al., 2014; Cuevas, Fiore, Salas, & Bowers, 2002). From this perspective, scientists can understand what aspects of successful information sharing are missing in CSIR, or what could be augmented to improve information sharing in the overall CSIRT system. In order to address this perspective of research, this dissertation collects some of this data with the overall goal of providing insights and functional requirements from CSIRT data. The results are aimed at managers, to help describe observed connections between organization and CSIRT operations, and at computer scientists and machine learning developers to help draw attention to organizational factors affecting human-automation teaming.

3.1.1.4. Systems Engineering (SE) Frameworks

Multiple approaches are often warranted for analyzing and improving system function (M. C. Jackson, 2000). Systems engineering provides approaches that are general and flexible in incorporating psychological and organizational factors of human-system interaction across levels of a system’s architecture (DeGreene, 1970). Two major frameworks from SE literature are presented below. Section 3.1.1.3.1 provides background on neurocybernetics, which was briefly introduced in Chapter 2. Section 3.1.1.3.2 presents a framework for system architecture development, which is applied to guide scope and process of determining functional requirements for potential system-based solutions.

3.1.1.4.1. Neurocybernetics as an Organizational Framework

Neurocybernetics is a branch of systems engineering that explores mapping information flows in the human nervous system. Subsequent work in neurocybernetics provides an organizational framework and methodology that can be generalized for analyzing other complex, multi-tiered teams (M. C. Jackson, 2000). Beer's "viable system model" (VSM), depicted in Figure 4, is one such cybernetic approach, which shows an organizational structure and information flow schema that reflects the nervous system functions in the human body. VSM describes a hierarchical representation of how systems coordinate together to remain stable, or 'viable', within their larger environment. As shown in Figure 4, the five-element system is broken down into implementation (System 1), coordination, control, development, and policy (System 5); channels through which information can flow connect these systems. Viability within the larger environment is established through efficient and effective information flow between these systems, and between the larger hierarchical system and the environment.

The neurocybernetics framework can be easily applied to human organizations, which are multi-team systems with similar organizational schemes (e.g. workers, supervisors, managers, directors, and executives). Considering that CSIRTs are commonly organized groups within a larger organization, neurocybernetics is a useful framework for understanding the operations and inter-team coordination of security groups in terms of stability and performance. The organizational scheme of a CSIRT can vary across companies and sectors. VSM within the neurocybernetics framework offers a methodology that can accommodate this variation and complexity. The methodology includes a diagnostic tool for determining information flows that will be useful in understanding information sharing in cyber operations and across various collaborating entities. Further description of the VSM as a methodology for system analysis is provided in Section 3.1.2.3.

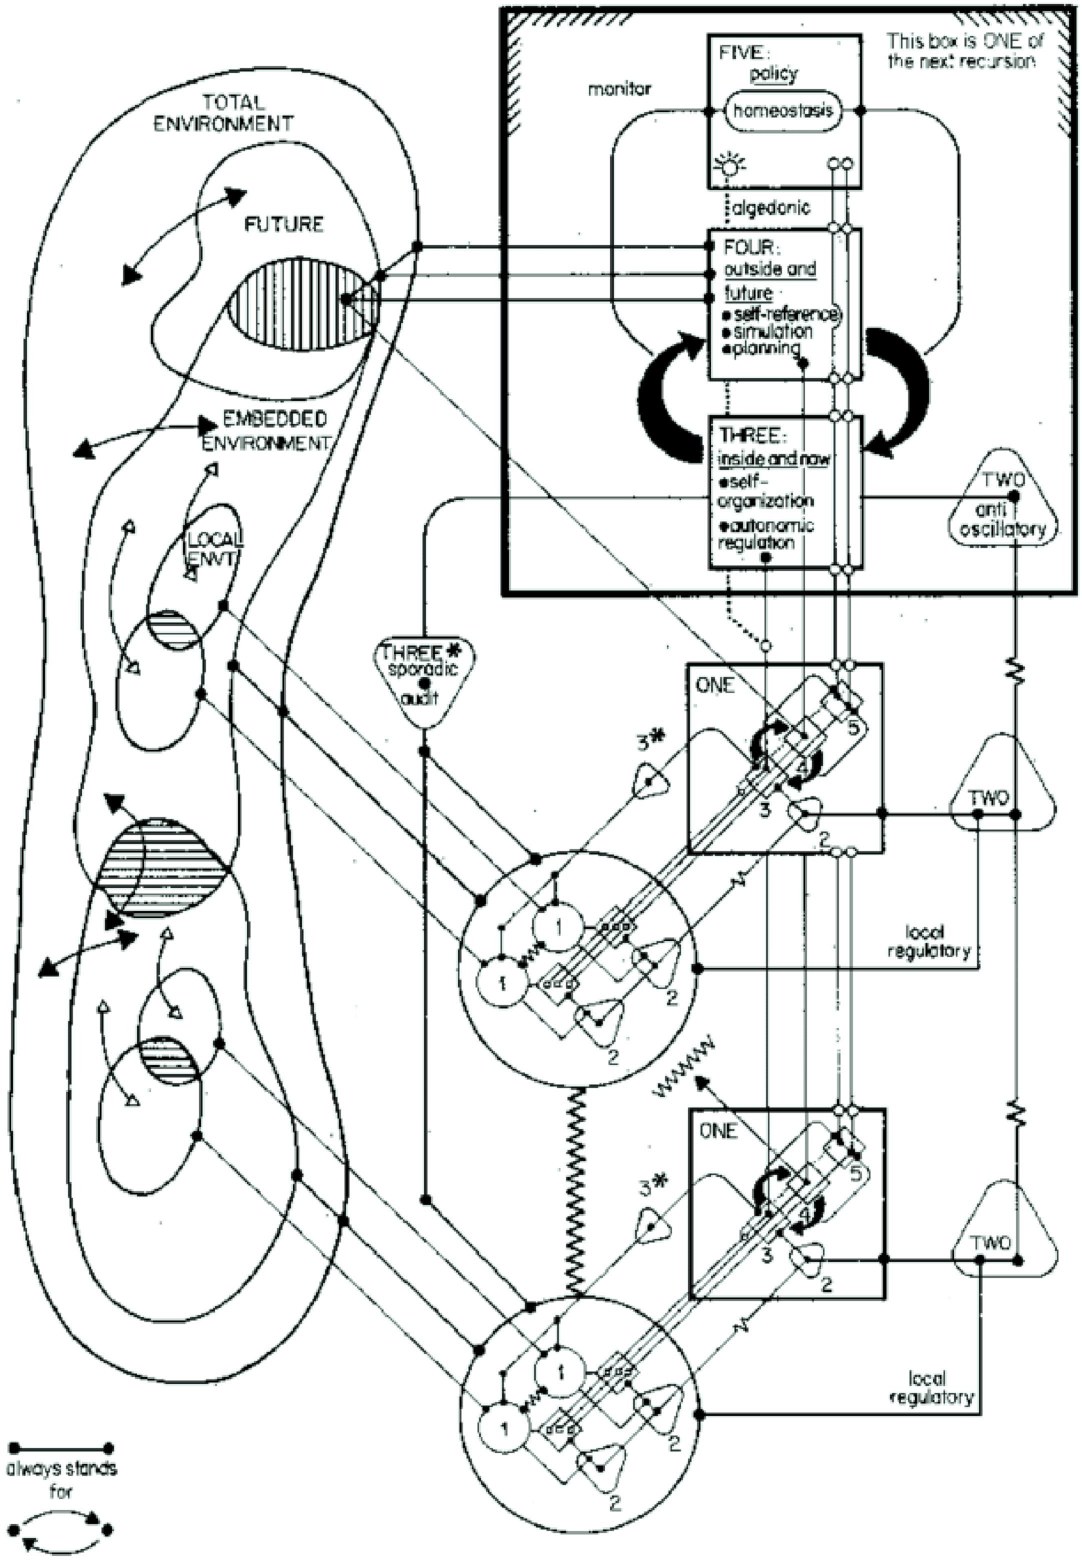

Figure.4. Beer's VSM (Beer, 1995)

3.1.1.4.2. Systems Architecture Framework

Systems architecting within the context of SE project management offers architectural and development methodologies for system advancement and improvement. The project management perspective on SE (Caldwell, 2009) incorporates developing engineering requirements and overall system design, as well as lifecycle management and task coordination during the actual building of the system (*NASA Systems Engineering Handbook*, 2007). Project management provides a useful, holistic approach for viewing the human-automation CSIRT problem in terms of multiple stages of design and development. More specifically, the research conceptual framework draws on the multi-level, multi-stage considerations of project management to analyze human CSIRT teams performing incident handling and produce system requirements for consideration during design and development of automation for these teams. Essentially, project management concepts contribute to framing not only the methods and analyses of this research, but also potential solution spaces.

This dissertation employs a framework developed within the project management perspective. Oftentimes, SE literature addresses system development with a lifecycle approach that allows for gradual increases in development over the course of design, use, and retirement of a product or system (*NASA Systems Engineering Handbook*, 2007). The first stages of this lifecycle include concept development and architecture development, which essentially determine the goal of the system and interpret how the system should reach that goal. One useful framework for applying the first stages of the lifecycle approach to information systems specifically is presented in Levis & Wagenhals (2000) and depicted in Figure 5.

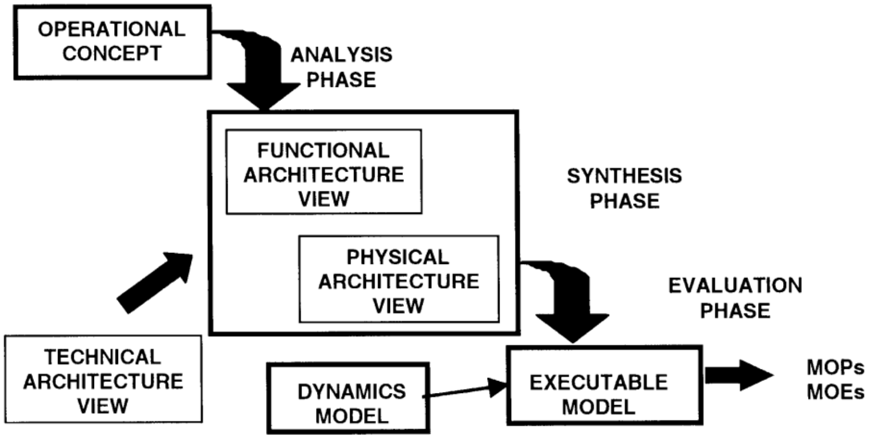

Figure.5. Three Phases of Architecture Development from (Levis & Wagenhals, 2000)

The concept of *architecture development* is the construction of functional and physical requirements, tradeoffs, interfaces, risk management, verification, and validation (*NASA Systems Engineering Handbook*, 2007). The old adage of *"form follows function",* adopted from the architecture domain for systems architecture principles, indicates the importance of first determining how the system should behave, then deriving the physical components to manifest those behaviors. Levis & Wagenhals (2000) describe the lineage of system architecture, along with a framework for developing architectures for information systems. Referencing this framework, the functional architecture in particular "defines the transformations of input flows into output flows performed by the system to achieve its mission" (SEBoK, 2017).

The functional architecture will focus on specific, defined functions after stakeholder input and concept development, as opposed to an "outsider view" of the entire system of interest. Systems architecting frames potential solution spaces involving automation, provides a structure for presenting functional requirements for those solutions, and ultimately helps direct development of system behaviors.

3.1.2. Methodologies

3.1.2.1. Ethnographic Field Research

One goal of this dissertation is to gain a better understanding of the variety of contexts in which CSIRTs operate to help guide automation design considerations. Qualitative research is one class of methodologies available to explore such contexts. Qualitative research includes a broad set of techniques and methods that are generally focused on and conducted in the field (Merriam & Tisdell, 2016); this approach aims to develop a better understanding of a particular problem or phenomenon within the context and respective interactions (Merriam & Tisdell, 2016; Patton, 1985). More importantly, it allows for a researcher to explore that problem or phenomenon without having a hypothesis (Auerbach & Silverstein, 2003). In employing a qualitative methodology, a researcher can learn more about what questions are relevant within that problem space and context.

Field research is a qualitative research methodology that includes observation and interview methods (Pelto, 2016). Field research is often associated with ethnography, and is useful in obtaining a rich understanding of context and organizational behavior that is otherwise not captured using quantitative methods (Schwartzman, 1993). User-centered data collected directly from the field is invaluable when considering highly detailed information about the context in which users engage with their environments and conduct their daily tasks (Holtzblatt, 2016). Thus, the field research methodology was employed to assist in constructing knowledge about this context in CSIR organizations with the goal of informing new design requirements.

3.1.2.2. Cognitive Task Analysis

In order to apply the expertise framework described in 3.1.1, this dissertation uses a type of task analysis to identify expertise needed to within the application of CSIRTs. Task analysis is a collection of methods that aim to ‘collect, classify, and interpret data on human performance in work situations’ (Annett & Stanton, 2000 p.1). There are multiple approaches within task analysis, such as hierarchical task analysis (Annett & Stanton, 2000) and goal-directed task analysis (Humphrey & Adams, 2011). Cognitive task analysis (CTA) is a subset of methods used to assess the knowledge and cognitive activities needed to perform a particular set of tasks based on subject matter expert experience (Crandall, Klein, & Hoffman, 2006). Particularly, CTA helps determine what knowledge an expert applies, when she applies it, and how she knows which knowledge to apply.

Literature indicates that CTA is the best suited for expertise-aimed studies interested in understanding what knowledge is required and how it relates to the overall task structure (Crandall et al., 2006). CTA is especially appropriate for developing technological solutions to support cognitive processes, and has been performed in cyber security to understand and improve team effectiveness (T. R. Chen et al., 2014) and situation awareness (D'Amico et al., 2005), as well as inform system design (R. R. Hoffman & Deal, 2008). Within the context of this dissertation, CTA was used to understand expertise requirements for potential automation development.

3.1.2.3. Beer's Viable System Model (VSM)

In the case of security operations, CSIRTs are embedded in larger organizations and are driven by the mission and policies of those parent organizations. Cybernetic analysis allows researchers to identify how CSIRTs coordinate response throughout an organization and adapt to the external security environment. Furthermore, cybernetics can reveal to managers potential policy and procedure opportunities for improving system function that do not require complex technical solutions. This dissertation focuses on one major methodology for cybernetic analysis: VSM.

As mentioned in Section 3.1.1, Beer's VSM has been proposed as a methodology within management cybernetics for mapping information flows in human organizations. Detailed in (M. C. Jackson, 2000), applying the VSM includes two segments, which include system identification and system diagnosis. This general procedure closely resembles that of the workflow diagram (Section 3.1.2.1), but adds in aspects of the environmental context that are valuable in understanding external system inputs and how they affect the system of focus, as well as constraints and feedback (e.g. accountability).

The first segment from (M. C. Jackson, 2000) prompts the researcher to determine the purpose of the system to be investigated and its respective viable parts. This step scopes the system of focus, as well as the larger system to which it belongs, and the environment in which it exists. The second segment uses cybernetic principles to explore specific aspects of the system, such as its environment, inputs, outputs, feedback loops, performance measures, and even conflict resolution between components and other systems. The second segment is also supplemented with specific questions for different levels of systems within the organization to help draw out these elements. The outputs of this method include (for each organization studied) a diagram and a description of different aspects of the organization from the diagnostic questions in (M. C. Jackson, 2000).

3.1.2.4. Workflow Diagramming

Workflow diagramming is a classic industrial engineering methodology that has been used for over a century to analyze and improve processes (Gilbreth & Gilbreth, 1922; Taylor, 1911). Creating a workflow diagram involves abstracting a process into an illustration that indicates process flow, including components of the process, such as information and resources, as well as aspects of the process, such as transportation and delay (Graham, 2004). There are many types of workflow diagrams, such as process maps used in business (Damelio, 2011) and data flow diagrams used in software engineering (Adler, 1988). These visual representations of processes or aspects thereof are useful for creating a shared understanding of sequence of events and potential areas for improvement (Gilbreth & Gilbreth, 1922; Mclaughlin, Rodstein, Burke, & Martin, 2014). This dissertation utilizes the process map methodology from (Damelio, 2011) as the sole methodology for workflow diagramming. Developing an understanding of the incident response (IR) process helped determine and focus the needs of analysts for the following Needs Analysis.

3.1.2.5. Needs Analysis

In order to help determine a path forward for new systems or system features, SE research and practice often employs Gap Analysis or Needs Analysis (see below) to better define gaps, user needs, and market opportunities. This section describes the difference between these two methods, and presents the background and justification for using the Needs Analysis for concept development within the context of this dissertation. As mentioned in Section 3.1.1, the Needs Analysis methodology provides the background and support for the functional architecting framework.

A *gap* is defined as the difference between current state and desired future state (Langford, Franck, Huynh, & Lewis, 2008), and can incorporate operational deficiencies or disadvantages in completing the mission at hand. Gap Analysis is the method of determining these differences, which is often done within a firm to determine future strategy. In essence, this process helps define "where do we want to go", after which the firm can further determine "how do we get there". One potential misuse of Gap Analysis is that, if the future state is defined by where an adversary or competitor is currently, it can mislead firms into becoming a lagging player in the overall environment or market (Langford et al., 2008). This is similar to the concept of "mirror chess" in cyber security (Curry, 2019), and translates to only being as good as your enemy, and always playing defense. Thus, when performing Gap Analysis, it is critical to define the future state by where the market or firm needs to go to advance the status quo, not to be as good as another actor.

In order to increase systematic practices and conclusions, gaps are often quantified in finance and engineering. Metrics to further assess risk, worth, and value have been proposed within the Department of Defense literature on Gap Analysis to produce more robust assessments (Langford et al., 2008). Multiple metrics become part of these assessments, including effectiveness and performance. However, gaps must first be identified before being quantified in this manner. Quantification efforts also may or may not be useful depending on the setting and goals of the study, and the relative maturity of the respective technology in the field.

Within SE literature, the conceptual phase of system develop uses a Needs Analysis (Kossiakoff & Sweet, 2003) to overcome the hurdle of quantifying specific gaps. That is, a Needs Analysis assesses operational deficiencies against technological opportunities to narrow the scope of a system and define how the system can meet the needs of the field. According to (Kossiakoff & Sweet, 2003, p. 124), a Needs Analysis is a three-phase process that includes an Operations Analysis, Functional Analysis, and Feasibility Definition.

Other systems engineering sources refer to this in terms of the *Concept Development Phase* for Systems Engineering projects. The major output of this is a Concept of Operations (CONOPS) (ISO/IEC, 2011; Office of the Deputy Under Secretary of Defense for Acquisition and Technology: Systems and Software Engineering, 2008), which defines the current system, gaps in terms of operational needs or shortfalls (*AIAA Guide to the Preparation of Operational Concept Documents (ANSI/AIAA G-043A-2012)*, 2012), and a high-level description of a new system, including how it meets user needs in example scenarios. A CONOPS document helps in developing operational requirements (Cellucci, 2008) that precede technical requirements.

Despite the difference in terminology, the Needs Analysis as described by Kossiakoff aligns with other literature regarding CONOPS in the sense that there is a process for determining user needs and developing a new concept based on those needs. The approach employed in this dissertation utilizes both sets of literature to propose new concepts for human-automation teaming in CSIRTs. Figure 6 below depicts the major deliverables of the Needs Analysis as conducted in this dissertation in relation to the final method of systems architecting.

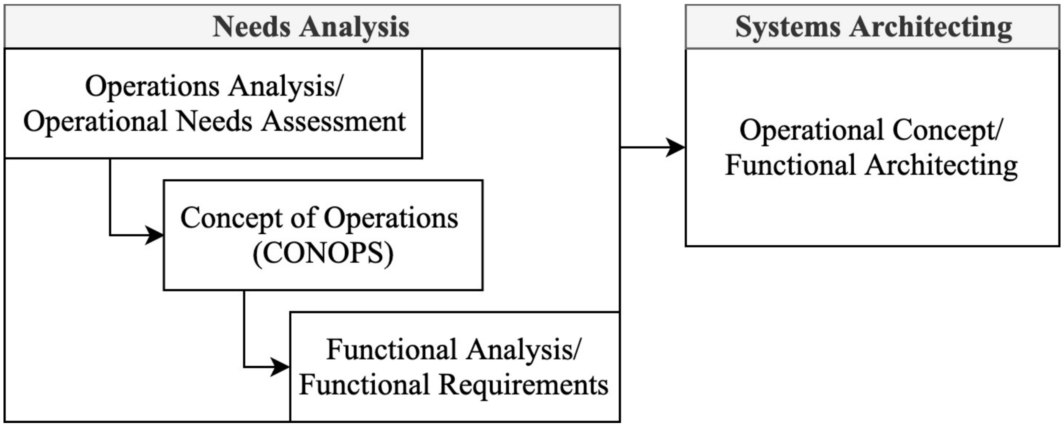
Needs Analysis
Operations Analysis/ Operational Needs Assessment
Concept of Operations (CONOPS)
Functional Analysis/ Functional Requirements
Systems Architecting
Operational Concept/ Functional Architecting

Figure.6. Needs Analysis and Systems Architecting

In cyber security, quantifiable metrics have started to appear in sales materials for security products. Performance measures include mean time to detect and mean time to respond, which are essentially aggregate indicators of the entire response system’s performance. Because these measures are aggregates, specific aspects of performance in different parts of the process may be lost, and underdeveloped effectiveness indicators may be eclipsed by measures of performance. Quantifying gaps in cyber security may not be a useful an activity due to the current maturity level of the field. However, identifying the gaps framed at conceptual and strategic level may help guide future direction of research and development. Thus, a Needs Analysis was employed to address industry gaps as well as answer the research questions in this dissertation.

3.1.2.6. Systems Architecting

The outputs of the Needs Analysis serve as the foundation of the new operational concepts borne from CSIRT needs. These new concepts were developed using systems architecting. Systems architecting is an inductive process largely aimed at determining qualitative aspects of a system versus measurable components (Maier & Rechtin, 2000). Systems architecting is less scientific and structured than the name implies (Maier & Rechtin, 2000), which has resulted in a variety of methods for approaching architecture development. Each standard and method has advantages and disadvantages, thus the method should be chosen based on aspects of the problem being addressed (Levis & Wagenhals, 2000). Levis & Wagenhals (2000) propose several methodologies for architecting information systems in particular, from the determining the operational concept to validating executable models of the architecture for integration and long-term deployment. The research proposed in this dissertation will address two aspects of this process: developing the operational concept and developing the groundwork for the functional architecture. Physical aspects of the architecture would require inputs and expertise from computer scientists who are able to properly determine appropriate algorithms to execute the desired functions.

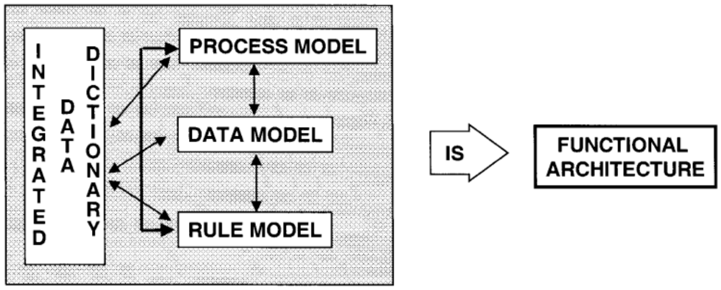

Figure.7. Functional Architecture Components from (Levis & Wagenhals, 2000)

The methodology presented in Levis & Wagenhals (2000) provides an adequately abstracted and iteration-based guide for developing new systems. First, the operational concept is derived from stakeholders at multiple levels of the system, including operators and managers. Next, the functional architecture is developed. Figure 7 shows the components of a functional architecture as described by Levis & Wagenhals, which include an overarching activity model, a process model, a data model, and a rule model. Combining these models forms an integrated data dictionary, which is a complete set of functions, states, rules and data needed for the system to meet its mission.

3.2. Areas of Interest and Research Questions

This dissertation investigated three major areas of interest: understanding workflow and expertise in CSIRT operational context, identifying automation opportunities, and determining functional requirements for future automation. The research questions below explore these areas of interest, and are divided into sub-questions that can be more easily answered across different methods. Answering the fourth and final research question helps translate findings from the first three questions into application.

3.2.1. Understanding Workflow and Expertise in the CSIRT

As there is limited information about the information sharing expertise profiles (and variations thereof) in CSIRTs, a portion of the research protocol is dedicated to investigating high-level aspects of team information sharing activities during routine and non-routine tasks in different CSIRT settings. Identifying critical functions in the workflow and understanding information sharing in real-world CSIRT context are also goals of this section of research. The follow research question aims to address these goals in three different CSIRTs.

[RQ1] Which dimensions of CSIRT team expertise in information sharing are required for a team to perform critical incident response functions individually and in coordination?

This research question can be broken down into two sub-questions that aim to address the two aspects of the scope of RQ1: workflow and expertise. These questions will be answered through different methods (described in Section 3.3), and help narrow the focus of the problem space.

[RQ1.1] What are the critical incident response functions with respect to information sharing performed by different members of each team?

[RQ1.2] Which dimensions of expertise, as they relate to information sharing, are required for tasks performed by the team members?

3.2.2. Identifying Automation Opportunities

After the teams and their respective processes have been sufficiently investigated in terms of information workflow and expertise, the researcher can start to determine the potential information sharing needs of CSIRTs that can be met with future automation development. The first step in this process is to develop understanding the current state of automation, where it persists in the process, and its effectiveness in CSIRT operations. The second research question explores the current state of automation (including new automation opportunities) in the teams being studied and the existing market of automation technologies.

[RQ2] What automation solutions might make sense to enhance the range of areas of information sharing and expertise on a CSIRT?

The sub-questions of Research Question 2 addresses automation from two different angles: existing automation, and future automation.

[RQ2.1] What automation currently exists for information sharing functions, such as handoffs?

[RQ2.2] Where do CSIRT members struggle the most in the incident response (IR) process?

[RQ2.3] What capability gaps currently exist in automation technologies with respect to user-identified needs in dimensions of expertise?

3.2.3. Determining Functional Requirements

After potential areas of automation have been identified, it is possible to examine how those functions are performed in practice. Further exploration of analyst needs help narrow the solution space for developing new automation to enhance information sharing and expertise in CSIRTs. Answering the third major research question provides guidance for future development.

[RQ3] What are the functional requirements of potential automated solutions, based on the human expertise and task elements of CSIRT response demands?

As with the previous research questions, Research Question 3 is broken down into four sub-questions to be answered through different methods. Investigation within different CSIRTs will provide a broader sample of considerations for requirements, which will be supplemented by CSIR expert inputs.

[RQ3.1] What expertise is needed to perform the task effectively?
[RQ3.2] What other information flows are needed to perform the task?
[RQ3.3] What are the expected outputs of the task?
[RQ3.4] Who uses the outputs, and how could the outputs be improved based on current processes?

The fourth and final research question investigates how the answers to the previous RQs might address the CSIRT problem space.

[RQ4] What is a feasible path forward in translating the collected insights and functional requirements into actionable solutions?

To answer this research question, one of the areas of need identified from RQ2 (supplemented by the functional requirements identified in RQ3) was selected as a use case for development. The use case draws on SE methodologies commonly used in industry and demonstrates the value of this overall approach. The results provide guidance for future system architectures that could aid in information sharing in incident response.

3.3. Methods and Tools Overview

The non-experimental research in this dissertation is composed of three main studies: an exploratory field study of current CSIRTs, a Cognitive Task Analysis study with CSIR experts, and a Needs Analysis for conceptual development of future automation. Methods employed include observations and semi-structured interviews. Tools include workflow diagrams, Needs Analysis and systems architecting. As each method addressed more than one research question, a relational matrix (Figure 8) depicts direct and indirect contributions of methods to research questions.

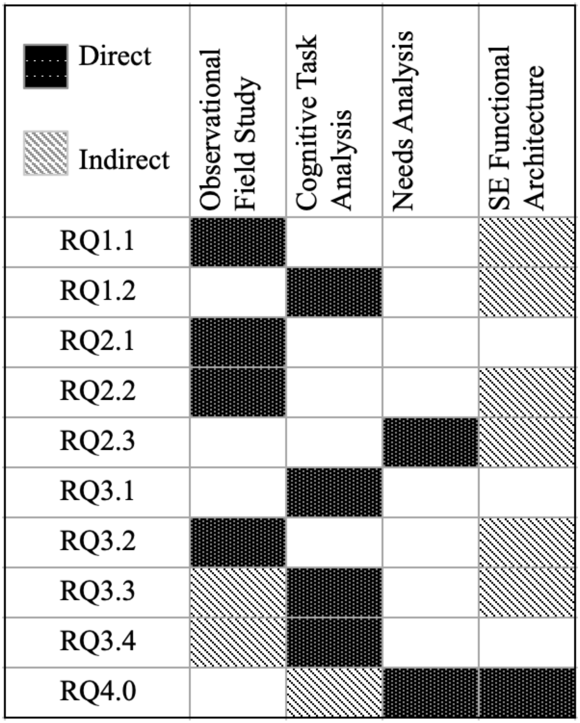

■ Direct ▨ Indirect	Observational Field Study	Cognitive Task Analysis	Needs Analysis	SE Functional Architecture
RQ1.1	■			▨
RQ1.2		■		▨
RQ2.1	■			
RQ2.2	■			▨
RQ2.3			■	▨
RQ3.1		■		
RQ3.2	■			▨
RQ3.3	▨	■		▨
RQ3.4	▨	■		
RQ4.0		▨	■	■

Figure.8. Relational Matrix of Research Questions to Methods

The order of the studies described in the following sections is constrained by dependencies between each study. Figure 9 depicts these relationships through the inputs and outputs for each study.

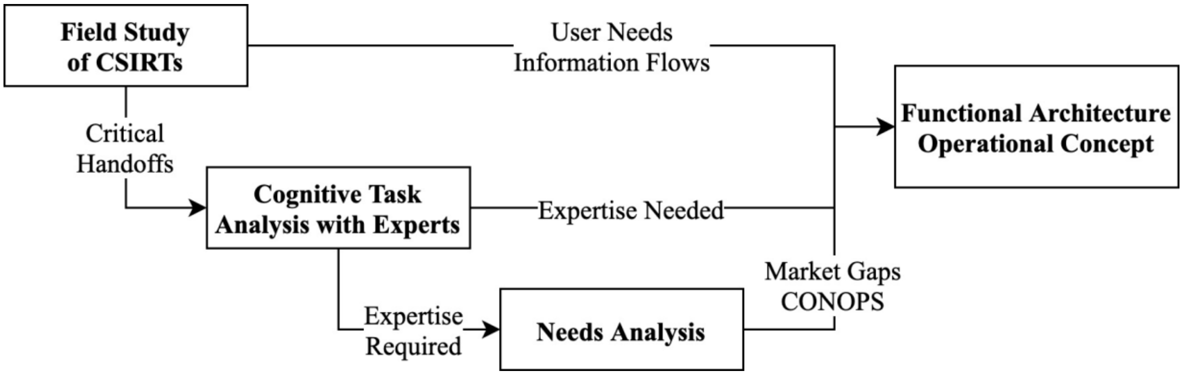
Field Study of CSIRTs
User Needs Information Flows
Functional Architecture Operational Concept
Critical Handoffs
Cognitive Task Analysis with Experts
Expertise Needed
Market Gaps CONOPS
Expertise Required
Needs Analysis

Figure.9. Study Input and Output Relationships

3.3.1. Observations and Informational Interviews in CSIRTs

The first study was an exploratory field study that aimed to reveal significant information sharing processes, variables, and relations among them through a basic qualitative approach (Merriam & Tisdell, 2016). Specific data collection methods employed in Study 1 included observations and process-oriented interviews with experienced employees and CSIRT leaders or managers. The value of Study 1 was threefold. First, by providing rich contextual data, it laid the foundation for the various analysis methods (Kerlinger & Lee, 2000), including workflow diagramming and cybernetic analysis, which help identify larger process and organizational themes across CSIRTs. Second, it helped identify critical functions on which to focus the Cognitive Task Analysis in Study 2. Lastly, data collected in Study 1 was used to illustrate an initial visual model of the information sharing process in CSIR organizations for shared understanding and systematic analysis.

Observations are a non-intrusive way of collecting data from humans, and can be very informative about behaviors in the populations of interest (Gillham, 2008; Kerlinger & Lee, 2000). Specifically within the context of organizations and teams, observations are a powerful way of collecting information regarding interactions and social context (Schwartzman, 1993). Observations were used to understand the incident response process to note potential problem areas via analyst behaviors. This method was also used to help build my understanding of CSIRT operations, as I had not been previously exposed to security operations. Relevant observational methods literature was used to help guide data collection and recording (Gillham, 2008; Merriam & Tisdell, 2016; Pelto, 2016; Schwartzman, 1993).

The informal, conversational interview is another common qualitative method that was conducted with participants to supplement observational findings within the research setting (Pelto, 2016; Schwartzman, 1993). Unstructured informal interviews are appropriate when the goal of the interviewer is to remain flexible and exploratory in inquiry, and learn more about phenomenon and relevant questions to ask in the future (Merriam & Tisdell, 2016). The informal interview also allows participants to provide rich information regarding perspective, context, and background. Informal interviews used in Study 1 aimed to obtain perspective views of information sharing with respect to procedures and protocols and help verify the observed process. The goal of these interviews was to determine a high-level sequence of events and identify critical functions related to information sharing. In order to help ensure coverage of topics in conversation (Pelto, 2016), a checklist of items was employed about which to ask participants, which can be found in Appendix A.

Though evidence points to wide variation between CSIRTs (Ruefle et al., 2014), literature describes a generic picture of CSIRT operations that may not reflect the range of actual structures and processes. A stratified cluster sampling strategy (Pelto, 2016) was adopted to help observe this variation . This study investigated a subset of the CSIRT population, where the unit of study was one team. Though mainly opportunistic, the sampling strategy included teams from different organizations or institutions with some variation in size or specialization. Three (3) teams were investigated as part of this study, which are described in Chapter 4.

3.3.2. Establishing Trustworthiness in Qualitative Research

Reliability essentially indicates consistency within data collection and analysis (Kerlinger & Lee, 2000; Noble & Smith, 2015). However, reliability in the traditional sense is very hard to achieve for qualitative research (Merriam & Tisdell, 2016), and replicability for qualitative research is low. Furthermore, literature has gone so far as to suggest that traditional concepts of rigor (reliability and validity) are inappropriate for qualitative research, and that “trustworthiness” is a better conceptual framing (Cypress, 2017; Lincoln & Guba, 2011). Thus, I adopted trustworthiness as the concept relating to rigor.

Merriam & Tisdell (2016) suggest that trustworthiness can be established through several strategies, including triangulation, documenting the investigator’s position, and establishing an audit trail. Other articles on qualitative research reliability report similar suggestions (Leung, 2015; Noble & Smith, 2015). Reliability of the personal accounts collected was established through *triangulation* of collected data and existing literature (Merriam & Tisdell, 2016). This was also a part of the assessment of saturation, described in Section 3.3.2.1.

3.3.2.1. Investigator Position

The *investigator's position* is a statement that helps establish the researcher's personal bias and lens through which she conducted the research. It is also one technique that aids in reflexivity of the researcher throughout a qualitative study (Birks, Harrison, & Bosanquet, 2014; Noble & Smith, 2015; Sultana, 2007). In order to address investigator bias, qualitative research often includes a personal statement of the researcher's position and potential bias within the context of the study.

At the time of this dissertation, I (the researcher) was a PhD student in my fifth year of study trained in the Industrial Engineering discipline before pursuing human factors. From a paradigmatic perspective, my education was largely based in positivistic philosophy. I had 7 years of prior experience working in industry, which mostly included manufacturing settings and included work-study programs. I also had some experience with software development and evaluation, but had no practical experience in security activities or settings. However, I did have family and friends who had worked in security-sensitive environments (including military, cyber security, and defense contracting).

My goals going into the dissertation were to increase my own understanding of the domain and differences in context for the purpose of addressing gaps or problems evident in the field. I needed to shift my research paradigm to adopt constructivist philosophy in order to achieve these goals. My role was both data collection instrument and researcher, tasked with capturing data from observations and interviews, transcribing and transforming said data (if it was in a different form than could be analyzed), coding, and analyzing. These indicate several stages of interpretation through which I needed to process the data. In addition to the already mentioned experiences and perspectives, I maintained the perspective that rich context collected directly from the environment could be valuable to both the research and her personal understanding.

As a student researcher who made direct contact with participants, I needed to be aware of how my personal and professional background could impact my interpretation of the data. Furthermore, I needed to consider that my position as a researcher in a field very sensitive to sharing information with those considered strangers could have impacted the data, especially during the first study. It was also important that I maintain social distance between participants and myself, which was a difficult balance to maintain considering the camaraderie needed to overcome the aforementioned boundaries. During the first study, I was careful to remain neutral regarding other interviews conducted within each organization, especially as participants tended to share their opinions about peers, managers, and other parts of the organization during the interviews. Being aware of the social and organizational context of the research was also important, as these factors made up the majority of environmental influences for each team.

3.3.2.2. Audit Trail

One important strategy in establishing rigor and trustworthiness in qualitative research is create an *audit trail*, or to document each step of the process (Leung, 2015; Merriam & Tisdell, 2016; Noble & Smith, 2015), including sampling strategies, decisions and definitions in coding and analysis, and reflections throughout the entire study. In doing so, the research illustrates a clear story and provides transparency of how the research was carried out and how conclusions were reached. Throughout this dissertation, details of the aforementioned areas are included to ensure clarity and understanding with the reader, and to help establish the rigor with the audit trail.

Reliability can also be thought of "as the extent to which research can be replicated" (Merriam & Tisdell, 2016, p. 250). Regarding the human as the collection instrument for interviews, reliability can be improved through additional practice, which will produce more consistent results. Before this dissertation, I did not have deep experience in interviewing methods. However, preparation for interviews included a literature review and practice interviews with non-participants. Furthermore, recordings of interviews (in Study 2) were used alongside interviewing techniques to reflect and improve skills in this area between each interview to improve within the timeframe of the study.

3.3.2.3. Data Saturation

Sampling in qualitative research is often driven by the amount of data needed to ensure that the research question can be adequately answered. This is often referred to as *saturation*, or the point at which no new data can be captured from more participants (Fusch & Ness, 2015; Lowe, Norris, Farris, & Babbage, 2018; Merriam & Tisdell, 2016). Past the point of saturation, the amount of effort and resources required to interview more participants outweighs the potential new information gained. Saturation was assessed for the Cognitive Task Analysis in Study 2 to evaluate how many participants were appropriate to include in the study. The Cognitive Task Analysis method is further described in Section 3.3.5.

The goal of the Study 2 was to obtain validation of expertise themes (the six dimensions of expertise and two more identified from Study 1 data) from expert interviews, data collected in Study 1, and existing literature. In order to evaluate if data collection needed to continue, it was first determined what type of saturation was needed to justify the study's conclusion. Saunders et al. (2018) describe four (4) types of saturation. The two most appropriate types of saturation for this study are *data saturation*, in which no new themes arise from further collection, and *a priori thematic saturation*, in which already-identified themes are validated within a particular population.

The six dimensions of expertise (Garrett et al., 2009) act as the 'a priori themes', representing areas or sub-categories of expertise that are required for analysts to investigate and escalate incidents. The framework provides themes identified from theoretical grounding in prior research across expertise literature, which can be validated within a particular domain, and in this case, for a particular task. In addition, two themes identified in Study 1 and triangulated with existing literature, were added to the set for saturation analysis. These two themes regarding *policy* and *self-awareness* are specific to this domain, and did not fit absolutely into the six dimensions framework. Furthermore, these additional expertise themes represent specific findings of Study 1, and add to the discussion regarding how these aspects of the task environment impact the overall goal of identifying automation requirements.

Data saturation was achieved for Study 2 by assessing number of data sources (interviews, observations, existing publications) that supported each theme. This technique, also known as triangulation, was previously discussed at the beginning of this section (Section 3.3.2) regarding reliability. Table 3 below denotes instances from both studies, as well as references from literature, that support the notion that the respective expertise theme is relevant for CSIRT analysts.

Considering the high degree of agreement across Study 1 participants, Study 2 participants, and existing literature on CSIRTs, the results of this saturation analysis concluded that five (5) participants for Study 2 were adequate to ensure content validity.

Table.3. Saturation Analysis

Expertise Theme	**S2 Participants (5 total)**	**S1 Teams (3 total)**	**Literature**
Subject Matter	P1, P2, P3, P4, P5	T1, T2, T3	(Assante & Tobey, 2011; Bada et al., 2014; Cobb, 2016; L. Hoffman et al., 2012; Newhouse et al., 2017; Ruefle et al., 2014; Steinke et al., 2015; Tøndel et al., 2014; Werlinger et al., 2010)
Communication	P1, P2, P3, P4, P5	T1	(T. R. Chen et al., 2014; Cobb, 2016; Y. Lee & Lee, 2004; Ruefle et al., 2014; Steinke et al., 2015; Tøndel et al., 2014; Werlinger et al., 2010)
Info Flow Path	P1, P2, P3, P4	T2, T3	
Expert Identification	P1, P2, P3	T2	(Ruefle et al., 2014; Steinke et al., 2015; G. White & Granado, 2009)
Interface/Tool	P2, P3, P4	T1, T2	(Silva et al., 2015; Tøndel et al., 2014; Werlinger et al., 2010)
Situational Context	P1, P2, P3, P4		(Ahrend et al., 2016; Bishop et al., 2017; Ruefle et al., 2014; Steinke et al., 2015; Tøndel et al., 2014; Yufik, 2014)
(Added) Policy	P1, P3, P4	T2, T3	(Bishop et al., 2017)
(Added) Self-Awareness	P1, P2, P3, P4, P5		(Cobb, 2016; Freed, 2014; Y. Lee & Lee, 2004)(Ahrend et al., 2016; Bishop et al., 2017; Ruefle et al., 2014; Steinke et al., 2015; Tøndel et al., 2014; Yufik, 2014)

3.3.3. Workflow Diagramming: Understanding Incident Response (IR) Processes

A workflow diagram is a type of process chart that provides information about work (Graham, 2004). Creating these diagrams can be useful in gaining better understanding of process components beyond standard operating procedure, identify opportunities for improvement, and gain a bigger picture view of the process as it flows through multiple entities of an organization (Damelio, 2011; Graham, 2004). The workflow diagram was one analysis method used during the exploratory field study (Study 1). The method generally followed the process mapping methodology in (Graham, 2004).

The IR process in cyber security operations has been flowcharted from a very high level (Ruefle et al., 2014). However, additional details are needed to understand how the people and systems within the process interact, allowing for richer context around activities and users needs of information systems in incident response. More specifically, data collected during Study 1 indicated that incident response occurs in tiered organizations in which different sub-teams react with different levels of response, which is further discussed in Section 3.3.4. The data also indicated that different information systems are used between entities, increasing the number of interfaces, and thus complexity, within the process. Considering the context gaps highlighted here, cross-functional process maps were selected for analysis. Cross-functional process maps, also known as swimlane diagrams (Damelio, 2011), highlight boundaries and how workflow crosses those boundaries; they are useful in determining where handoffs occur between entities. Similar to a cross-functional process map is the operational sequence diagram, which can be used to create more detailed representations of man-machine systems for the purpose of better system design (F. A. Brooks, 1960; Kurke, 1961). As the goal of Study 1 was go gain better understanding of the overall process and organization, the cross-functional workflow diagram had a level of detail sufficient for that purpose. The operational sequence diagram is much more useful at later stages of development of a new system.

Using data generated from observations and interviews, cross-functional process maps were created to depict the process flow across organizational entities for each location. The entities are assigned a horizontal space (or lane), and the process is mapped from left to right. Entities can be divided into sub-entities to show handoffs that occur within the same group or department. The general procedure closely followed (Damelio, 2011), using consistent symbols and conventions to illustrate the process based on available data.

3.3.4. Management Cybernetic Analysis: Understanding IR Organizations

Another analysis method used in Study 1 was a management cybernetic analysis of the three teams studied. The Viable System Model is one such method that has been proposed as for understanding information sharing in security organizations (Gokhale & Banks, 2004). The goal of using VSM as an analysis method for this study was to understand what information sharing is present versus what is needed for a viable security organization. Moreover, by comparing three teams from different sectors, the analysis can provide insights regarding impact of environment on information sharing, system organization, and viability. Findings from this analysis are triangulated with other data sources in the dissertation, as well as external sources in literature.

Concerns from literature highlight corporate complacency as a common issue that affects security organizations (KPMG International, 2015; Rhee, Ryu, & Kim, 2012; Tøndel et al., 2014). Gokhale & Banks (2004) note that VSM acts as a framework to consider these issues, which is especially appropriate in incident response settings. More specifically, VSM presents algedonic signals, or pain/pleasure signals in the neurologically inspired cybernetic framework, as alarms or rewards that escalate up through an organization when certain events occur. Within CSIRTs, algedonic signals represent alerting mechanisms from CSIRTs to the corporate level, which provide a level of security awareness to the control functions, strategic decision makers, and policy makers of the corporation (Systems 3, 4, and 5, respectively).

Other information flows in the form of metrics have also been cited as particular pain points in computer security settings (Ruefle et al., 2014). VSM metrics act as indicators of performance of the CSIRT or larger security organization, and can provide different levels of management with a better understanding of the operational effectiveness of the System 1 components. Thus, VSM is a valuable tool for highlighting potential strengths and weaknesses in metrics reporting, as it focuses on feedback, or return signals, as an adaptability feature of a system.

I applied Beer’s VSM methodology to analyze data collected during Study 1. This analysis method helped to map information flows at different levels of the organization, scoped to include System 1 (operators) to System 3 (senior management), described below. Published articles for applying VSM (M. C. Jackson, 2000; Walker, 1998) guided my implementation of this method.

Using methodologies like VSM, the neurocybernetics framework can represent an organization broken down into five (5) systems at different levels. Each level serves a unique purpose towards the overall system goals, and they all must work together to adapt to environmental changes (M. C. Jackson, 2003). This section provides an overview of these system descriptions, with some orienting descriptions of the context within security operations. Results of the cybernetic analysis are presented in Chapter 4.

System 1 is comprised of a group of components, or *subsidiaries*, that work to deliver the main outputs of the overall system. These subsidiaries are not separate companies, but rather operational units within the same organization. Each of these subsidiaries has its own local management structure (M. C. Jackson, 2000). Within the context of Security Operations, System 1 subsidiaries are responsible for maintaining stable and secure operations for the overall organization. Subsidiaries may be different IT groups embedded in separate business units, or different support groups for specific appliances in the network defense systems. Each subsidiary provides some key function within the incident response cycle. Thus, not all incident response functions are contained with an organization titled "security operations". Security Operations functions can be spread across a company. For instance, a subsidiary may be the Security Operations Center (SOC), which is responsible for detecting a threat, then handing off threat information to affected business unit IT groups for them to handle according to their respective procedures. Alternatively, the SOC subsidiary may perform all containment and mitigation tasks after a threat is detected, then task other IT or business groups with investigative or remediation tasks.

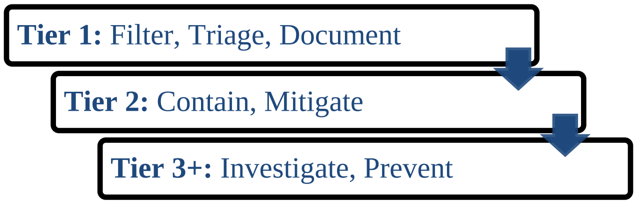
Tier 1: Filter, Triage, Document
Tier 2: Contain, Mitigate
Tier 3+: Investigate, Prevent

Figure.10. Tiered Organization in Information Security

Within a given System 1 SOC subsidiary, organization schemes have recursive properties. One key feature of SOCs is that many are organized in a tiered fashion, with different levels of responders grouped by level of expertise and system access (Figure 10). When a potential threat is detected, the lowest level of responder (Tier 1) will perform some action to filter, verify, and document the threat. The actions are procedure-driven at Tier 1, and, after complete, the ticketed threat will be passed to the next tier (Tier 2) for additional activities to contain and mitigate the threat. If additional activities are needed from security analysts, Tier 2 would then escalate the threat to a specialist in Tier 3 (or higher). If there are no further security actions, the Tier 2 analyst would task other subsidiaries, such as an "exchange team", or a "hardware team", with activities required within particular appliances or physical locations.

Referencing VSM, these incident response tiers are not always represented as separate subsidiaries. They may report to the same local management structure, use the same information systems, and use the same environmental data to drive threat awareness and response decisions. As stated, the recursive properties of SOCs would allow for deeper analysis to further investigate SOC organization and operation. As literature indicates that CSIRTs might differ in organizational structure (Killcreece et al., 2003; Ruefle et al., 2014), the cybernetic representations were "normalized" to be able to analyze the teams at the same level of recursion. Recursion is essentially the principle that there are multiple layers within a larger system that are similar in structure (Beer, 1984); this analogous to reference of "grain size", which distinguishes which layer is being analyzed or addressed. Thus, the security operations organization is the system of focus, with System 1 subsidiaries being bounded by the smallest managerial unit across all three teams and System 3 representing senior management over the security operations organization.

System 2 facilitates collaboration and alignment amongst the affiliated groups, and is commonly identified as the information systems between the subsidiaries (B. Williams & Hummelbrunner, 2010). System 2 allows System 1 subsidiaries to communicate and coordinate activities. Applied to Security Operations, System 2 includes email, instant messaging, phone systems, ticketing systems, and other security software by which different groups can pass information. However, in a dynamic setting, it is important to consider how System 2 can facilitate or negatively impact coordination between subsidiaries. In some instances, System 2 may be limited to email, which can create a communication funnel through a widely used channel. In other cases, System 2 may include redundant communication systems that are not integrated, which can create multiple instances of information that need to be maintained or acknowledged.

System 3 and System 3* act as the audit and control functions, respectively (Flood & Jackson, 1991; B. Williams & Hummelbrunner, 2010). System 3 is generally referred to as management, and is responsible for accountability and resource allocation of System 1 components. It acts as a control function, ingesting information regarding policy (from System 5), strategy (from System 4), and performance (System 3*), and translating these into actions to direct System 1, effectively providing organizational stability (B. Williams & Hummelbrunner, 2010). As scoped here, System 3 should have visibility over Security Operations as a whole, including performance and effectiveness of each subsidiary. System 3* is generally regarded in literature as the audit mechanisms in the organization that allow metrics or performance measures to flow to System 3. This is denoted as System 3* (and not as System 4) because it is the channel for the cohesion and monitoring function (Espejo & Reyes, 2011; Hoverstadt, 2010) fulfilled by System 3. Figure 6.6 in (M. C. Jackson, 2003) illustrates parallel channels of feedback for System 3*, as well as between other Systems in the VSM.

System 4 is the strategic level of management that has connection to external information feeds and makes longer-term decisions regarding the direction of the organization. System 4 has direct communication with Systems 3 and 5, and provides intelligence on how the overall system must adapt in order to remain viable (Vidgen, 1998; B. Williams & Hummelbrunner, 2010) This system must act as both a filter and a switch (M. C. Jackson, 2000), determining which information from Systems 1 through 3 is appropriate and necessary for System 5, as well as which intelligence from the environment and which information from System 5 need to be pushed down through to System 3 and lower. Literature in computer security generally regards the Chief Information Security Officer (CISO) as the intelligence function capable of maintaining System 4 (Hale, 2017). However, not all organizations allow this position to fulfill its intended purpose, which will be discussed in later sections.

System 5 is the policy-making body in the organization (M. C. Jackson, 2000). This may be a group of individuals, and may or may not all be elected by shareholders, appointed by owners or members in the case of a non-stock corporation or LLC, or appointed by other legally authorized entities in the case of a non-profit such as a university or professional society. For instance, System 5 may be an executive board comprised of chief officers for various functional areas, and it may not always include a security-minded leader with security expertise. In this case, this board may itself appoint a committee to specifically make decisions regarding security risks and priorities. System 5 must prioritize what is in the best interest of the overall organization when making decisions regarding policy. System 5 is notably weak with respect to recognition and expertise in computer security in both for-profit and non-profit organizations (CompTIA, 2018; DCRO Cyber Risk Governance Council, 2018; Kral, 2018) and within the federal government (Dunne, 2018; Newman, 2018).

3.3.5. Cognitive Task Analysis: Eliciting Knowledge from CSIR Experts

Within the context of this dissertation, CTA was used to understand expertise requirements for potential automation development. The second study used Applied Cognitive Task Analysis (ACTA) (Militello et al., 1997) as the main data collection method to explore some of the critical functions identified in the exploratory field study. ACTA is one appropriate method within CTA as it was developed for practitioners to conduct CTA in applied settings by creating usable procedures and aids (Militello et al., 1997). Due to the structure and availability of ACTA, this dissertation aimed to promote ACTA as part of the methodology and as a replicable procedure for cyber operations, which may help others apply human-centered approaches in complex cyber settings.

The ACTA methodology consists of three (3) main stages, which were all included as part of the methodology of this dissertation. First, a high-level task diagramming activity was conducted consisting of a task and 3-6 subtasks. The task of interest was selected from data collected in the exploratory field study (Study 1), and focused on critical information sharing functions within the CSIRT. The next step of ACTA is a knowledge audit, which helped determine which expertise is needed to perform the task and subtasks. Questions in the knowledge audit included dimensions of expertise as they relate to information sharing. For instance, the audit aims to elicit knowledge regarding cues, strategies, and difficulties in the tasks being investigated. Thus, the results of the knowledge audit provided details regarding the “situational context” expertise for the key information sharing tasks. Lastly, a simulation interview was completed with the expert to elicit decisions and judgments made during the task. Supplemental questions were added during the CTA with each participant to help address specific aspects of information sharing with respect to the six dimensions of expertise (Appendix B).

3.3.6. Needs Analysis & Functional Architecting: Concept Development

HAI is quickly making strides in the computer security domain as new platforms enter the market to address operational and resource challenges. In order to situate the first two studies in relation to these developments, existing automation technologies and capabilities were explored to identify gaps between current technological capability and needed expertise (or expertise augmentation). This included research regarding existing popular products and platforms to collect information on technology, supplemented with automation-related data from Study 1 and Study 2.

The process described in (Levis & Wagenhals, 2000) was used as a guide to initialize the functional architecture of specific critical functions determined with previous methods. As mentioned in Section 3.1.2.6, the architecture development process starts with a concept. The process of developing an operational concept is more on the less-scientific, interpretive side of architecting (Levis & Wagenhals, 2000). It is a goal statement for the system to be developed. By analyzing observational and interview data from stakeholders collected in Study 1 and Study 2, a goal statement was developed to reflect the needs of users. The needs defined during process corresponded with concepts for new functions.

After this concept foundation was determined, I performed a functional decomposition of the system using products of the CTA (Levis & Wagenhals, 2000). This step resulted in a breakdown of functions into sub-functions, inputs, outputs, mechanisms, and constraints; in order to avoid unnecessarily increasing complexity at this stage of conceptual development (Levis & Wagenhals, 2000), an exhaustive approach was intentionally not employed for the decomposition. Finally, the functional architecture was constructed. The functional architecture included four major components to reflect different questions an operator might ask: the activity model, the dynamics model, the data model, and the rule model. A high-level perspective was adopted to avoid restricting analysis to only an information system level as described in the Levis & Wagenhals article.

One function identified in the Needs Analysis (Study 3) was selected based on estimated feasibility; the purpose was to provide a clear example of functional architecting for new concept development in CSIR. Using a systems engineering standard for IDEF (*ISO/IEC/IEEE Std 31320-1:2012*, 2012) and results of Studies 1 and 2, *the activity model* was generated for one function identified in the Needs Analysis. The activity model is a visual representation of the functional decomposition, and can help determine separate algorithmic assets that need to be developed. The *dynamics model* for this function describes the behavior of the function; it included definitions of states and the transitions between states. The dynamics model helps developers and operators understand the modes of operation and what triggers transitions between those modes. Definition of inputs, transformations, and outputs satisfy the general requirements of the *data model* as described by Levis & Wagenhals (2000). The data model was created based on what data flows (and attributes thereof) would be required by the function to execute properly; sources of data are also discussed that might need to be established. The *rule model* has notably strong parallels with the outputs of the CTA, in which experts designate inputs, conditions, outputs, and cues for decisions. Thus, the results of the CTA informed the construction of the rule model, which includes decision trees and tables, and highlights relationships between data sources and conditions under which functions need to be executed.

CHAPTER 4. OBSERVATIONAL FIELD STUDY OF CSIRTS

Chapter 4 focuses on the first of three studies presented in the research conceptual framework. The data collection method was an observational field study based on ethnographic methodology. Analysis methods included qualitative thematic analysis, workflow diagramming, and cybernetic analysis. Data collection and analysis details are described in Section 4.1. Study 1 was used to answer four (4) of the research questions presented in Chapter 3; together, these questions help establish the currents state of CSIRTs across three different environments. Findings and answers to these research questions are presented in Section 4.2.

4.1. Data Collection and Analysis

4.1.1. Sampling

As described in Chapter 3, one goal of Study 1 was to capture variation across different CSIRT settings. To help observe this variation, a stratified cluster sampling strategy (Pelto, 2016) was used to study a subset of the CSIRT population, where the unit of study was one team. Though mainly opportunistic, the sampling strategy included teams from different organizations or institutions with some variation in size or specialization. Three (3) teams were investigated as part of this study, each with a different operational setting. These settings included a company in industry, a university, and a state government.

Table 4 provides a general description of the three teams observed, including purpose of the security team in each setting, the specific operations in focus, and the ‘viable parts’, based on the Viable Systems Model. Analysis methods included aggregation across these teams, as well as comparing between them, which are discussed in depth in Section 4.2.

Table.4. Organizations Included in Study 1

	For-Profit Company	**Public University**	**State Government**
Objective	Stable and secure operations, customer support, production	Stable operations (uptime)	Stable operations (uptime) and secure storage of information
Mission/ Scope	Shared responsibility with IT Operations; Entire company network	Within central IT Operations; Limited scope to certain incidents / traffic	Within central IT Operations for one branch of gov't; Limited scope to certain gov't departments / traffic
Viable Parts	Groups, IT Teams (shared management structure)	Colleges, Schools, IT Teams	Departments, Groups, Teams

4.1.2. IRB Compliance and Participant Recruitment

The application package for Study 1 was submitted to the Institutional Review Board (IRB) and included a narrative, protocol, and consent form (Appendix C). The protocol was approved for research (IRB protocol number 1801020155, approval date 03/27/2018). For each of the three locations, personal contacts were sent IRB-approved recruitment language via email to potential organizations. Due to on-site data collection, two additional levels of IRB approval were needed (after initial conditional approval) at the company level and the individual participant level. First, IRB guidelines state that the researcher is required to obtain a signed statement from a manager at each location stating that research personnel would be allowed onto the respective premises to talk to employees directly. This signed document, which acted as the company consent for research on premises, was obtained and submitted to IRB to be attached to the protocol. This step of approval preceded recruitment of individual participants; although the unit of study was at the team level, IRB requirements state that individual consent is needed with an approved consent form. To recruit individuals, manager acted as a conduit for recruitment documents. *IRB communicated that some level of committed participation was needed before full approval could be granted,* as the manager acted as a conduit for recruiting and could not commit employees to the study by him/herself. IRB then approved the protocol for data collection.

Upon arrival at each location, each CSIRT manager was invited to discuss the study and interviewed to collect background information regarding their organization, its structure, processes, and characteristics. After this orientation exercise, direct interactions with participants and other employees occurred to reorient teams to the study. If the participants had not signed the consent form prior, they were given the opportunity to sign at this point. One detail worth noting is that, at 2 of 3 locations, the majority of participants consented to the study in person rather than beforehand via email. Signed consent forms were scanned and uploaded within 48 hours to secure cloud storage.

4.1.3. Data Collection Procedure

A general description of the procedure is included in the application narrative, located in Appendix C. Upon arrival, a tour of the security operations center (SOC) was provided during by the manager; this provided a general idea of layout and environment as well as introductions with members of the incident response team. After the tour, an unstructured interview was conducted with the manager in an office or conference room, which lasted 2-4 hours. This interview was not audio or video recorded per request of the managers.

The one-on-one interview provided insights into the organization from the managers' perspectives, which provided better understanding of the organizational structure and some tensions between different parts of the company. The manager also shared general security threats that the organization faced regularly. The "stack," or general security architecture was also described: this sometimes included a general inventory of appliances used by the firm to protect the network and remediate threats. The interview also included descriptions of SOC team responsibilities regarding incident response, with some coverage of procedures and protocols. Finally, the manager also described reporting practices of incidents throughout the organization.

Direct interactions with analysts commenced after the initial meeting with the manager. This included conversations in the operations center, in conference rooms, in other shared spaces, and also sitting with analysts at their stations to observe live incident handling. In some locations, being "on the floor" was not acceptable, as not all teams were comfortable with direct observations in this setting. Accordingly, adjustments were made to the IRB protocols to accommodate this concern by using conference rooms and other spaces to talk to analysts about more general aspects of their work. After several hours of live observations (which the organizations referred to as "shoulder surfing"), it became clear that, without more context regarding the systems being used, and why, the exercise was not fruitful in producing a higher-level understanding of the team's processes. However, the interactions with analysts did provide a better understanding of tool/system complexity, general ergonomic and usability considerations for incident responders, and some of the "native language" used in this setting.

Analyst interviews included some of their personal backgrounds and experiences, daily duties in their respective roles, explanations of certain processes, issues within those processes, and even why some of those processes were in their current state. The informal, unstructured interviews also included the analysts describing how automation was used in their organization, and issues or concerns about automation. Analysts also described shifts, shift changeovers, and escalations. This interview process was repeated at different levels of analysts or different sub-teams if they were available. If participants were especially helpful or reflective, they were revisited for further questions. A description of the participant profiles per location can be found in Table 5. Due to the relatively low number of participants in each level of the organization, the data were aggregated (Table 5) across all three locations to prevent identifiability issues.

Table.5. Participant Profiles across Locations

Participant Profiles	**Sum of Participants Across All Locations**
Management	4
Tier 1	9
Tier 2 or Higher	8

For direct observations, each analyst was asked to use think aloud protocols (Hartson & Pyla, 2012) when handling an incident. They would perform their tasks while describing what they were doing while handwritten notes were recorded. Interview questions also probed for more information regarding why they would do certain activities. Notes included aspects of analyst work, including workstation setup, pivoting activity between systems, and even some keystroke-level trends.

For organizations that did not allow direct observations, simulation interviews similar to that discussed in the ACTA methodology (Militello et al., 1997) were used instead. In this simulation interview, managers or analysts would essentially engage in a cognitive walkthrough (Polson, Lewis, Rieman, & Wharton, 1992) for past or hypothetical incidents, noting step by step what would happen at different levels of organization. Many times, these incidents were considered routine in the eyes of the participants. Participants were then asked about what they would consider non-routine incidents; responses indicated that many of these 'non-routine' incidents were simply escalated, and that higher tier analysts were needed to describe the response. The distinguishing factors between routine vs. non-routine for the lower tier analysts were: their familiarity with the incident type and if policy allowed them to address it. Data indicated that, the higher the analyst, the less routine their work became. That is, in all organizations observed, higher tier analysts would select systems and actions based on unique information to a given incident, making it difficult to describe any standard procedure at a useful level of detail.

4.1.4. Data Processing

After each day of data collection, all generated data were reviewed and copied into more organized notes into a separate notebook. These notes also included more detail and reflection for later use. Data were transcribed from handwritten notes into MSWord, and then imported into NVivo. Many of the statements generated during data collection were in the form of bullet points about questions or targeted areas from the observations and interviews. Diagrams or illustrations regarding information flows were translated into statements and then transcribed for analysis.

4.1.5. Qualitative Data Analysis: Coding and Theme Development

Using an analysis approach from qualitative research, themes were generated during two rounds using a bottom-up coding scheme (Auerbach & Silverstein, 2003; Saldana, 2009). Round 1 started with a systematic review of the previous level codes to group into categories. The goal was to group the codes based on the higher-level theme to which they relate. For example, the codes in Table 6 were identified in the prior step and grouped into a theme on the right.

Table.6. Round 1 Theme Development Example

Grouped Codes	**Round 1 Theme**
Shift leads have private meetings after a handoff to discuss higher level issues	Handoffs establish shared awareness
Shift handoffs create shared situation awareness	
Shift handoff includes reviewing open items in system	

Round 2 followed a similar process. Once more, codes generated in Round 1 were systematically reviewed and grouped to extract the next level of theme. Analytical questions were used questions to guide the categorization: “What are the common threads between some of the themes? Do they relate? What seem to be mentioned together?” In the same spreadsheet, another column was added to accommodate the 2nd order theme in the rightmost column. Count sums of the original code references were used to get a picture of quantitative breakdown of codes to themes. This quantitative approach could be biased based on the amount of time that people chose to talk about a particular topic in semi and unstructured interview settings. The second processing step reduced the number of data points from 99 to 11 themes. These themes are presented in Section 4.2.1 and located in Appendix D.

4.2. Results

4.2.1. Identifying Themes in CSIRTs

The results of the observational data analysis included 11 themes about SOCs and computer security incident response. Themes developed with qualitative analysis provide potentially useful insights driven by semi-structured data collection instead of specific assumptions and questions about the setting and participants. These insights allow researchers and managers alike to consider additional variables or factors outside a given set of concerns, as well as a broader view of potential issues in security operations settings beyond a single team or issue.

The themes represent findings from across all three teams and are presented by rank order (highest frequency to lowest) in Table 7. These themes were used to answer research questions, and tied to other data analysis findings.

Table.7. Themes Developed from Observational Data

Rank	Freq.	Theme
1	38	Communication, feedback and accountability are necessary for IR, awareness, and learning; If lacking within or between levels of org, issues arise
2	36	Organizational alignment on security priorities and awareness of IR issues is important for "full-cycle" IR process
3	35	Continuity of awareness and documentation around incidents is important
4	23	IR requires a wide range of skills and flexibility; Workforce may not be able to maintain if not designed to do so
5	22	IR requires a wide range of activities, including filtering and decision making; These can be split based on expertise or authority of analysts
6	20	Automation is seen as a potential solution for low-level tasks and coordination, but considered out of reach for teams who don't have the support resources
7	14	Knowledge sharing (in a repository, in person, or through other channels) may be important for learning and shared awareness
8	12	Formal and informal roles emerge to meet an organizational need for management, communication, and decision making
9	11	Identity and culture of the team affect communications and responsibilities
10	10	Handoffs are varied in terms of procedure, formality, and documentation; In whatever form, they are important for continuity in several contexts
11	8	Incident handling methods may be indicators of organizational maturity; Maturity as a focus may drive incident handling methods

4.2.2. Mapping CSIRT Processes

As described in Chapter 3, workflow diagrams are one method of process mapping that allows illustration of a given process. These diagrams can be useful in gaining better understanding of process components beyond standard operating procedure, as well as a bigger picture view of the process as it flows through multiple entities of an organization (Damelio, 2011; Graham, 2004). In particular, “software systems” were denoted as separate entities in the workflow diagram to help frame the role of technology as more than a tool, but also as a critical component that could someday be automated and transform into another “teammate” entity. Finally, the workflow diagrams were able to help identify critical components of information sharing in CSIR, and thus focusing Studies 2 and 3 to those critical areas.

Components of the visualization draw upon participant knowledge of procedure, and include additional steps that participants included in their recollection of how the process works. That is, some components are not part of the procedure, but rather the representations of how the analysts think about their daily tasks in incident response. Observations supplemented participant inputs, and acted to verify steps. Task diagrams created in Study 2 validate some cognitive aspects of this process (Chapter 5). In general, the process was an amalgamation of all participants in each location. Participants of Study 1 were not asked to review the diagrams, as they were fully constructed post-study. However, experts in Study 2 were able to verify major decision points across the general process flow for incident response (Appendix F).

An example CSIRT anomaly response and escalation process (Figure 11) begins with alerting of an anomaly by an information system (appliance or information management platform) and stops at closing or resolving an incident ticket. One item to note is that the cross-functional workflow diagram includes information systems as separate entities to show where systems play a role, especially considering the system separations noted in Section 3.3.4. The below diagram is an example of the full process flow diagrams located in Appendix E.

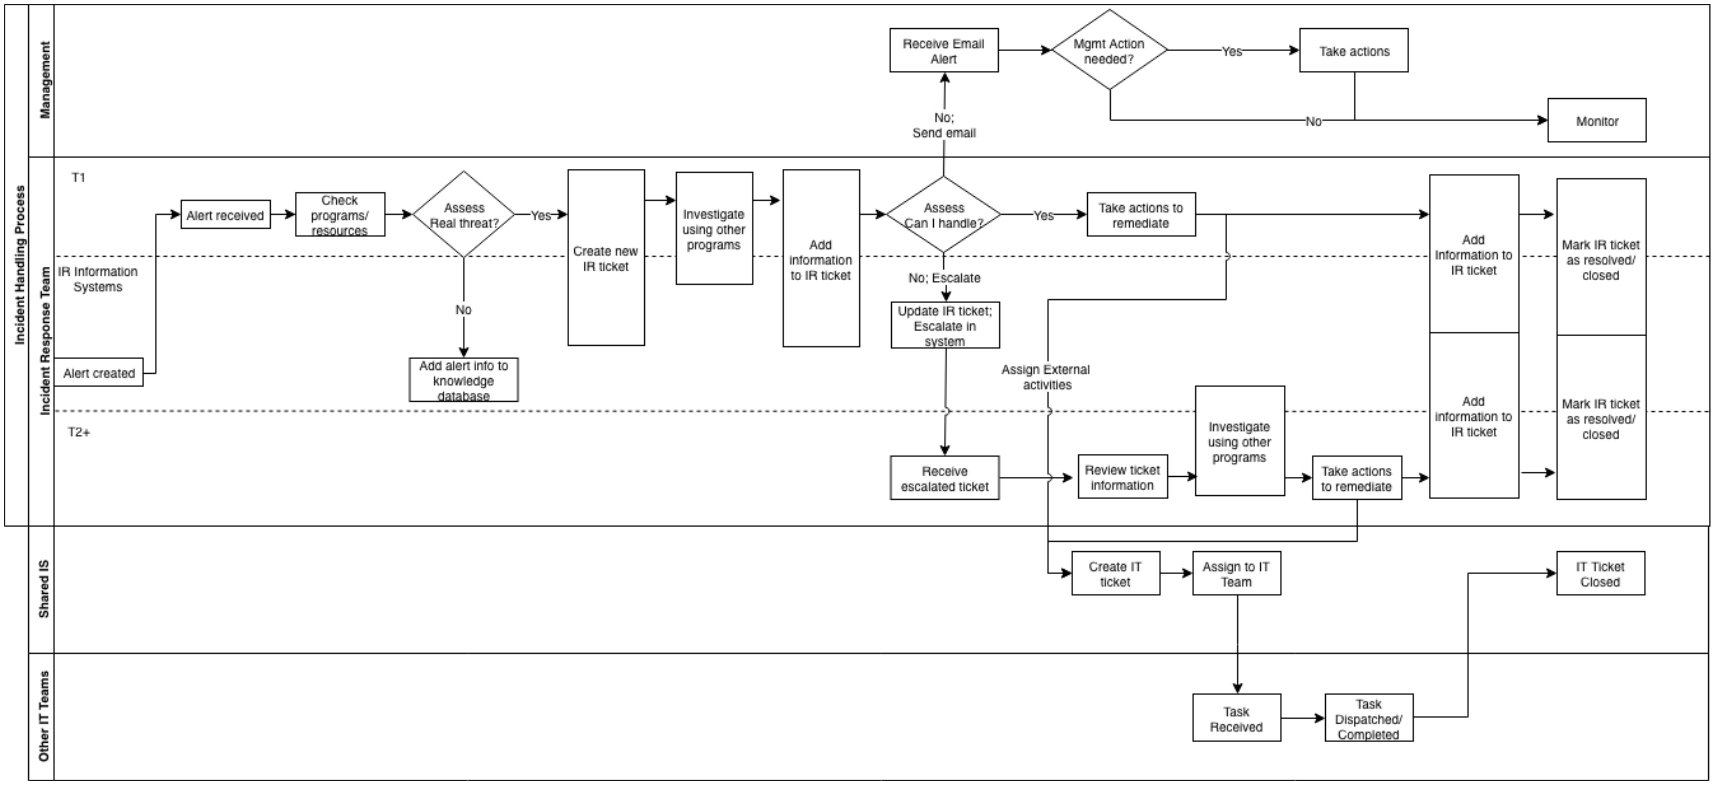
Incident Handling Process
Management
Incident Response Team
Shared IS
Other IT Teams
T1
IR Information Systems
T2+
Receive Email Alert
Mgmt Action needed?
Yes
Take actions
No
Monitor
No; Send email
Alert created
Alert received
Check programs/ resources
Assess Real threat?
Yes
No
Add alert info to knowledge database
Create new IR ticket
Investigate using other programs
Add information to IR ticket
Assess Can I handle?
Yes
Take actions to remediate
No; Escalate
Update IR ticket; Escalate in system
Assign External activities
Add Information to IR ticket
Mark IR ticket as resolved/ closed
Receive escalated ticket
Review ticket information
Investigate using other programs
Take actions to remediate
Add information to IR ticket
Mark IR ticket as resolved/ closed
Create IT ticket
Assign to IT Team
IT Ticket Closed
Task Received
Task Dispatched/ Completed

Figure.11. Example of Workflow Diagram from One CSIRT

The workflow diagram highlights these particular entities, in order to address several issues. First, all teams had separate security operations IR teams and IT operations groups. Most teams reported through separate lower management, but interacted regularly for incident response and so are considered separate entities accordingly. Second, as mentioned, sometimes the entities interact through shared systems, but also have separate systems to collect and store their relative information. For instance, within the CSIRT entity, analysts at different tiers may share ticketing systems, or not. Within the same entity, analysts may not have shared access to other systems used during investigation. Lastly, sub-teams (or subsystems) address different tiers of response and pass between each other. Accordingly, the incident response group is split into different sub-entities (tiers) that pass information between each other. These tiers were previously described in Section 3.3.4.

Information systems were separated into only two entities: IR systems and IT systems. Sometimes members of different entities have access to both systems. For instance, the IT ticketing system is typically owned by IT, but IR analysts have access to create, assign, and track tickets within that system. However, there is often another ticketing system owned by IR that is solely managed and maintained by the IR entity. While not all of information systems (with their owners and interactions) are represented in the diagrams shown, it is important to note that the complexity in information systems is a major driver in the system as a whole, and that there are a variety of interfaces between shared and unshared systems to consider. The information systems across locations also varied greatly, some utilizing less than 10 systems at the T1 level, some utilizing over 40 systems at the T1 level.

4.2.3. Viability of CSIRT Organizations

This section describes the three CSIRTs observed (in three respective subsections) using the VSM framework for terminology and visualization, known within the methodology as ‘system identification’. A review of the VSM methodology and how it was applied in Study 1 can be found in Chapter 3. Recalling the remark about recursion in Chapter 3, each organization was analyzed at the same “grain size” or layer, such that Security Operations is at the System 1 level, and the board of directors (or similar body) occupied System 5. The initial interview with managers helped to determine who or what occupied the systems in between Systems 1 and 5. VSM principles were used to determine 'who does what', which did not always align with organizational charts. That is, an organizational chart of a company does not necessarily directly translate to VSM Systems 1 through 5.

The VSM descriptions of CSIRTs can help draw out similarities and differences across the various organizations, aiding in later discussions. The most notable difference between the organizations is in the System 2 and 3* connections between System 1 subsidiaries and System 3, which help with information flows among them. Appendix G presents the results of the diagnosis (the second part of the VSM methodology) and delves deeper into questions about viability in these organizations.

4.2.3.1. Location 1 – A Public Industrial Company

The first organization was a company with the overall purpose of providing quality products and uninterrupted services to external customers. This company experienced security threats from foreign nation states, and aimed to protect its own network as well as the security of their products deployed to customers. Some specific types of incidents common in this organization included phishing and malware. Security Operations was the system of focus within this organization. The goal of Security Operations was to maintain stable and secure operations, internal customer support, and external collaboration with similar organization external to the company.

In this organization, System 1 can broadly be described as the segments that support stability and security within the company's network. Generally speaking, these segments would be IT and Security that act as independent structures (Ruefle et al., 2014). There were more potential groups under this scheme within System 1 and subsystems of A and B that are not represented here. As shown in Figure 12, (internal) Subsidiaries A and B are illustrated as reporting up through parallel, but aligned, executive channels to Level 3. Based on cybernetic representations (Herring & Kaplan, 2000; M. C. Jackson, 2000), each subsidiary has some control mechanism that oversees its daily operations and functions. This is represented by the respective management positions over the departments.

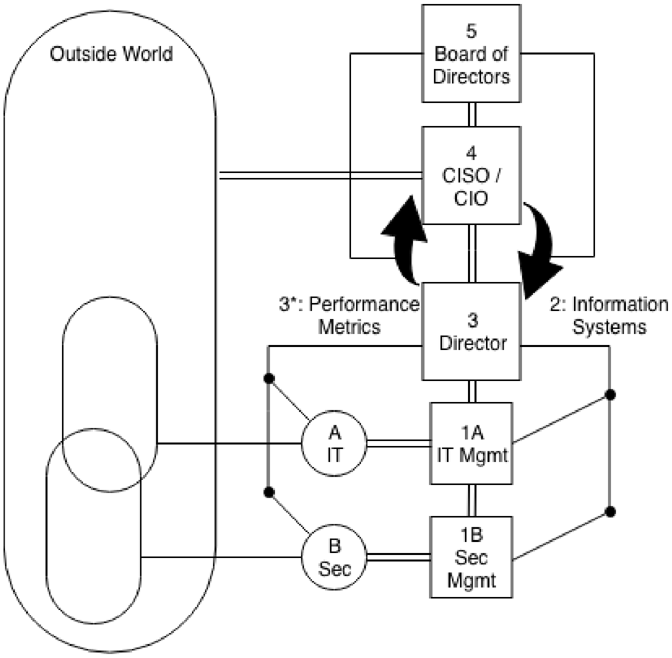

Figure.12. Location 1 Cybernetic Diagram

System 2 represents the information systems between the subsidiaries that facilitate coordination. These are shown as solid lines to depict that these systems interact in a consistent manner, including between the subsidiaries. That is, most information systems were shared, but there were some separate systems. Lastly, the subsidiaries in System 1 have a shared awareness of network operations, allowing for more fluid information alignment between them.

System 3 represents the “line management” (Hoverstadt, 2010), which oversees all System 1 subsidiaries and assesses general performance and resources. In this organization, the System 3 is represented as the Director of Security Operations, but also includes security representatives embedded in other divisions of the business. Together, these comprise System 3, though the division representatives are responsible for communicating security policies, goals, and objectives directly to other parts of the business on behalf of System 4. System 3* allows metrics or performance measures to flow to System 3 (senior management). Within this organization, performance metrics were tracked in terms of time to mitigate incoming incidents, and were reviewed by System 3 management.

System 4 within this company is represented by two entities: the Chief Information Security Officer (CISO) and Chief Information Officer (CIO). These individuals strategically lead the IT and Security organizations in the larger company as peers. Lastly, System 5 is represented here as the Board of Directors: a body of individuals from different aspects of the business who assess the company's goals and direction, and create policy to support those initiatives.

4.2.3.2. Location 2 – A Public University

The second organization (Figure 13) is a university whose main goal is to educate students and conduct research for internal and external funding agencies. Though the most common threats noted in this organization were related to phishing and invalid account logins, participants noted that the value of data stored at a research institution is high, especially considering sponsors and goals of potentially sensitive projects. The system of focus within this organization was Security Operations within IT. The goal of Security Operations was to support stable operations with uninterrupted network uptime and monitor network activities.

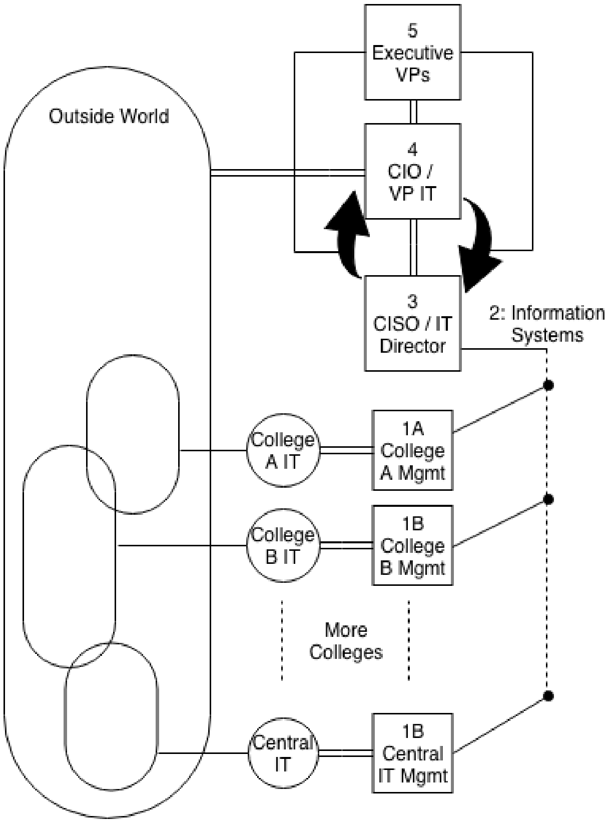

Figure.13. Location 2 Cybernetic Diagram

Within the university security organization, System 1 is represented by the separate college subsidiaries that have their own IT groups, and who are managed by separate control mechanisms. There were multiple colleges with IT groups, not all are represented here. The specific subsidiary studied in this research was Central IT Security Operations. These subsidiaries reported up through the CISO and IT Director levels, which were organized hierarchically under the same executive. System 2 is represented by dotted lines, as the information systems were present, but most were not shared. Email and general IT ticketing were shared systems, but there was limited network visibility by Central IT into the operations of other subsidiaries.

The institution did have an internal audit function that represents System 3*. The internal audit website states that security of data is included in the scope of the audit function, though often in the context of financial and personal information. The internal audit office oversees the audit function over the entire university system, and includes a wide range of goals, services, and risk categories. However, evidence and impact of this function were not observed within Security Operations. Another mechanism that can provide feedback to System 3 is performance measurement. Within this organization, participants did not indicate that metrics were actively tracked between different IT groups.

The CISO and IT director represent System 3. The CISO was directly responsible for day-to- day operations within the team, and often covered on incidents to balance system oscillations between the System 1 subsidiaries and from above. From the limited observation scope, System 4 includes the CIO and/or Vice President of Information Technology, who were ultimately responsible for network stability. Lastly, the Executive Vice Presidents collectively represent System 5. It was unclear from the observational study how much security-related policy was actively and critically determined at this level.

4.2.3.3. Location 3 – A State Government

The third organization was a state government agency (Figure 14), whose purpose was to maintain stable operations across hundreds of constituent departments, bureaus, and sectors. Some common security threats noted in this organization included phishing and malware. Similar to the university setting, participants indicated the potential value of information and access / control within a state system to hackers, including the Department of Revenue, Department of Transportation, and the Department of Corrections.

The system of focus within this organization was the Security Operations group within IT, whose goal was to support stable operations with uninterrupted network uptime and monitor network activities, as well as security support to other subsidiaries in System 1. Within this organization, the separate department, bureau, or local government subsidiaries represent System 1. These entities had their own IT groups, and were managed by separate control mechanisms. There are dozens of other System 1 subsidiaries with IT groups; not all are represented here. The observed impact of these other subsidiaries was increase scope of SOC responsibility, but without increased authority to act upon threats to those subsidiaries. The specific subsidiary studied in this research was executive branch IT Security Operations. These subsidiaries report up through the CISO and IT Director levels, which are organized hierarchically under the same executive (the CIO).
System 2 is represented by dotted lines, as the information systems are present, but most are not shared. Email and general IT ticketing were shared systems, but there was limited network visibility by executive branch IT into the operations of other subsidiaries. System 3 is represented by the CISO and IT director. The CISO was directly responsible some level of daily operations within the team, and often covered on incidents to balance system oscillations between the System 1 subsidiaries and from above. System 3* took the form of cost-related metrics. These were rigorously tracked to assess cost influencers in the larger organization by other subsidiaries. Other security-related metrics were not observed at the SOC level.

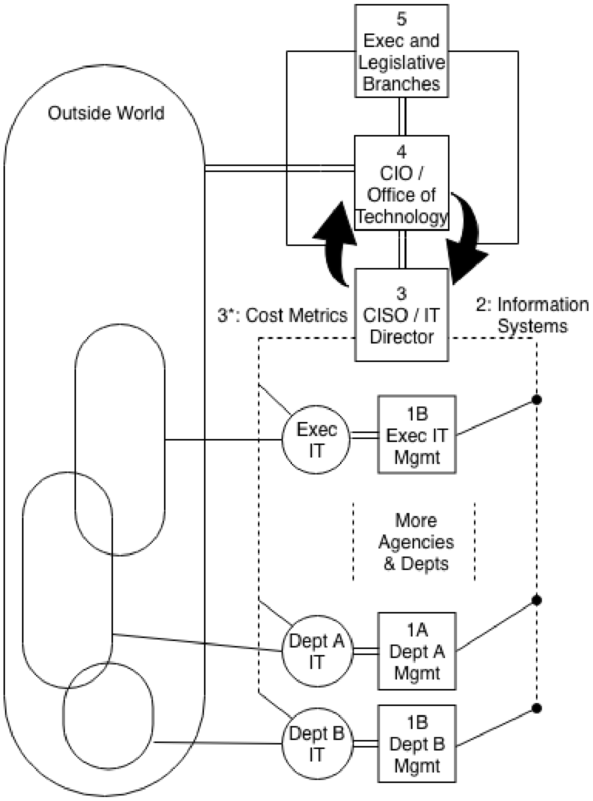

Figure.14. Location 3 Cybernetic Diagram

From the limited observation scope, System 4 involves the Chief Information Officer and the Office of Information Technology, who were ultimately responsible for network stability. Lastly, System 5 was difficult to ascertain for this team. System 5 seemed to be represented by the heads or policy-making constituents of the Executive and Legislative branches. It was unclear from the observational study how much security-related policy was actively and critically determined at this level.

4.2.4. Findings

This section is dedicated to answering the four (4) research questions pertaining to Study 1 (originally presented in Chapter 3). Answers to the research questions are presented through addressing sub-questions in progression; the sub-questions were explored during data collection and analysis.

4.2.4.1. Critical CSIRT Functions: Research Question 1.1 and 3.2

As described in Chapter 3, RQs 1.1 and 3.2 were intended to identify information flows and critical incident response functions with respect to information sharing processes performed by different members of the CSIRT. In so doing, the scope for Study 2 could be sufficiently narrowed, and generate insights could be generated across different CSIR settings with respect to influencers on these functions.

4.2.4.1.1. Information Sharing Flows and Processes in CSIRTs

Workflow diagrams and the VSM analysis were two methods that helped identify information sharing processes in SOCs. These tools were useful in understanding the complexity and structure of the SOC, and how incident response spans multiple subsidiaries and layers of an organization, as well as potential information sharing interfaces with external stakeholders.

The information sharing processes that were observed were: strategy or team meetings by tier level, team meetings at the Operations Center level, and a variety of handoffs. Based on interviews with experts in Study 2 (but not observed in Study 1), there is an additional process called an After Action Review (AAR). AARs have been identified as potentially valuable to CSIRTs, especially in creating shared awareness and learning opportunities p. Observed modes of information sharing in the participating teams included phone, instant message, email, shared information/ticketing systems, and face-to-face.

The participants viewed knowledge sharing and information sharing as separate activities, and they seemed to focus on information sharing as more important than knowledge sharing (depending on the organization’s aims and goals). Because of these distinctions made by the participants, information sharing in the incident handling process was deliberately separated from knowledge sharing for themes developed from the study. The modes of knowledge sharing, in addition to those listed above, include archived documents and analyst-managed wikis. General information sharing activities, such as communication, were grouped into the above processes, as communication is basic to all information sharing.

4.2.4.1.2. Critical Information Sharing Process: Escalation

Within the context of this dissertation, criticality of the process is defined through two parameters: occurrence rate (of the process) and sensitivity to errors. If the information sharing process occurs often and/or minor errors in the process can cause disruption, criticality is considered to be high. As the incident response process can involve many analysts at different levels of skill or access, information is often passed across tiers as often as incident alerts are created. Depending on entity size and the organization's ability to detection intrusions, participants in Study 1 indicated that hundreds of alerts could pass through a single response team in a given day. Based on these factors, the most critical information sharing process identified in incident response was a *handoff*, or the passing of an incident ticket from one person/team to another (Guinery, 2011). Literature supports the performance aspect of this criticality, stating that many errors occur during this process (*Growing the Security Analyst: Hiring, Training, and Retention*, 2014; Steinke et al., 2015) between individuals and teams. Interviews with experts also investigated the potential for error propagation during handoffs. During a handoff, a more experienced receiving analyst may verify the information received, thus catching potential errors of lower-tier analysts. Despite the necessity, some participants viewed this verification step as a waste valuable time for the verifier, which can impact the time available for their respective investigation tasks.

Handoffs were observed in several forms: (1) escalations, (2) passing between entities in the same tier, (3) shift handoffs between crews of analysts working around the clock, and (4) external handoffs outside the company or organization. As external handoffs were outside the scope of the research proposal, these were not included in the bridge to the second study. Shift handoffs were not observed in all teams, as not all observed teams had multiple shifts of analysts to hand off between. Passing within a tier is less common than collaboration within a tier, so there were not as many opportunities to observe this activity.

Escalation handoffs were observed in every team, and were critical to completing the incident response process due to the tiered nature of computer security operations. Thus, *escalation handoffs* were the primary focus for the second study. Team members did not always consider meetings as important across the three sites observed, which was evident in frequency and attendance. Moreover, social factors impacted analyst participation in meetings. For instance, if analysts at the lower tier levels were contractors, they were not always considered as part of the team, and thus not included. Both of these influences impacted how meetings were conducted, and indicated that the observed information sharing process was limited in scope and purpose.

4.2.4.1.3. Critical Functions in Escalations

Within the escalation handoff process, two categories of critical functions were identified: *formal* and *informal*. Formal functions were documented processes that, when asked about escalation, analysts consistently broke down into these steps. These were: *open* ticket, *handle* as much as able, *document* in ticket, *assign* to next tier/analyst, and *close* ticket. Informal functions were not documented procedures or listed by analysts as "steps in the process". However, these functions were often present or sought after by higher and lower analysts during and after an escalation. These were: *collaboration* between analysts during incident handling and escalation decisions (often desired in some sort of synchronous two-way communication), *confirmation* that a ticket assignment was received, and *accountability* regarding analyst performance and incident resolution.

Literature has described the entire incident response process in general terms (Ruefle et al., 2014). The formal and informal functions above provide further detail into the "analyze incidents" and "respond to incidents" elements of Figure 1 in Ruefle et al (2014). Element comparisons of said figure and the findings from Study 1 are shown in Figure 15. Elements from literature are denoted above the dotted line, while elements identified from Study 1 data analysis are below the dotted line. Formal functions are denoted in white boxes; informal functions are in grey boxes.

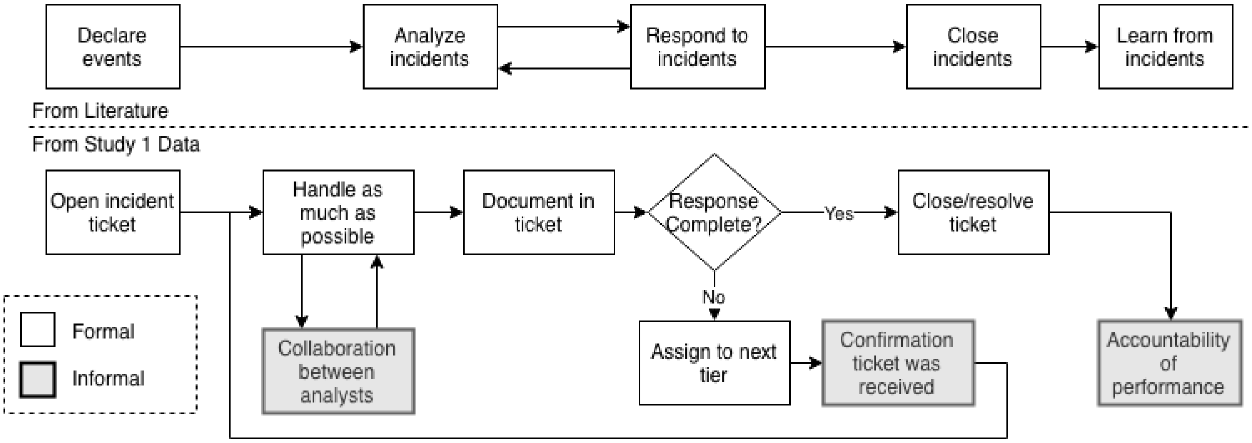

Figure.15. Formal and Informal Functions vs. Elements Identified in Literature

Theme 10 (Section 4.2.1) from the generated observation data directly supports the above findings: *Handoffs are varied in terms of procedure, formality, and documentation. In whatever form, they are important for continuity in several contexts***.** As already discussed, varying forms of handoffs occurred in each of the teams. Colloquially, CSIRTs only refer to shift changeover processes as "handoffs". However, the way a handoff is defined in general indicates that there are other points of the incident handling process that meet the criteria, but are viewed differently by analysts. That is, the amount of attention that an escalation handoff receives from analysts is much less than that of shift handoffs.

Handoffs may happen between teams from different organizations (back shifts managed by a different group, incidents handled by a different organization altogether), as incident response often involves multi-team systems (Tetrick et al., 2016). This increases the 'distance' to the feedback information needed, both organizationally and sometimes temporally and geographically. From the organizational perspective, the external escalation seems to sometimes result in a complete lack of feedback altogether; this could be due to the separation noted above, but also potentially due to the belief that the lower tiers do not need feedback. That is, the separation of authority/level results in a one-way communication channel during escalation processes.

Shift handoff information sharing processes were documented in 2 out of the 3 organizations. Each team had a different shift-to-shift procedure. Team 1 (Location 1) had three consecutive shifts with regular crews. Shift handoffs were performed between these crews, documented, and the handoff documentation was reviewed by a lead. Team 3 (Location 3) did not have a backshift crew, but had some sort of incident handling coverage managed by a separate team in the same company. This team performed a less formal handoff (with less formal documentation) between those two teams. Team 2 (Location 2) had only one shift with no backshift coverage, and thus no shift handoff. The teams who did perform this kind of handoff stated that they held actual meetings or verbal handoffs supported with paperwork for continuity of documentation and accountability.

4.2.4.1.4. Supporting Factors in Escalations

In addition to the functions described in 4.2.4.1.3, other supporting factors were identified that can affect escalation handoffs in different organizations. These factors include: documentation, feedback, communication protocols, organizational distribution, and organizational culture. This section discusses each of these factors in turn.

Inefficiencies and disconnects in handoffs can result in errors in other domains, and have been studied in depth in medicine (Dhingra, Elms, & Hobgood, 2010; Pesanka et al., 2009; Solet, Norvell, Rutan, & Frankel, 2005; Starmer et al., 2014), and acknowledged in aviation (Billings, Lauber, Funkhouser, Lyman, & Huff, 1976). Incident response literature has drawn parallels with other domains, specifically regarding handoffs and error prevention (Steinke et al., 2015). ***Documentation*** is one supporting process in CSIR to help ensure data capture during handoffs, thus creating traceability. Handoffs are often accompanied by some sort of documentation practices to ensure data capture. This is not unlike other domains in which capturing new knowledge has been identified as crucial to responsiveness (Garrett & Caldwell, 2002). Documentation regarding handoffs ranged from using particular dedicated functions in the ticketing system (to express the logic behind the conclusion and following escalation) to separate MS Word documents compiled by the shift crew and reviewed during formal shift handoffs. Escalation handoffs, as the most frequent type of handoff and the most common information sharing event between the operational tiers, did not include formal documentation outside the ticketing system for any of the three teams observed.

Documentation was viewed as necessary in all organizations for *procedural needs*, but necessary in only two of the organizations for *accountability*. Only one organization identified documentation as a *learning tool*, beyond the general acceptance that training documents were required for lower tier analysts. In this particular team (Team 1), the entire crew would run through the shift handoff document (which they created during their shift) to discuss incidents and how they were handled. Moreover, the handling analyst would present each incident, creating some forum for discussion and learning directly to the learner.

Theme 1 identified in Section 4.2.1 indicates that feedback is an important part of information sharing in CSIRTs, and can support escalation handoffs in particular. Within the context of incident response, ***feedback*** is the class of outputs that encompasses communication back down the tiered chain to acknowledge receipt or provide comments in response to performance or actions. It creates shared awareness and improves task coordination (Caldwell, 2008). Feedback can be provided in the form of any of the modes expressed earlier: phone, email, instant message, shared information system, or face-to-face.

Feedback supports informal functions in escalations, including collaboration, confirmation, and accountability. It has also been identified in literature as important to improving collaboration in CSIRTS (Tetrick et al., 2016). Feedback is evidence that two-way communication persists, allowing for more effective system performance. In CSIR context, this means analysts can receive new data, information, or knowledge, and incorporate it into handling and decision making functions. In addition to performance feedback, tactical incident handling feedback also provides a shared awareness that a higher-tier analyst has successfully received an incident, and even how it may be handled in the future. This refers to the confirmation function discussed in Section 4.2.4.1.3. Confirmation, feedback allows the sender to identify who is handling their former task, and, if they are privy to the information, how the incident was resolved. (Note that authority and clearance sometimes dictate how much feedback can flow down to lower tiers regarding resolution.) Finally, feedback provides a channel for accountability, allowing higher tier analysts to provide feedback to lower tier analysts on performance, and, (perhaps through a push notification in the ticketing system), providing a general channel for any analyst to know that a particular incident was resolved.

Feedback in at least one of the above forms was missing from each team. Based on observational data, this lack of response may be due any number of reasons, including lack of protocols to support feedback, organizational disbursement (managerially, temporally, or geographically) (Tetrick et al., 2016), or general culture of accountability and how team-oriented the organization is.

Analysts expressed that communications ***protocols*** were sometimes unclear regarding *who* they should contact in the first place for actions or authority to perform an incident handling task.

Additionally, they were sometimes unsure *how* to contact with respect to channel. Protocols in the form of communication channel and receiver were not necessarily shared at all tier levels (some more experienced analysts had this implicit knowledge, but it may not have been documented), or they were not established. One commonality was that there were “generic” escalations to a tier group, which higher tier analysts would then select from the incoming channel. However, direct feedback to the escalating analyst was not part of this process.

Organizational distribution was also a factor affecting how information sharing occurred. In all of the teams observed, there were analysts that were not collocated with their counterparts, either due to a flexible work schedule (working from home), or due to different parts of the same organization being geographically distributed. In this case, signal/feedback are restricted to certain channels between sender and receiver. Furthermore, if the chosen communication channel is unavailable or “incorrect” in terms of what was appropriate (Garrett et al., 2009), the result can be a broken communication chain.

Referencing VSM diagnosis terminology (Appendix F), System 2 represents the information systems that connect System 1 subsidiaries. In a distributed organization, System 2 becomes even more critical in facilitating information flows to ensure adequate coordination of activities. The VSM diagnosis presents additional insights regarding perceptions of System 2 and effects it has on the larger organization. One important finding from this analysis, corroborated by the other analysis methods in Study 1, was the number of information systems used and the lack of integration between them. This can adversely affect an analyst’s ability to perform well in this temporally sensitive setting.

Lastly, ***organizational culture*** was identified as a factor regarding availability and popularity of providing feedback. Some organizations did not deem this as necessary for various reasons, including the fact that development of lower tier analysts was not a priority, or shared awareness of incidents was not important. As observed in the three teams in Study 1, if feedback was not prioritized or enforced as part of the protocol, then it is less likely to occur. Moreover, if the organizational culture cultivates a mindset of “this is no longer your problem”, feedback is removed from the norms of the team altogether.

In addition to these team-level insights, the VSM diagnosis (Appendix G) also sheds light on the impact of organizational culture on incident response processes. Though not the main focus of the analysis, accounts of how the Systems 4 and 5 operate in the three organizations indicate tensions at the higher levels of management within a firm that could manifest in CSIRT operations. These include unclear priorities about security, insufficient focus on developing security capacity, and lack of autonomy of the CSIRT as it conducts its daily activities.

4.2.4.2. Automation in CSIR Information Sharing: Research Question 2.1

As described in Chapter 3, the intent of RQ2.1 was to understand the current role and availability of automation for information sharing functions in CSIRTs. Answering this question helps improve understanding of the problem space, as well as framing the potential solution space as automation continues to be developed for CSIR applications. The following subsections describe the progression towards the larger RQ.

4.2.4.2.1. What is "Automation"?

Literature in human supervisory control defines automation as "the automatically controlled operation of an apparatus, a process, or a system by mechanical or electronic devices that takes the place of human organs or observation, decision, and effort" (Sheridan, 1992, p. 3). After completing Study 1, data indicated that analysts in CSIR have a different, narrower perception of what automation is. Interviewed CSIRT analysts defined automation as "something that eliminates the human task, or greatly reduces the time to do it", and oftentimes seemed to refer to as an entire system or technology, not necessarily a feature. One key distinction to note is that, when asked to describe automation applications, they often referred to some formal function or task-directed activity, such as filtering, correlating, and executing, and did not mention the informal functions such as coordinating, notifying, or confirming.

Referencing concepts from function allocation in Chapter 2, automated technologies in observed security operations are currently at high degrees (levels) of automation (Sheridan, 1992) for low-level T1 tasks, but low degrees for more advanced (higher tier) tasks. That is, simple, bounded, repetitive or vigilance tasks (that might difficult for humans to sustain over time) have been allocated to technology. More cognitively complex and problem-solving tasks remain the responsibility of human analysts. Observed automated capabilities, mainly involving filtering, include monitoring defined network traffic and detection of rule violation. Additional capabilities identified in literature include automated response with some level of supervision by a human to validate activities (Albanese, Cam, & Jajodia, 2014).

One interesting finding is that, when asked to list current automation tools, most analysts did not include installed systems such as intrusion detection systems (IDS) and related appliances, or the security information and event management system (SIEM), despite the fact that both of these classes of systems are considered automation based on Sheridan's definition above, and market reports indicate widespread reliance on these types of tools (Cisco Systems, 2018). Most responses to the automation questions were more focused on new automation capabilities, such as orchestration and complete Tier 1 response capabilities that could be used to replace low-level analysts and improve systems integration.

4.2.4.2.2. Observed Automation and Available Tools

As mentioned above, some automated tools exist in CSIRT settings based on the adopted definition of automation in the previous section. One well-known application that has some level of automation is IDS, which is programmed to detect anomalies and create alerts based on large amounts of network log data (Killcreece et al., 2003; Schultz, 2012). This class of systems is essentially an alarm system that notifies human operators of potential problems. Intrusion prevention systems (IPS) go one step further and block network traffic that meets certain user-defined or signature-based criteria (Faysel & Haque, 2010). These systems can be host-based or network-based, essentially monitoring two different perspectives of the same network.

Two of the three teams observed had a SIEM to filter and sort incoming alerts from various appliances (sensors), which lower tier analysts would monitor and address. The third team used specialized appliances that were monitored individually. The SIEM can have dashboards (configured by the company) and managed rules to perform its main functions. Some SIEMs can also do base-level data correlation, and are often the main systems used for monitoring system activity, sometimes being displayed on a large shared screen in the SOC. Many analysts have the

SIEM open on their desktop, pivoting between the SIEM and other tools and applications as they investigate alerts. SIEMs are especially helpful in collecting and filtering the vast amounts of data produced by computers and logged by security systems, which scales with the number of endpoints in the network.

Even with these installed systems, there are still tedious aspects of incident response that plague analysts on a daily basis. For instance, some ticketing systems may not have properly defined fields for the business, so the general text box is used for all information that needs to be documented. In so doing, the search features within the ticketing system become extremely confined in terms of searchable data, as many are only capable of searching predefined fields.

Another aspect of automation observed in incident response settings is the user-created tool or function created to streamline the process. For instance, observation data noted scripts made by individual analysts that help “scrape” the ticketing system and create a draft document for shift handoffs. Other instances of scripting were discussed, but with the caveat that these scripts are easily disrupted by software updates, and require maintenance (and the skills) to maintain the functionality of those scripts. Leads and managers brought up future automation applications, such as orchestration (Koulouris, Mont, & Arnell, 2017), as potential solutions for improving incident response operations and coordination efforts. However, none of the teams observed had purchased or installed these tools.

As described in Chapter 2, ‘orchestration’ is a fairly new technology and considered a developing market in incident response. This class of tools, sometimes referred to as security orchestration, automation, and response (SOAR) platforms, are currently sold commercially, and offer a wide range of integration services. A recent report by Gartner describes the need and selection of SOAR technologies (Neiva et al., 2017). These tools were not observed during Study 1 data collection, but were discussed at the senior management level as a potential path forward to alleviating labor concerns.

4.2.4.3. CSIR Information Sharing Challenges: Research Question 2.2

The intent of RQ2.2 was to identify within different CSIRTs where the teams struggle with respect to information sharing. Again, answering this research question helps narrow the scope of Studies 2 and 3. It also provides insights to CSIR managers regarding potential gaps in their processes. Much of the answer to this RQ was informed by the qualitative analysis, and supported by VSM. Study 1 identified three main areas of challenges: feedback, discontinuity of information systems, and need of documentation. While these have been mentioned before as factors affecting information sharing, these issues are further described in this section as pain points in CSIRTs.

4.2.4.3.1. Feedback as a Struggle in Information Sharing

Expanding on the answer to RQ1.1, team members across all teams seemed to struggle with the lack of feedback within the incident response process, or knowing what happened to an incident they touched. *What is received by anyone? Who is working on it? Did they resolve it? How did they resolve it?* Though this is not a specific point in the process, Theme 1 (Section 4.2.1) indicates that it is an aspect of information sharing that is critical to shared awareness and learning. For example, escalations happen because of reasons related to expertise and/or protocols. There were observed struggles about to whom to send escalate a particular incident and how that person should be notified. In the teams studied, these struggle points were often mitigated by protocols, by an analyst or lead with more organizational knowledge, or with self-selection systems in which higher tier analysts choose the incidents they work on from a general list of incoming incidents.

Lack of feedback to create shared awareness has previously been identified as a pain point in cyber defense (Gutzwiller et al., 2015b). Feedback facilitates shared awareness, which allows different levels of the CSIRT and management to fully understand where the incident is in the handling process, and how to react accordingly. Expanding on the use of the term, development-driven feedback was an additional category of feedback that was missing from the investigated teams. From a learning perspective, feedback might educate lower tier analysts on how to approach similar incidents in the future. However, the mode of feedback is an important consideration, as this was found to be a barrier in some of the teams observed. Feedback can be provided in several modes, as discussed above. One such mode is through the information system itself, which for some organizations, seemed to be the crux of the issue. Information systems are discussed in the next section.

4.2.4.3.2. Discontinuity of Security Incident Information Systems

As indicated by Theme 3 in Section 4.2.1, continuity of information sharing was also an issue observed in two of the three teams. For instance, one team had completely separate ticketing systems between tiers, which effectively created parallel information streams (or “stovepiping”) that needed to be managed: one between the lower and higher tier, and one just within the higher tier. The systems house some of the same information (ticket number, incident type, priority), but did not necessarily include other details that lower tier analysts were interested in seeing (how it was resolved). This lack of awareness seemed to cause some struggles for those analysts that were looking to learn more about how those escalated incidents would be approached.

Recalling the VSM background presented in Chapter 3, System 2 is comprised of information systems that help System 1 subsidiaries collaborate and communicate during operation. Lack of integration within System 2 was observed in the CSIRTs studied, indicating potential concerns about viability (presented in Appendix F). Separate systems created not only difficulty in communication, but also a discontinuity in documentation such that the entire “case” of an incident could not be viewed or managed within the same system. This point was a struggle for management and for higher tier analysts looking to gain more holistic views and understanding of incidents.

From a productivity and workload perspective, this system separation seems counterintuitive, and yet this was a business requirement for one observed organization. Needs were expressed regarding why these separations exist, and that there are currently tools are out there to manage those needs (i.e. case management systems). These systems require complex architecture and credential-based user profiles, and can be expensive to purchase and deploy. Due to the cost and effort, completely new systems are not always an option for teams that do not have the resources or explicit needs to justify the investment.

4.2.4.3.3. Inconsistent and Deficient Documentation

Documentation plays a role in knowledge retention (and thus, organizational adaptation to that knowledge) and accountability. Themes 3 and 11 (Section 4.2.1) suggest that documentation is an important aspect of information sharing in incident response. This finding is further supported with the VSM diagnosis, particularly regarding System 3 and 3* that discusses the audit function of the VSM.

While all observed teams documented incidents in some way, the quality of the documentation was inconsistently viewed in terms of importance. This issue has been identified in incident response literature as an overall challenge (Cusick & Ma, 2010; Kurowski & Frings, 2011; Tøndel et al., 2014). Paradoxically, within these organizations, incident documentation was identified as a main source of data for metrics to assess the performance of the security operations centers.

Metrics are one mechanism through which System 3 entities in the VSM can monitor performance of System 1 subsidiaries. Metrics in general were identified at the management level as a struggle, which is consistent with current literature on security operations (Ruefle et al., 2014; Tetrick et al., 2016). Existing metrics were identified on some teams as insufficient, inappropriate, or lacking altogether. Despite the importance of metrics at different levels of the organization, only one team expressed a desire to perform macro-level analysis on security metrics.

Documentation quality was viewed from two perspectives: that all fields were filled out, or that the depth of the documentation was appropriate and sufficient for the incident that was handled. The latter, similar to the idea of quality control, requires more expertise to evaluate, and was not assessed by leads or management in two of the three teams observed. Quality control was not an identified organizational priority in two of the three teams observed. Tensions around documentation and the time-value tradeoff were observed. Poorly designed or configured input forms increased the time and effort needed to establish ‘quality’ documentation. Limited utility and execution of incident meta-analysis were also indicated during data collection. These findings indicate deeper issues within documentation practices and suggest immediate opportunities for addressing analyst needs.

As mentioned previously, only one team had documentation addressing shift handoffs, but no teams had documentation addressing escalations aside from the ticketing systems. This may be sufficient in the future should some of the above issues regarding separate systems and lack of feedback be resolved. A formal assessment of these needs presents ideas for future research.

4.2.4.4. CSIR Outputs: Research Question 3.3 and 3.4

The intent of RQs 3.3 and 3.4 was to identify, from a systems perspective, the classes of outputs of the overall process, and the respective stakeholders of those products. Answering these RQs helps define the current state as well as potential expectations for future improvements or new systems. The following subsections are dedicated to answering each segment in turn.

4.2.4.4.1. Types of Outputs

Three types of outputs were identified for the overall process using workflow diagrams and the VSM methodology. Incident handling outputs are outputs of the entire process. Escalation outputs are the products or messages that are delivered as a result of escalating an incident from a lower tier to a higher tier. Finally, feedback outputs encompass communication back down the chain to acknowledge receipt or provide comments in response to performance or actions.

The primary ***incident handling output*** is a contained incident. Ideally, this incident would also be thoroughly investigated, eradicated, and steps would be taken to add knowledge to the information systems for prevention. However, this is not always the case. A major expectation (from management) around these outputs is that the incident is handled as fast as reasonably possible. An informal output is knowledge regarding an incident, where it came from, how to handle, and how to prevent. Again, this informal output is not always captured, which will be discussed further below.

Escalation outputs are specific to the handoff between one tier/group and another. In particular, these outputs include an incident ticket, reference to the ticket, and information relating to the incident from T1 analysis. Expectations of these outputs include proper escalation (actions by the higher tier level are warranted), and timely escalation, especially if the T1 analyst does not have the authority to take containment actions.

One expected output that was missing from all representations was ***feedback***. Within the escalation process, feedback from the receiver to the sending analyst (message that the incident was received, or if the information provided was sufficient for their use) was missing overall. Depending on the organization, this feedback could be complex considering the geographic and, sometimes, organizational distribution of sender and receiver. In general, channels exist to support this kind of feedback, but it is generally not utilized. Within the overall incident handling process, feedback in the form of knowledge was also missing across the sites. While some had mechanisms to support knowledge sharing within tiers/groups, the knowledge sharing across these tiers or groups was relatively weak. That is, post-incident reviews, or how an incident was handled, were not observed during data collection. Evidence that these exist was identified at one location, but the scope of the meeting invitees was limited to include only one representative from T1, who was then responsible for sharing with shift leads.

4.2.4.4.2. Customers of the outputs

The primary and direct customers of incident response outputs are ***team leads***; leads may or may not review completed tickets during a shift to ensure quality. The leads do not necessarily use the outputs, aside to assess the quality of their shift’s performance. Additionally, different levels of management comprise the largest group of ‘output users’ (security operations, director, and chief information security officer). ***Management*** uses incident response data in different forms; for example, high levels of management may receive raw incident data directly from an analyst. Alternatively, they may receive amalgamated data of all incidents in a particular category or time frame to understand the larger security profile of incidents coming in and how they were handled. According to participants, the use of amalgamated versus raw may be tied to organizational maturity. Lastly, ***IT groups*** also receive incident response outputs in the form of tasks to be completed for full remediation. Within their shared ticketing system, IR analysts may assign actions to different segments of the IT organization for this reason.

The recipients of escalation outputs are mainly ***higher tier analysts***. As the incident is worked on, more information will be added to the ticket, which becomes the passed product during escalation handoffs. Higher tier analysts use these outputs to continue the investigation. Management can also be a recipient of escalation outputs, as mentioned above. In some cases of escalation, management is notified through email with details of the incident up until that point. Management may just become aware of the incident, or the email may trigger intervention or action.

Lastly, ***lower tier analysts*** are the would-be receivers of feedback outputs should they be produced. These outputs would be used to increase lower-tier knowledge and performance, as well as complete communication loops across the levels of the organization. Given the analyst shortage and the substantial gap between lower and higher tier analyst skills and expertise, feedback could be especially important in increasing overall capability of developing analysts as they progress through the tiered response organization. While higher levels of analysts and management could benefit from feedback, the observations indicate that there is general lack of information backflow in the process that could most benefit lower-tier analysts in their growth.

4.2.4.4.3. Improving outputs

Based on some of the above descriptions, there are several ways in which outputs could be improved. Though cross-organization consensus of these improvements was not established as part of the study, the resulting list of considerations and insights for managers is presented in Chapter 7 as a way of allowing organizations to select what is useful to their respective environments. These recommendations should also be considered within the industrial development community as security settings change with respect to technological evolution. These potential improvements include decision verification, quality checks, and feedback improvement.

One output not explicitly described is the decision to escalate an incident (which results in a ticket being assigned to a higher tier). One potential improvement is ***decision verification*** to ensure that the escalation was appropriate (for the higher tier or, for those organizations that have separate groups, for the entity receiving the escalation). While this could create more work for higher tier analysts (or the chosen verifying party), experts from Study 2 indicate that verification is standard practice at higher levels of incident response, and that this activity may already occur. Furthermore, should automation continue to be developed for Tier 1 incident response, verification will be needed in order to train machine learning algorithms regarding the appropriateness and quality of escalation in terms of channel, receiving analyst, and so on.

Next, ***quality checking*** of incident response tickets could be improved to ensure that the information collected and used for decision making was correct and complete, which helps validate novice analyst investigation skills. From a technical perspective, this might be especially beneficial to those teams who evaluate the ticket quality based on completeness alone. Review protocols driving these outputs to management might also be improved to filter and focus outputs that reach Systems 3 and higher (referencing VSM terminology). While covering management on all incidents can increase their awareness, it also increases the signal-to-noise ratio at higher levels of management. The VSM indicates that this particular type of effect does not support the overall system viability (Beer, 1984). By creating more specific attenuation through these channels, management may be more effective in driving system adaptation.

Lastly, ***feedback*** could be improved altogether, either inside information systems or outside, to ensure complete communication loops and improve the knowledge and performance of lower levels of incident response. Feedback improvements can be incorporated into automation development or installation within an organization, but also in general protocols between humans if automation is not available. Feedback should be considered not just a completion of a communication sequence, but also as an opportunity for learning and development. Improvement to System 2 integration and focus may naturally facilitate better feedback.

4.3. Study 1 Summary

Results of Study 1 helped narrow the scope of Studies 2 and 3 by identifying critical areas of information sharing in CSIR, as well as other contextual factors that influence their performance. Three teams in different settings were included to help capture environmental differences that affect a team’s mission, structure, communication, and use of technology in relation to incident handling. Results indicated deeper implications of these relationships, especially as it relates to overall security effectiveness and organizational maturity. Outputs from Study 1 informed the task selection for Study 2, as well as key considerations for developing new operational concepts in Study 3.

CHAPTER 5. COGNITIVE TASK ANALYSIS WITH CSIR EXPERTS

Chapter 5 focuses on the second of three studies presented in the research conceptual framework. The data collection method was a Cognitive Task Analysis; this interviewing method with experts was used to elicit expertise regarding escalation handoffs and general information sharing in CSIR. Analysis methods included a qualitative coding exercise with two raters to categorize results within a dimensional expertise framework. Data collection and analysis details are described in Section 5.1. Results from Study 2 answered two (2) of the research questions, and validate findings from Study 1, specifically regarding RQs 3.3 and 3.4. This set of research questions determined expertise needs in CSIR, specifically in relation to information sharing. Findings and answers to these research questions are presented in Section 5.2.

5.1. Data Collection and Analysis

5.1.1. Sampling and Recruitment

5.1.1.1. IRB Approval

A separate IRB protocol was created for this study to accommodate a different population from the first study and the use of different methods. The protocol was largely based on the ACTA methodology from Militello & Hutton (1997), and defined the population of interest as experts with 5+ years of experience in computer security incident response. The IRB package for this study is located in Appendix H. After the protocol was approved (IRB protocol 1802020208, approved 03/09/2018), recruitment commenced. The original intent of the protocol was to complete the interviews in person to increase camaraderie and participant comfort with the researcher. However, after 5 months of stale recruitment, the protocol was amended to include video conferencing as a medium for data collection (amendment reference 1802020208A001, approved 01/23/2019). This greatly alleviated recruitment issues, and allowed the study to progress accordingly.

5.1.1.2. Recruitment of Experts

Experts were recruited using two main methods. First, I attended a security symposium at Purdue University in April 2018 to present the contents of the proposal for this dissertation. During this event, experts attending the symposium were recruited for the study and became contacts for snowball sampling. Many cyber security professionals expressed a level of hesitation in participating in a recorded interview study, even after IRB approval. In some cases, their respective employers would not allow them to participate, despite the fact that the study did not focus on particular companies or sectors. This hesitation became a major inhibitor to recruitment. Thus, the second recruitment method employed was snowball sampling, which is commonly used to gain access to hard-to-reach populations (Bernard, 2006; Pelto, 2016). The participants who had already completed the interview became connections in the industry that vetted the study to other cyber security professionals, making it easier to gain credibility and trust with potential participants.

The majority of participants were approached about the study through face-to-face interactions, after which the prospective participant would share personal contact information. Formal recruitment then followed with a direct email from myself to the participant, either through standard email systems or through networking platforms with direct peer-to-peer messaging (such as LinkedIn). The purpose of this email was to share formal recruitment language approved by the IRB, and present more information about the study. Due to some of the concerns highlighted in this section, generic, wide-scale recruitment through social media may not be effective for this population.

Email and direct messaging correspondence often resulted in a short, scheduled phone call to further explain the study, needs, and assurance about confidentiality. The phone calls lasted no more than 20 minutes, and from that point, the prospective participant could decide whether or not they wanted to continue with scheduling an interview. Scheduling was coordinated through email. WebEx was used for participants using web-conferencing. All participants were asked to provide a signed consent form before recording started.

5.1.1.3. Sampling Approach

Wide sampling is an important aspect of validity in qualitative research to help increase generalizability of findings (Leung, 2015; Merriam & Tisdell, 2016). Qualitative literature in grounded theory commonly points to sampling strategies with respect to the population size, or with respect to data saturation (Goulding, 2002). However, as the goal of this study was not to develop grounded theory, but rather to validate known themes in the domain, other literature bodies were consulted for guidance regarding sampling. Though an exact number is a contentious subject amongst qualitative researchers, some recommendations from seasoned veterans in the field recommend between 6 and 12 participants (Baker & Edwards, 2012).

Literature for CTA literature does not offer guidance with respect to sampling, but does note that data collection is extremely costly, and that access to experts is hard to achieve (Zachary, Techologies, Crandall, Miller, & Nemeth, 2012), resulting in a "take what you can get" approach. Time constraint was a risk identified in CTA literature (Crandall et al., 2006) as one of the main challenges in getting expert participants. This constraint ultimately did affect the recruitment for this study. Smaller numbers of participants are common in studies that employ CTA (Plumptre et al., 2017; Read, 2013; Roberts, Flin, & Cleland, 2016; A. R. White, 2019; Yates et al., 2012), but there is no "right answer" (R. R. Hoffman, 1987). Furthermore, other types of studies that involve experts typically include a small number; (Nielsen & Landauer, 1993) recommend a sample size of five (5) with respect to expert evaluators in order to achieve a high cost-benefit ratio of cost of time and participation to number of issues identified.

Based on the above, the initial goal was to recruit between 4 and 6 participants. It took 7 months to recruit and collect data from 5 participants, all while employing snowball sampling to help access the population. Recruitment was extremely difficult in part due to the level of caution adapted by many security professionals. Moreover, participants were not compensated for their time, and interviews ranged from 90 to 120 minutes – a relatively long period of time for an important expert to be engaging in non-productive work.

Three out of five participants were contacts made during the security symposium discussed in Section 5.1.2.2. Follow-up for participating in Study 2 occurred between August 2018 and

February 2019, and usually required multiple emails, advanced scheduling, and a high degree of flexibility for cancellation. Each of the five participants had between 7 and 38 years of experience in cyber security roles, and had experience from two or more sectors explored in Study 1 (industry, academia, government). During interviews, some participants commonly referred to other experiences that helped them in cyber security; these were not prerequisites to participation. A table of these facts can be found in Table 8 below.

Table.8. Summary of CTA Participant Backgrounds

Participant	**Years in Cyber Security**	**Government**	**Industry**	**Academia**	**Related Experiences**
1	10	Yes	Yes	Yes	
2	9	Yes	Yes	No	Law Enforcement
3	7	Yes	Yes	No	Military
4	38	Yes	Yes	Yes	
5	15	Yes	Yes	No	Military

In order to narrow the scope of the CTA, I selected the largely structured ACTA protocol (Militello et al., 1997). This selection ensured that interview questions would be very specific with a particular task area, but still wide enough to include the abstract aspects of the task targeted for this study. That is, keystroke level task breakdowns were too specific for the scope of this dissertation; ACTA could capture the larger context and environment that drives expertise in security incident response.

5.1.2. Data Collection Procedure

For in-person interviews, I met with participants at some agreed upon location that was convenient for the participant. These included offices, coffee shops, or conference rooms. For remote interviews, a WebEx conference call was scheduled to allow video capability. This was not efficient in all instances, as participants would call from their cubicle or desk, which was shared or in an open office space. While Cognitive Task Analysis is preferred to be face-to-face for building rapport and better communication and interview cues, there is no evidence that in-person is a requirement. Precautions were taken to verify visuals created during interviews with participants, verbally or visually.

The interviews began with general greetings and introductions, and verification that a consent form was signed and submitted. A recording device [Olympus digital recorder WS-852] was used for all interviews, verbalizing when the recording would start or stop. The recorded content started with some background about the person's experience in cyber security, including how many years in the industry, which sectors they worked in, and what roles they filled during that time. This was done to get some context regarding the person's responses and perspectives, but is not included in reported data as a precaution for confidentiality.

Next, I reiterated to each participant that the goal of the method: to understand more specific expertise applied up to and during escalations, and set the scenario as Tier 1 to Tier 2 incident response activities and interactions. From this point, provided job aids from the ACTA Methodology (Militello et al., 1997) were used. Lastly, additional questions were formulated that were more specific to communication and included them in the overall interview set to help provide specific context around communication that might be left out of general task descriptions. The interview aids developed by me, along with the ACTA aids for the task diagram, knowledge audit, and simulation interview, can be found in Appendix I. Each interview concluded with thanking the participant, and verifying that there might be some short follow-up questions via email.

5.1.3. Data Processing

Within 24 hours of each interview, the .mp3 file was downloaded to a personal computer from the device, and uploaded to secure cloud storage. Audio files were then transcribed through a paid third party service (Rev.com). After transcriptions were completed, the MS Word documents were downloaded and listened through the audio file while reading the transcription. This was done to verify correctness and completeness of transcription. This was also an opportunity to remove identifiable indicators (personal references) from the transcription.

The first segment of ACTA included a diagramming activity that involved the interviewee depicting the task of interest in the form of an abridge diagram, which is further elaborated through discussion. The task diagrams drawn during the interviews were digitized with draw.io (an online tool for creating visualizations). Task diagram content was supplemented with content from the interview, in case details from the process were missed during the drawing activity. Task diagrams can be found in Appendix F.

The ACTA method uses Cognitive Demands Tables (CDTs) to condense a 90-120 minute interview (with the three aforementioned parts) into a digestible and usable format for analysis. To create the Cognitive Demands Tables (CDTs), I started with extracting out elements from the interviews that participants (experts) said are difficult for novices. Transcripts were used to populate the details of the table. As some items came up multiple times throughout the Task Diagrams, Knowledge Audit, and Simulation Interview, I employed iterative processing to fill in the table such that all inputs were represented. This was completed for each interview, with separate CDTs for each participant. These CDTs can be found in Appendix J.

5.1.4. Analysis

Analysis for Study 2 employed a top-down coding scheme based on the dimensions of expertise framework in (Garrett et al., 2009). As mentioned, these codes were identified a priori, and were chosen based on applicability to the knowledge work conducted by security analysts and the holistic approach to expertise described in the six dimensions of expertise framework. Much focus in security literature is on identifying specific subject matter expertise, or broad scale knowledge, skills, and abilities (Assante & Tobey, 2011; Bishop et al., 2017; Newhouse et al., 2017), which are ever-expanding lists meant to help with hiring and training. Business literature has pointed to broader hiring strategies regarding specific attributes that ‘can’t be taught in a classroom’ (Cobb, 2016; van Zadelhoff, 2017), but attributes such as curiosity, ethics, and understanding of risks are difficult to quantify or assess in a standard interview and hiring process.

Though these approaches to human expertise requirements are both useful in their own respects, this dissertation focuses on applying expertise to automation development. It is critical to create alignment between expertise frameworks and the area to be automated to ensure fluid translation into technology requirements. Dimensional approaches to expertise are generally uncommon, but are used here to help understand well-rounded expertise needs in cyber security and focus development efforts.

The codebook for Study 2 analysis was developed from Garrett et al (2009), with additional examples given to help coders deduce the code definitions. The codebook was reviewed by one of the paper's authors to verify alignment and consistency with the original article. To ensure thorough understanding of each code (Goodell, Stage, & Cooke, 2016), codebook definitions included specific examples, and, if applicable, exclusion criteria, for each code. The codebook can be found in Appendix K.

CDTs acted as the data to be analyzed for Study 2. In order to increase trustworthiness of findings, a second rater (in addition to myself) was included for data analysis of the CDTs. Both raters were CITI certified for human subjects research with prior experience in qualitative data coding. I also developed a training procedure to ensure alignment and consistency between the two raters (Goodell et al., 2016), which included a description of the research background and instructions on how data was to be handled during and after analysis. Training lasted 45 minutes, and also included a "practice round" with CDTs developed in other research to acclimate both raters with the format and presentation of data. The examples were used as an exercise to code and discuss thoughts and alignment for each code. Training documentation can be found in Appendix L.

5.1.4.1. Coding Procedure

I chose a fully crossed design (Hallgren, 2012) for coding, meaning that both raters (which included myself) coded all participants' CDTs Crandall et al. (2006) suggest in their CTA guide that multiple coders are recommended to ensure "soundness of research method and conclusions drawn" (p.100). With two raters and five participants, a fully crossed design was feasible and allowed me to assess systematic bias between coders (Hallgren, 2012), thus improving the overall IRR estimate.

After training, the other rater and myself independently coded the participant CDTs, using the codebook as a guide. We both used de-identified CDTs to code separately. The decision regarding medium was left up to the rater, though instructions for both electronic and manual were included in the training documentation.

The CDTs were organized in a table format, and within each cell, text was broken down into smaller segments separated by a semicolon. Each of these segments was a statement translated or summarized from the interviews, and acted as a unit for coding. We each coded the segments using the codebook as a guideline. Units could be coded with more than one category (that is, segments could have any number of codes that applied). We each drew a box around the segment (or multiple segments, if applicable), and labeled the box using the designation in the codebook for the original six dimension of expertise (C1 – C6), and the two additional codes I added (C7 – C8) to reflect findings from literature and studies in this particular context. Any additional codes identified by either rater were denoted with an asterisk (*) or keywords and were written in the margins with lines connecting the note to the applicable segment. Example of a coded CDT can be found in Appendix M.

When complete, both sets of documents were collected and the coding results were transcribed into a spreadsheet format. Extra codes and notes were also copied into the same spreadsheet so that they could later be grouped and analyzed alongside the rest of the codes. Different examples of coding tables and schemes were evaluated from literature (Crandall et al., 2006; Pelto, 2016); CTA literature commonly applies schemes much like the CDT from ACTA to help organized data. A sample of the spreadsheet can be found in Appendix N.

Next, a joint coding exercise between the other rater and myself was conducted for the extra codes that did not fit completely into the codebook. We had a teleconference with a shared spreadsheet that allowed for synchronous edits. Notes and comments were incorporated in to the exercise. We discussed each extra code in turn, creating secondary codes to group primary codes. Finally, we agreed on themes to represent the code groups accordingly.

5.1.4.2. Inter-Rater Reliability

Another way to define reliability is the extent to which a set of scores is random (Frey, 2016), or how much of the variance is due to variability in participants being scored. To establish trustworthiness, a second rater was invited to code the data, and assessed inter-rater reliability (IRR) between both raters. As percent agreement is not considered an acceptable measure of inter-rater reliability (Hallgren, 2012), Cohen's κ (Cohen, 1960) was used as the main measure. This coefficient was designed for fully crossed design with exactly 2 raters, and includes probability of agreement by chance in addition to rates of agreement. In order to compute κ, a contingency table was created to compare ratings by code and by rater (Table 9). As shown, the contingency table is a 6x6 table in which full agreements are tallied in the diagonal and disagreements are tallied by rater and by code.

Table.9. Contingency Table of Agreement

	Rater 1						
		C1	C2	C3	C4	C5	C6
Rater 2	C1	61	1	1	8	5	26
	C2	0	49	1	0	0	3
	C3	3	3	20	2	0	5
	C4	4	7	1	60	2	4
	C5	6	0	0	3	30	7
	C6	2	3	0	5	0	56

5.1.4.2.1. Assumptions

In order to consistently apply agreement and disagreement evaluation, I made some assumptions about certain scenarios that presented uncertainty. I applied these assumptions in efforts to further increase the conservativeness of the overall result.

1. If there was any code agreement per segment, the agreement for that code was tallied (ex: C1 vs C1). This also applied for more than one agreement (ex: C1 C2 vs C1 C2 would result in two agreements for the respective codes).

2. If a segment was coded by only one rater, the segment was removed from the overall count of agreements and disagreements. These are hereto referred as “orphan codes”.
3. If a disagreement occurred between two or more codes, disagreements were tallied piecewise (ex: C1 vs. C3 C4 would be a disagreement between C1 and C3, and C1 and C4)
4. If a code outside the framework was applied (C7, C8, or *), agreements or disagreements relating to those codes were ignored.

5.1.4.2.2. IRR Calculation

The κ coefficient for this dataset between two raters was $\kappa = 0.51$, or “fair agreement”. According to literature on inter-rater reliability, this result indicates “moderate agreement” (Landis & Koch, 1977), but should not be used as a predictive measure (Krippendorff, 1980). The interpretation of κ varies from literal (it is what it is) to more conservative cutoffs ($\kappa > 0.80$), but depends on how κ is being used. In the case of this study, κ is not being used to predict ratings or conduct hypothesis testing, but rather understand how consistently different raters applied the six dimensions of expertise within each participant dataset.

The κ coefficient is already considered a conservative measure, and asserts that, the closer κ is to 0, the more chance or probability is involved in ratings. However, the above results indicated that further training and use of the codes could help increase consistency of ratings between and within raters, especially for more complex coding frameworks like the six dimensions of expertise. From the Table 9 above, the disagreements show consistent discrepancy between the use of C1 and C6 for the two raters. A high κ value could be an indicator that further time and materials for training is required, and additional clarification is needed regarding inclusionary and exclusionary criteria for these codes. Further exploration of this discrepancy is discussed in Chapter 7.

5.2. Results

5.2.1. Qualitative Research Outputs

5.2.1.1. Summary of Codes

Figure 15 indicates summarized code counts for each of the codes in the codebook, including policy and self-awareness. These counts are indicators of frequency in the interviews as interpreted by the two raters, and include single votes (one rater coded) and double votes (full agreement between raters, or a double vote) counted each as a single instance. Due to the nature and scope of interview, analysis results could not definitively conclude that frequency indicates actual importance of particular dimensions of expertise within incident response. However, the frequency may suggest *perception* of importance amongst experts who have deep experience in the field.

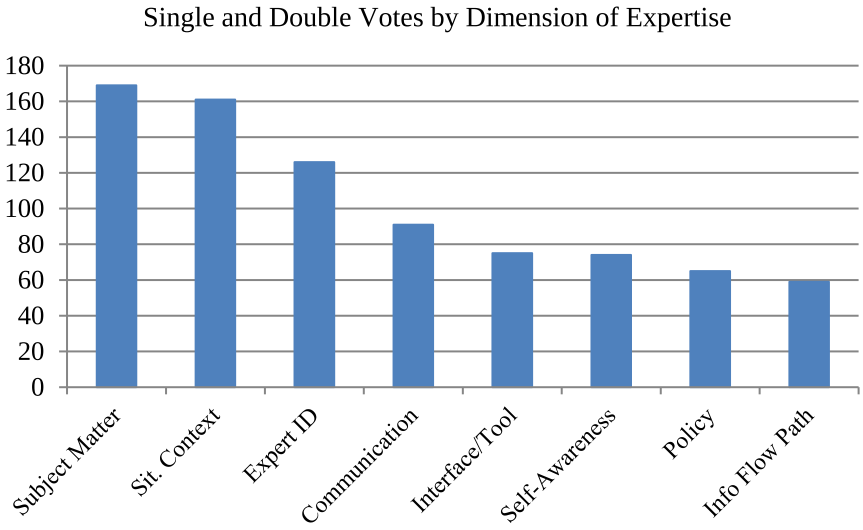

Figure.16. Rater Vote Comparison of Codes

Figure 16 shows a comparison of votes across the eight (C1 – C8) codes used in data analysis. Subject matter expertise (169 votes) and situational context expertise (161 votes) were the top two dimensions, followed by expert identification expertise (126 votes) and communication expertise (91 votes). Interface or tool expertise (75 votes), self-awareness (74 votes), policy (65 votes), and information flow path (59 votes) were relatively close in frequency.

5.2.1.2. Developing New Themes

As previously mentioned, there were two codes identified in addition the original six dimensions framework: policy and self-awareness. These codes were included in the independent coding exercise, and themes were developed based on all data related to these codes. Additional codes identified by the raters were grouped and analyzed after individual coding was complete. We used an online, shared spreadsheet such that changes could be made by either rater real-time, including sorting, notes, comments, or formatting. We discussed each keyword in turn, and determined a shared term to describe the concept. After this step, the terms were grouped into larger themes, and added a description to further elaborate on the concept identified from the original CDTs. The new themes are listed below in Table 10. In order to validate the two extra dimensions from Study 1 (policy and self-awareness), data from Study 2 was used to add more clarity to these dimensions, which are also included in Table 10 below.

Table.10. Dimensions / New Themes Developed from Study 1 and Study 2 Interviews

Theme	Extra Code	Theme Description
1	Job Scope/ Boundaries	Organizational structure and culture can impact how individual expertise is shared, developed, or cultivated within the environment
	Collaborative Culture	
	Org/ Policy	
2	Collaborative Problem Solving	Collaboration is a bidirectional process in incident response problem solving and learning processes, and is facilitated by shared awareness/common operating picture and networking within the SOC; Expertise sharing and development is enabled by collaboration practices and environment.
	Collaborative Learning	
	Shared Awareness	
3	Person Trust	Deeper social factors between entities, such as trust, prior experience, and professional relationship, affect how C2, C3, and C4 types of expertise are developed and utilized.
	Communication / Relationship	
4	Readiness	There's an additional aspect of expertise that is built upon aptitude, thinking style, and self-facilitated learning in CSIR; This impacts how novices overcome naiveté (over-trusting data/systems) and employ more flexible (non-linear) problem solving.
	Continuous Learning	
	Info Trust	
	Linear Problem Solving	
C7	Self-Awareness	Self-awareness relates to Theme 4, and is important for a "feedback loop to self", or the ability to self-evaluate current state and a need to adjust or seek assistance
C8	Policy	Policy is related to Theme 1; Many firms try to remove the need for C1, C3, C4, and C6 by creating procedures and policies, which effectively removes some amount of need for applying expertise for decision-making. Policy can force information paths and escalations, create rules/guidelines based on context, and provide clear procedures to follow.

5.2.2. Findings

Two research questions addressed in Study 2 directly highlight aspects and dimensions of expertise within the CSIRT setting. The first question (RQ1.2) focuses on the original six dimensions of expertise, while the second question (RQ3.1) broadens this to include other aspects of expertise that may not be included in the six dimensions framework (policy and self-awareness). These questions and the related findings are presented below.

5.2.2.1. Dimensions of Expertise in CSIR Information Sharing: RQ1.2

Research question 1.2 aimed to answer which of the six dimensions of expertise was required to perform escalations (as the main information sharing task). The intent of this question was to help identify and potentially prioritize expertise requirements in order to 1) help understand analyst needs beyond the existing literature on knowledge, skills, and abilities (KSAs), and 2) to frame requirements for technology as its capacity transitions from tool to teammate. Answers to this research question provide insights to CSIRT managers regarding analyst development and informs the analysis of Study 3.

Findings of Study 2 indicate a difference in representation across the dimensions of expertise and additional codes used for the analysis. I could not identify previous efforts in literature to identify or assess dimensions of expertise in real-life settings, though application in simulation models (Nyre, 2016; J. D. Onken, 2012) did lay out a framework for the measurement aspect and attempt to connect the framework to performance outcomes. The ranking and representation of the six dimensions from Study 2 validate the importance of expertise altogether in novice roles of incident response, indicating that further investigation and development is warranted in this domain.

All of the original six dimensions of expertise (C1 – C6) were represented in the data, with different strengths; the two extra codes (C7 – C8) were also represented. Not surprisingly, subject matter expertise and situational context expertise were most strongly represented. Subject matter expertise is commonly cited in cyber security literature in relation to the skills shortage (Assante & Tobey, 2011; Bishop et al., 2017; Cobb, 2016; Ruefle et al., 2014), and situational context is a key aspect of incident response that is already being investigated as cyber situation awareness (Healey et al., 2014; Mancuso, Minotra, Giacobe, McNeese, & Tyworth, 2012; Oyewole, 2016; Tyworth, Giacobe, & Mancuso, 2012; Vieane et al., 2016). Within cyber security, experts indicated that situational context in particular is required to assess all relevant signals and determine the appropriate action going forward. Furthermore, they stated that this is often a challenge at the lower-tier level due to physical, policy, and other dimensional expertise limitations. This finding supports that the idea that dimensions investigated in isolation may not adequately identify correlations and connections to other dimensions or relating factors.

Data indicated that expert identification expertise, or knowing where (or with whom) to find information, was also identified as important for T1-T2 investigation and escalation. Within incident response teams, expertise is often distributed amongst the organization (T. R. Chen et al., 2014; Ruefle et al., 2014) or embedded within different information systems or databases, making navigation of those resources critical to timely and accurate response. If a novice does not have adequate wayfinding strategies, techniques, or resources, the organization runs the risk of interrupted signal (not getting the right people involved) or delayed response, which affects the baseline metrics of the SOC.

Communication expertise was commonly mentioned in relation to daily tasks for analysts. This finding aligns with recent literature that highlights the need to expand beyond traditional KSA approaches to expertise analysis (Cobb, 2016). Though there are messaging tools that help analysts overcome physical, geographic, and temporal barriers to communication, participants indicated that basic expertise regarding 'communication skills' are necessary for effective coordination when communicating to entities inside and outside the SOC and larger organization. Issues connected to lack of expertise in this dimension included improper escalation, inappropriate language (i.e. too technical, not technical enough, emotional) to a given receiver, inappropriate detail to a given receiver, or insufficient coordination techniques (i.e. not asking or telling in an effective manner to get the needed information).

5.2.2.2. Expertise Requirements for Escalations: Research Question 3.1

Research question 3.1 intended to expand the scope of RQ1.2 to include other aspects of expertise required for T1-T2 escalations that might not fit well into the six dimensions framework. These additional expertise requirements should be considered context-specific, and not generalizable to every type of incident response team beyond computer security. The following section describes these themes, identified during data analysis of Study 1 and Study 2.

In addition to the findings indicated for RQ1.2 above, there were two additional aspects of expertise that were uncovered during investigation and analysis bridging from Study 1 to Study 2. These codes, referenced as *self-awareness* and *policy* did surface in the expert interviews as relevant aspects of incident response at the T1-T2 levels. Self-awareness was referenced more commonly in the interviews, especially in relation to learning, performance, and interacting with other analysts. Self-awareness was previously identified during Study 1 as a quality that is needed to learn quickly and remain adaptive in this environment. This code was not part of the original framework, and might not be its own dimension of expertise. However, it is conceivable that this attribute is the needed feedback loop to evaluate and improve a given skill. That is, self-awareness is the ability of an individual to objectively evaluate one's environment, actions, and performance. This code could apply to any dimension of expertise, and is thus a 'meta' concept of the dimensions of expertise.

Policy was also mentioned often throughout the interviews. More specifically, it was emphasized by experts that one needs to know what the policy is and how to interpret it correctly in a given situation in order for his/her actions to be consistent with the company mission. Many novices do not know policy or interpretation thereof, thus it tends to be integrated through procedures as institutionalized knowledge. However, participants indicated that the explicit recognition of policy within procedures is not a required part of T1 training, and learning why certain policies exist is not common until later stages of analyst development.

Policy is decided at the upper echelons of organizations (Hoverstadt, 2010; M. C. Jackson, 2000), and may be influenced by environmental and legal stressors. Policy can be complex and undefined regarding reasoning and potential impact if not followed. Experts indicated that it is inappropriate for novices to try to interpret policy at the lowest level of incident response under high temporal pressure, especially because they may not know the impact of their own decisions during mitigation and response. This is a significant driver in creating procedures to standardize response based on rules created at higher levels of the organization, and perhaps a reason why novices are not presented with the opportunity to understand policy at a deeper level.

5.2.2.3. Validating Findings for RQ3.3 and 3.4

Study 2 findings further support answering RQs 3.3 and 3.4 that were addressed in Study 1 (Chapter 4). Expert interviews regarding expertise during escalation provided some validation and verification regarding task outputs, customers, and potential improvements, which are discussed below.

5.2.2.3.1. Validation of Expected Outputs

In addition to validating the findings in Study 1 regarding steps in the incident response process, Study 2 data supports additional expectations regarding escalation task outputs. These include decisions, investigative materials, and completed checks. Specifically regarding steps up to an escalation, experts indicated that there were three main outputs. First, the investigation involves a series of decisions made by the novice analyst regarding what they should do with a particular alert. Sometimes these are guided by procedures; other times, analysts are left to apply their expertise with no aid. The decisions might include:

- "Is this a signal or noise?"
- "Can I handle this myself?"
- "Should I handle this myself?" and
- "How should I handle it?"

As part of the investigative process, analysts may check various systems, alerts, open cases, and user data to understand the incident. All of this information should be included in the ticket to provide an audit trail and traceability for other analysts who receive the handed off ticket. Part of this process includes filling out all necessary components of the ticket, and verifying or checking specific sources or pieces of information. To a certain extent, the receiving analyst may review all three of these components as part of their self-briefing process. However, it is not guaranteed that these outputs will be checked or verified.

5.2.2.3.2. Validating Customers of Outputs and Potential Improvements

Study 2 participants indicated that the customers of the outputs are usually determined by policy, procedure, and organizational setting (internal vs. external customers). For CSIRTs in companies that are classified as managed security service providers (MSSPs), there may be a number of external customers outside the SOC and outside the firm that need to be notified or involved in an incident response activity. Experts explained that the MSSP receives a notification to investigate, and that it sometimes requires external customer inputs to help navigate their specific infrastructure and policies. Internal receivers of escalation include higher tier analysts or specialists, which can be embedded within procedures. The most common type of escalation involves an internal customer, who is often the next tier-level analyst who is monitoring a given email account for escalated tickets. This could be rotational, meaning the receiving individual is not always the same.

Outputs of the task are typically checked to ensure the criteria discussed in RQ3.3 are verified: sound decisions, correct materials, and complete record of events. Study 1 indicated that not all organizations check tickets for these items, especially due to high volume and low analyst availability for quality checks. To compensate, some organizations conduct ticket audits instead of full-scale quality checks. One potential improvement is to automate these quality checks, as rule-based comparison is within automation capability, and this could increase wide-scale ticket quality and confidence. Criteria such as completeness and material verification could be relatively easy to check using an automated system. Decisions could be verified by comparing to procedures (if they exist), or playbook libraries.

Two other potential improvements in output delivery are decreasing the time to delivery and improving the language and presentation style to the customer. While the data indicate that these are not high priority, they could improve overall efficiency as it relates to overall process and communication. Escalated tickets could sit in an email queue for some amount of time before being worked by the next analyst. Considering the above statements regarding quality, time to delivery of a *quality response* adds another level of verification and validation that could slow down the overall mean time to respond. Finally, participants indicated that communication expertise is an important dimension for analysts to develop, especially because of the variety of people with which they might interface during response. Improvement to this aspect of escalation may help information flows be more efficient and effective, as it could reduce the time to ‘decode’ a message, and thus the time needed to react.

5.3. Study 2 Summary

Results of Study 2 provided validation for findings in Study 1 regarding the incident handling process and challenges in information sharing observed in different CSIRTs. Experts also indicated that a wide range of expertise is needed to perform escalation and related information sharing tasks beyond traditionally recognized subject matter expertise. Considering that subject matter expertise is often the focus of training education programs for cyber security, these findings corroborate some of the tensions experienced in the industry when recruiting and hiring candidates who are both qualified and capable of integrating quickly into the CSIRT. Furthermore, the expertise identified by experts helps highlight a broad range of needs for analysts, especially as technology develops to augment and assist them in their tasks. These results are applied accordingly in Study 3 to develop concepts for future automation that can address expertise requirements.

CHAPTER 6. NEEDS ANALYSIS AND REQUIREMENTS DEVELOPMENT

Chapter 6 describes the last of three studies presented in the research conceptual framework. The data collection method was a Needs Analysis; data from previous studies and available market data were used to determine gaps between CSIRT needs for information sharing during incident handling and current capabilities of automation for these tasks. Data collection and analysis details are described in Section 6.1. Results from Study 3 answered two (2) of the research questions, which further described gaps and developed a path forward for addressing information sharing needs in future automation development in incident response. Findings and answers to these research questions are presented in Section 6.2.

6.1. Data Collection and Analysis

6.1.1. Overview of Needs Analysis Methodology

In Study 3, a Needs Analysis was conducted; this included an assessment of automation gaps with respect to industry needs. This Needs Analysis methodology applied the six dimensions of expertise construct as a framework to identify certain needed capabilities on behalf of human analysts in the system. Overall, the Needs Analysis in Study 3 provided justification for pursuing new automation development through three phases: Operations Analysis, concept of operations (CONOPS), and Functional Analysis. These phases help to collect data and identify needs, define new concepts for the next generation of technology, and produce some ideas of automation functions, respectively. Results of the Needs Analysis informed an operational concept for one function using systems architecting, as presented in Chapter 3. An overview of steps followed in Study 3 is located in Figure 17 below.

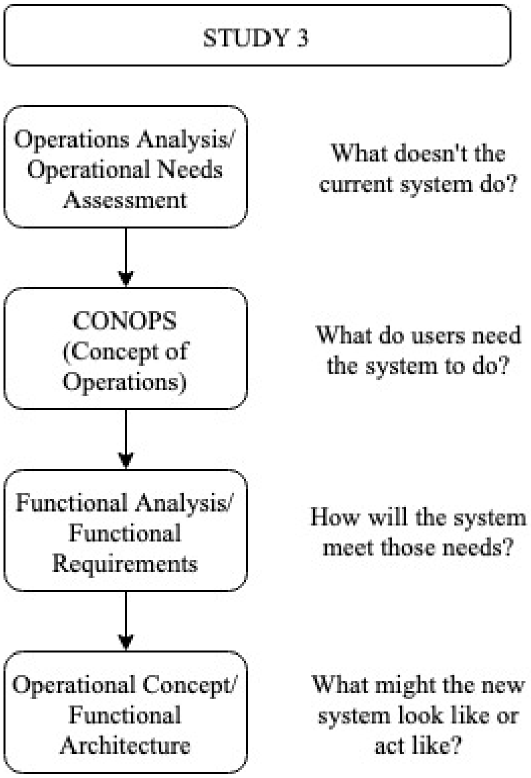
STUDY 3
Operations Analysis/ Operational Needs Assessment
What doesn't the current system do?
CONOPS (Concept of Operations)
What do users need the system to do?
Functional Analysis/ Functional Requirements
How will the system meet those needs?
Operational Concept/ Functional Architecture
What might the new system look like or act like?

Figure.17. Research Steps for Study 3

6.1.2. Data Collection

As described in Chapter 3, the Needs Analysis started with an operational needs assessment of current automation technology in incident response: Security Orchestration, Automation and Response (SOAR). SOAR technologies are platforms that mainly integrate various tools across incident response, conduct monitoring and automated response via pre-programmed rules, and provide guidance in the form of playbooks for analysts to follow during investigation. Data collected about these SOAR platforms acted as a starting point for identifying gaps between new technology and expertise needs of analysts, and informed the CONOPS and subsequent steps.

SOAR platforms included in this study were selected using a combination of two different techniques. First, a well-cited and influential report on SOAR technologies (Neiva et al., 2017) was used to identify some of the platforms included in the analysis. The report highlighted 16 different SOAR vendors including in the in-depth analysis of SOAR capabilities. Second, a generic Internet search was conducted for “SOAR, technology, cyber security” to identify other prominent tools that might not have existed at the time of the report, or were not included in Gartner's analysis. From these two techniques, nine (9) platforms were chosen (out of 19 total identified) for analysis based on the amount of feature information available on their respective websites; this was essentially a convenience population sampling.

- Cybersponse (https://cybersponse.com/)
- Demisto (https://www.demisto.com/)
- Siemplify (https://siemplify.co/security-platform-overview/)
- Swimlane (https://swimlane.com/platform/)
- Phantom (https://www.splunk.com/en_us/software/splunk-security-orchestration-and-automation.html)
- D3 Soar (https://d3security.com/platform/)
- LogRhythm (https://logrhythm.com/solutions/security/security-automation-and-orchestration/)
- Syncurity (https://www.syncurity.net/)
- Resilient (https://www.ibm.com/security/intelligent-orchestration/resilient)

In order to gain information about each platform, each platform's website was evaluated as the main source of data. If available, technical reports, white papers, and sales information were also included from the respective platform website. Collection included a line-by-line capability and feature assessment of each platform.

6.1.3. Data Analysis

The Needs Analysis provided the foundation for investigation and further pursuit and development of new systems from a research perspective. Kossiakoff & Sweet (2003, p. 57) state that the main outputs of Needs Analysis are the answers two major questions: "is there a valid need for a new system?", and "is there a practical approach to satisfying such need?" The following analysis for steps help answer these questions in terms of identifying operational deficiencies and technological opportunities to address gaps in CSIR. This section is organized to first describe the procedure of the operations analysis, the results of which are summarized in a CONOPS document (Appendix P). Next, the procedure for the functional analysis is described, followed by an outline of what is included in the concluding feasibility statement.

6.1.3.1. Operational Needs Assessment Procedure

The goal of this analysis (also called an Operations Analysis) was to identify operational gaps with respect to dimensions of expertise represented in SOAR platforms. Essentially, this study identified trends in current technology regarding which dimensions they tend to address or not address. By comparing unaddressed dimensions of expertise to needs identified in Studies 1 and 2, the research outlines operational deficiencies in current SOAR platforms.

Existing examples of gap analysis were used to develop the format for the Operations Analysis (Mineraud, Mazhelis, Su, & Tarkoma, 2016), especially those that were qualitative in nature (G. Reed, Philip, Barchowsky, Lippert, & Sparacino, 2010). This method allowed me to follow more formalized steps and meet existing standards that could not be identified in literature referring to 'gap analysis'. Systems engineering resources also helped guide analysis outputs (*ISO/IEC/IEEE 29148:2011*, 2011; *ISO/IEC/IEEE Std 31320-1:2012*, 2012; Kossiakoff & Sweet, 2003).

From the data collected about the SOAR platforms, each feature was evaluated against the six dimensions of expertise (Garrett et al., 2009). This was done to identify which dimensions the feature could augment for a human user. See Appendix K for additional definitions of the dimensions used to classify each feature. A matrix was developed to map features to dimensions; an example from one platform is shown in Table 11. The full table is located in Appendix O.

Expertise columns were ordered using empirical results from Study 2. Market reports, expert assessments, and academic literature (Table 3) were also included to support expertise needs, but were not quantitatively summarized for direct comparison. The format of the matrix allowed for fast visual assessment of gaps, much like a heat map, to provide preliminary results of where capabilities did or did not address needs.

Table.11. Example of Technology Capability Matrix

SOAR Platform / Technology	**Features**	**Situational Context**	**Subject Matter**	**Policy**	**Expert ID**	**Interface/Tool**	**Communication**	**Self-Awareness**	**Info Flow Path**
Demisto	Determine paths or flows				x				x
	Real-time workplan review	x							x
	Codeless playbook creation		x			x			
	Incident repository / knowledge database	x		x	x				
	Evidence board for information presentation during investigation	x							
	Multi-tenancy (data segregation by role)	x		x			x		
	Unified platform - integrated technologies	x				x			
	Incident or user-based reporting / Analyst tracking							x	
	Auto-documentation of incident activities			x					
	Virtual "war room" / ChatOps	x	x		x	x			x
	Correlations & Related Incidents	x							
	Machine Learning Chatbot				x		x	x	

Next, instances in each dimension were summed across all SOAR platforms to understand total platform capabilities compared to the six dimensions of expertise. This allowed each dimension to be addressed accordingly, discussing the capabilities advertised in each dimension, and how each dimension is represented quantitatively. From this point, findings could be organized to answer two questions: "*how does the current system not meet projected needs?*", and "*what is the value of fulfilling these needs?*". Expertise literature regarding software development relates these activities to requirement analysis (Sonnentag, Niessen, & Volmer, 2006).

6.1.3.2. *Defining the Concept of Operations*

A concept of operations (CONOPS) is a document created to highlight user needs and provides a voice for key stakeholders. It describes aspects of the user's environment, including organization, mission, and objectives (*ISO/IEC/IEEE 29148:2011*, 2011). A CONOPS is commonly used in systems engineering projects to set the stage for how a new system can best meet the needs of a user, and provides some high level detail regarding characteristics of the system. Essentially, this document was developed to help summarize and synthesize findings, acting as the main input for the functional analysis. The CONOPS is meant to be a high-level description of needs that can be addressed in future SOAR development, and does not yet address specific aspects or features. Literature differentiates between the CONOPS and the operational definition, which is a subcomponent of CONOPS that is developed later in the systems engineering process (*AIAA Guide to the Preparation of Operational Concept Documents (ANSI/AIAA G-043A-2012)*, 2012)

A well-used standard on developing CONOPS was reviewed to better understand methodology and deliverables (*ISO/IEC/IEEE 29148:2011*, 2011). According to this standard, a CONOPS includes several components. First, the document provides a clear description of the existing system, including where it meets and does not meet the needs of users. In identifying gaps, it provides the justification for a new system, which is further elaborated with some level of discussion about ideas for the feasibility and lifecycle of the new system. Next, the CONOPS provides a description of the proposed system, defining operational objectives and how the new system directly meets user needs. This should include how the new system can or should overcome changes in environment. Lastly, the CONOPS defines scenarios that highlight use of the system with in the user's environment. The procedure generally followed the IEEE standard (*ISO/IEC/IEEE 29148:2011*, 2011). Parts of the CONOPS overlap with what AIAA defines as the operational concept. Thus, the CONOPS document (Appendix P) briefly touches on the concept and operational requirements regarding how the future state of the system can help meet the user needs defined in the CONOPS. These components are then further elaborated in the Functional Analysis and Operational Concept sections.

6.1.3.3. Conducting Functional Analysis

The Functional Analysis provides the functional definition of the proposed system. First, operational requirements were translated from the CONOPS into high-level functions that corresponded with each requirement. For instance, the need for explainability to the analyst results in two basic functional requirements: the system should a) have additional background logic beyond how an expert programs the rules, and b) be able to trace and explain that logic to a novice analyst. Each function was supported with examples from literature that describe instances or outlines of similar functionality to support feasibility and definition. Some high-level functional requirements were defined by using a general systems thinking approach, breaking down each function into inputs, process, outputs, and feedback loops. Finally, a tentative allocation of functions was developed between subsystems by defining functional interactions and organizing them accordingly. A general depiction of this process is located in Figure 18 below.

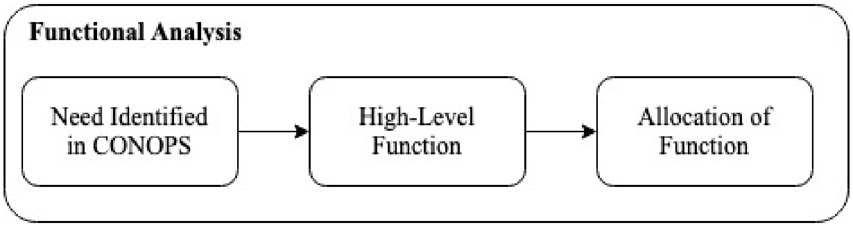

Figure.18. Functional Analysis Steps

6.1.3.4. A Statement On Feasibility

One important component of new system development is determining if the proposed system can feasibly be developed with an acceptable amount of cost and risk (Blanchard & Fabrycky, 2006). Literature proposes metrics to help guide measures of effectiveness for the new system in efforts to identify quantified needs early in the design process (Blanchard & Fabrycky, 2006), but also to create better understanding of capability and feasibility. Kossiakoff & Sweet (2003) discuss the Feasibility Statement as a key component of Needs Analysis to help further conceptualize the physical design of the system, otherwise known as the “form”. In short, functional design on its own cannot help define feasibility.

The goal of this study was not to provide a fully validated function and form of a new system for SOAR 2.0, but rather help conceptualize new capabilities for future development. As form-related questions were not part of the original scope or research questions, the focus for this Needs Analysis did not include physical aspects of the proposed system. However, the first two components of the Needs Analysis from (Kossiakoff & Sweet, 2003) provide a solid foundation further pursuit of new systems from a research perspective.

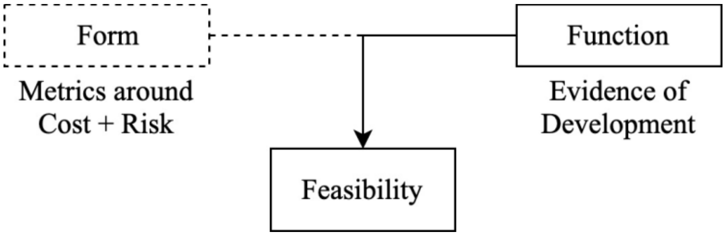

Figure.19. Feasibility Definition

Figure 19 summarizes this discussion of feasibility, and the distinction of form vs. function. Reiterating that this dissertation focuses on function, a narrowed definition of feasibility was adopted that more closely aligns with Technology Readiness Level (NASA, 2017). Here, feasibility is defined as *the extent to which literature provides evidence that the capability exists or is currently in development, such that the likelihood of realization in industry within the next 5 to 10 years is high.* That is, research and development literature provide evidence that other entities are currently working on similar capabilities, indicating that the functions are indeed feasible within the cyber security technological domain. The current estimated level of development of these capabilities is at or above a Technology Readiness Level 4.

6.1.3.5. Defining the Operational Concept And Functional Architecture

The final output of Study 3 helps answer the research question about a feasible path forward. Two products were created in succession to address this: the operational concept and the functional architecture. This section provides additional information to build upon that in Chapter 3 regarding how these activities were conducted.

As described, the CONOPS provides a larger overview of the needs and potential performance of a new system. In order to define additional details of certain functional components of said system, the operational concept is created. After determining operational needs from user data and translating them into functional requirements in the CONOPS, the operational concept explored technological capabilities to further conceptual development and establish proof of feasibility (Figure 20). Levis & Wagenhals (2000, p. 228) define the operational concept as “a concise statement that describes how the goal [of the system] will be met”, and state that the operational concept is a precursor to developing functional and physical architectures. Alternatively, another source form systems engineering literature (*AIAA Guide to the Preparation of Operational Concept Documents (ANSI/AIAA G-043A-2012)*, 2012) describes a more in-depth document with respect to description, and less emphasis on form-focused outputs, as the operational concept is meant to be a ‘living document’ that is continuously revised as more information is gathered from stakeholders and various analyses. Though literature diverges on the depth of information needed in an operational concept document, this study aimed to clearly define one of the concepts proposed in the Functional Analysis document. The operational concept document (Appendix R) expands upon the concept of “facilitating collaboration” (one function identified during the Functional Analysis) for SOAR technologies with respect to interfacing with and assisting human analysts. Following the guidelines in the AIAA guide, further detail was provided regarding what the functions would do and a potential functional architecture for the concept.

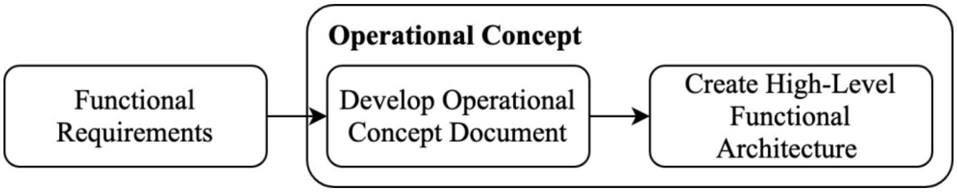
Operational Concept
Functional Requirements
Develop Operational Concept Document
Create High-Level Functional Architecture

Figure.20. Operational Concept Steps

Included with the operational concept is the functional architecture, which helps define different aspects of the concept in terms of activity, process, data, and rules. Developing the functional architecture components largely followed the procedure in (Levis & Wagenhals, 2000), drawing upon other systems engineering literature to help provide additional clarity and perspective (Blanchard & Fabrycky, 2006; Wagenhals, Shin, Kim, & Levis, 2000). Outputs from all previous steps in Study 3 informed this product, which again focuses on one function identified during Functional Analysis.

6.2. Results

6.2.1. Current State Of SOAR

The current state of the system (in terms of SOAR capabilities) is broken down into two subsections: alignment with the six dimensions of expertise, and gaps in capabilities versus the six dimensions. These effectively answer RQ2.3, and help determine the answer to the underlying question of the Operations Analysis regarding the need for a new system. This section ends with a concept of operations for the next generation of SOAR technologies to help frame results of the Functional Analysis.

The resulting matrix broke down each platform into capabilities and features, and then assessed each feature against the criteria for the dimensions of expertise. This matrix can be found in Appendix O. The quantitative view of total counts for each dimension, in order of rank by the Study 2 experts, is shown in Figure 21.

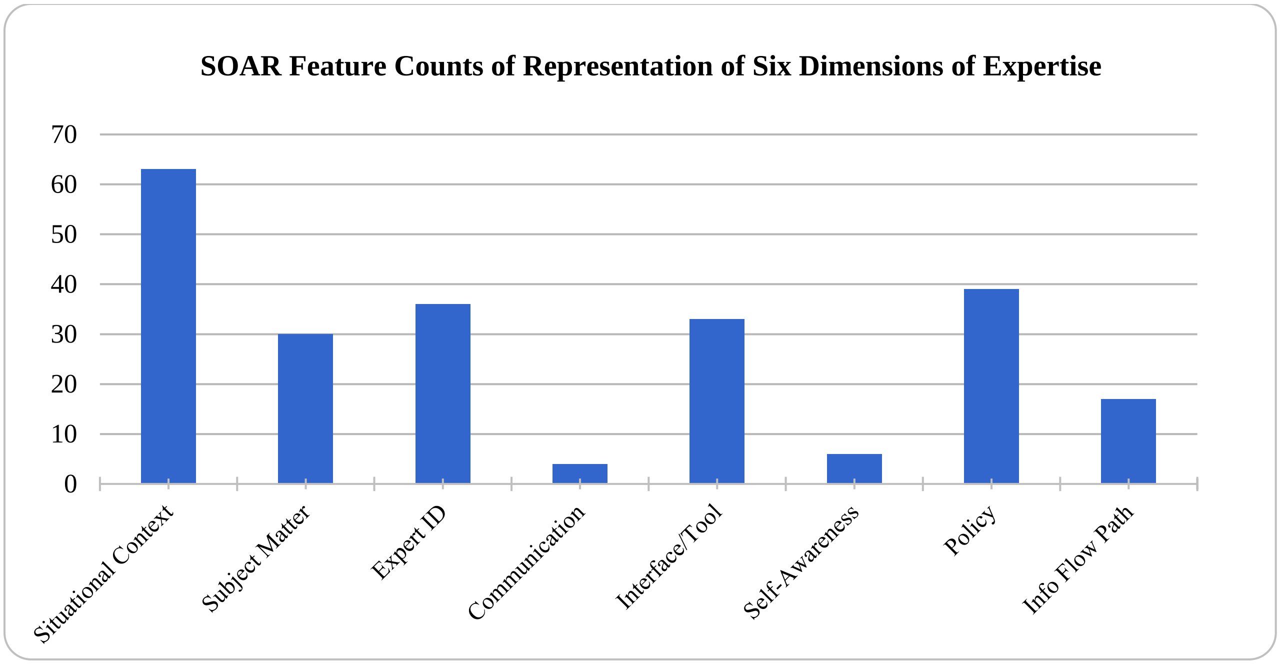

Figure.21. Sum of Features in Each Dimension of Expertise

6.2.1.1. *SOAR Features Aligning With Needed Expertise In Escalations*

The strongest alignment between SOAR features and the six dimensions of expertise is regarding situational context expertise. Considering how analysts need to pivot continuously between screens and platforms in order to gain context about an incident, developing technology to address these inefficiencies feasibly takes high priority. Many SOAR platforms focus on bringing to the context to the analyst by fetching data from different appliances and displaying them to the user.

Furthermore, the direction of the field seems to be trying to achieve a ‘single pane of glass’ (Neiva et al., 2017; Oltsik, 2018b), or a screen / interface window in which an analyst can obtain all needed context. The goal of many features relating to this dimension is efficiency, and in relation to escalations, providing situational context might help reduce the time to decide that an escalation is needed. Situation awareness literature (Endsley, 2018) specifically discusses how systems can support multiple levels of SA in relation to expertise, which can act as guidelines for system development and design. Additional cyber situation awareness (CSA) literature and ongoing research (Albanese et al., 2014; Mancuso et al., 2012, 2016; Tyworth, Giacobe, & Mancuso, 2012) also may help guide application design specifically for cyber security.

The next two dimensions of expertise most commonly mentioned by experts in Study 2 regarding escalations were subject matter expertise and expert identification expertise. This finding does not align with the next highest dimension that SOAR features address. The analysis indicates that *interface or tool expertise* is more commonly addressed than subject matter expertise.

Many SOAR platforms advertise that subject matter expertise can be embedded into the rules of SOAR protocols (C. Brooks, 2018; Oltsik, 2018a), essentially modeling automation and orchestration decisions off of experts. Expertise can come from the SOAR developer itself, or from the purchasing firm. The trend of pulling expertise from experts and embedding it into a system is not unlike the development of ‘expert systems’ (Buchanan, Davis, Smith, & Feigenbaum, 2018). Some of the lessons learned from the evolution of expert systems in other fields may be useful in helping guide SOAR development with respect to this dimension of expertise.

Augmenting expert identification expertise includes helping analysts determine where information or knowledge might exist, whether it is a person or some other non-human source. Some platforms, in bringing the context to the user, actually bypass this dimension of expertise altogether. Others claim to provide recommendations about who might be able to help with a particular ticket. One trend observed was the feature of playbooks or runbooks (Bedell, 2019; Neiva et al., 2017; Oltsik, 2018a). These effectively create predefined paths for incidents based on different indicators, which may also include to whom the incident should be escalated.

The last dimension of expertise that marginally aligned between salience in SOAR features and expert opinion was interface and tool expertise. SOAR platforms aim to overcome the learning curve of individual tools to allow the analyst to work seamlessly between appliances with the goal of reducing overall time to respond. This effort is both practical and needed, especially considering the complexity of the environment and software.

Information flow path was not strongly represented in expert opinions from Study 2. However, there is evidence that SOAR platforms are augmenting this dimension of expertise similarly to expert identification. Playbooks are predefined, rule-based procedures of what to do for a given incident, and help determine workflows and standardize system information flows. Essentially, SOAR platforms are standardizing and automating information flow paths, such that this dimension of expertise could be fully augmented by new technology.

6.2.1.2. Operational Deficiencies in SOAR: RQ 2.3

The Gartner report (Neiva et al., 2017) highlights certain technological capabilities that SOAR platforms should include in order to meet industry needs. More specifically, the report summarizes and discusses minimum requirements for orchestration, automation, and subsequent capabilities. Findings in Study 3 provide evidence of alignment between the capabilities identified by platform and this report, including playbooks / automated workflows, case management, abstraction layers, and documentation. One goal of this dissertation was to identify further gaps and areas for development to supplement and expand existing reports on SOAR capabilities. More specifically, RQ2.3 aimed to identify capability gaps that currently exist in automation technologies with respect to user-identified needs in dimensions of expertise.

One dimension of expertise that was identified as important for novice analysts was *communication expertise,* also identified as *interactional expertise* within expertise literature (Collins & Evans, 2002, 2018). Many SOAR platforms advertise that they help overcome communication barriers by providing chat features within the tool, as well as the ability to share other artifacts and documents. This aims to not only facilitate collaboration, but also to document the process for auditable record. However communication expertise extends beyond the mode of communication (which would, in fact, be information flow path expertise) or the auditability of collaboration. Instead, it builds on the idea that everyday communication skills (i.e. knowing how to talk to people and interact with them in different situations) are critical to analyst success. Very little evidence of true augmentation was founded regarding communication expertise in SOAR platforms.

In addition to the dimensional expertise gaps, *self-awareness* was a concept identified by experts as important, but was relatively unaddressed by SOAR platforms. Self-regulation and reflectiveness are both important aspects of learning (Breed, 2003; Ertmer & Newby, 1996). Self-awareness may be an important underlying aspect of *building* expertise and progressing with personal development, particularly in helping an individual know his or her boundaries, observing one's own performance, and making adjustments as they learn. Within ACTA (Militello et al., 1997), the prompts identify aspects of self-awareness in relation to expertise (M. S. Cohen, Freeman, & Wolf, 1996; Glaser & Chi, 1988; Klein & Hoffman, 1993). Furthermore, findings from Study 1 support the idea that feedback acts as a performance signal, which can trigger opportunities to self-reflect.

As mentioned, the goal of current SOAR platforms is mainly to alleviate labor shortages and provide some level of protection to companies inundated with data, false alarms, and a complicated array of software. Clearly this goal is focused on current operational stability. SOAR platforms offer some level of solution to immediate problems relating to low-level response, consistency, and tool integration. However, the next steps of the field should progress towards long-term development and retention of cyber security professionals to build sustainable capacity. By automating T1 activities, companies have reduced potential training opportunities for new analysts such that the traditional path to becoming a cyber security expert may fundamentally change. Thus, the findings indicate that *there is no explicit need for an entirely new platform, but expanding design considerations to build expertise, in addition to augmenting it, is warranted.*

6.2.2. A Feasible Path Forward: RQ 4.0

One major goal of this dissertation was to apply the conceptual framework and chosen methods in a replicable, comprehensive, and productive methodology. RQ4.0 was intended to deliver useful insights to CSIRT managers for direct use, and to identify a feasible technological path forward by translating the collected insights and functional requirements into actionable solutions for development. Insights for managers are provided in the amalgamated answers to previous research questions; some specific examples of actionable items are presented in Section 7.3.3.1. The following subsections describe the technology-focused outputs of Study 3 as products that help define this path forward for SOAR development.

6.2.2.1. Concept of Operations for SOAR 2.0

The CONOPS document (Appendix P) describes how a new system, or at least new system capability, is justified to support continued growth and advancement of analysts, as well as create better shared, distributed SA between human-machine teams. The document describes user needs, translates these needs into operational requirements to support current and future analyst activities. An operational requirement is a statement identifying essential capabilities (Kossiakoff & Sweet, 2003). The following figure (Figure 22) is also located in the CONOPS document. The operational requirements prompt the next steps of the Needs Analysis: the Functional Analysis, or development of the functional requirements of the system.

Current SOAR platforms do not advertise goals of development as synthetic teammates. However, some research has proposed how to pursue this path (Lathrop, 2017). The focus on the CONOPS is to continue driving development of SOAR into the domain of human-automation teaming. SOAR currently is able to automate tasks and guide analysts down pre-defined routes, but is not necessarily able to “share knowledge or expertise” with analysts in ways that analysts currently do so with each other. In order to make the jump to “teammate”, SOAR capabilities should expand to accommodate knowledge sharing, collaboration, and communication practices beyond a basic chat platform currently seen in SOAR technologies. The operational requirements below highlight potential capabilities that will elevate SOAR capabilities such that interactions between humans and the system are *bi-directionally value-added* and more realistically emulate current human organization and development.

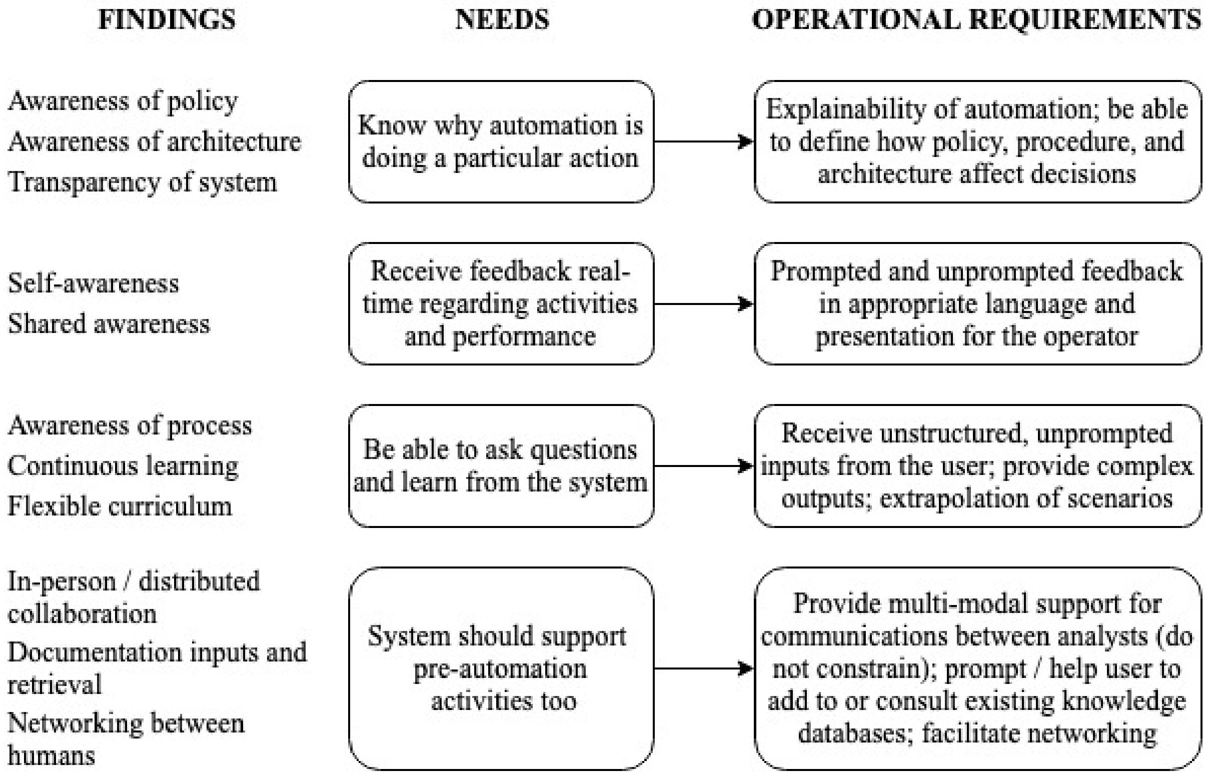
FINDINGS
NEEDS
OPERATIONAL REQUIREMENTS
Awareness of policy
Awareness of architecture
Transparency of system
Know why automation is doing a particular action
Explainability of automation; be able to define how policy, procedure, and architecture affect decisions
Self-awareness
Shared awareness
Receive feedback real-time regarding activities and performance
Prompted and unprompted feedback in appropriate language and presentation for the operator
Awareness of process
Continuous learning
Flexible curriculum
Be able to ask questions and learn from the system
Receive unstructured, unprompted inputs from the user; provide complex outputs; extrapolation of scenarios
In-person / distributed collaboration
Documentation inputs and retrieval
Networking between humans
System should support pre-automation activities too
Provide multi-modal support for communications between analysts (do not constrain); prompt / help user to add to or consult existing knowledge databases; facilitate networking

Figure.22. CONOPS Content: Needs to Requirements

6.2.2.2. Functional Needs of the System

Reiterating the goals of the Functional Analysis, this step aimed to provide functional definition to a new system by translating operational objectives into functions that must be performed in order to meet user needs. Three main classes of functions resulted from this analysis. The full analysis document can be found in Appendix Q.

The first set of functions aims to ensure ***explainability and transparency*** of the system, which are essentially measures of effectiveness for the system. The system must be able to show on-demand the activities it has been performing behind the scenes, as well as a summary of decisions made, such that the user understands what the system is doing *and why*. This is also referred to in literature as explainable artificial intelligence (XAI). System-prompted decisions should be able to be aggregated and analyzed by higher tier responders and analysts who can evaluate the enacted policies and rules over larger data sets and across new potential scenarios, effectively validating system decision-making. In future iterations of this functional concept, the system should be able to aggregate with some level of pattern recognition as well, creating some ‘self-awareness’ of its own efficiency and correctness. Finally, the system must be able to provide explainability to the user regarding playbooks, policies, and activities that it prompts the analyst to follow during incident response. The system should be capable of answering questions regarding why certain steps are taken, when they are or are not appropriate, and extrapolation into different or future scenarios. A summary of specific functions in this class can be found in Figure 23 below.

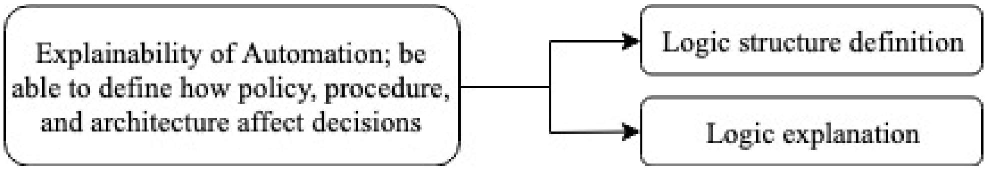
Explainability of Automation; be able to define how policy, procedure, and architecture affect decisions
Logic structure definition
Logic explanation

Figure.23. Explainability Functions

The second set of functions aims to create ***bi-directional, value-added human-machine interaction*** between the human analyst and the system (Figure 24). The system must be able to engage in collaboration *with* the analyst user by providing interactive feedback of incident activity and performance, collecting and analyzing competency progress of analysts, and performing after action reviews with analysts. This includes some level of sophistication regarding parsing inputs from the human, formulating a response, and delivering the response in an effective and appropriate manner, all while remaining dynamic over time. Generating this level of explanatory feedback is not a trivial function, and requires some level of understanding of the mental models of the human regarding the problem at hand, as well as determining what level of explanation is needed (Hoffman, Klein, & Mueller, 2018). The ability to provide abstracted explanations through analogies or scenario extrapolations would display definitive educational opportunities while also advancing the current state of XAI. Additional functions to enhance sensing of human comprehension and feedback might also provide additional capabilities in measuring effectiveness of system outputs.

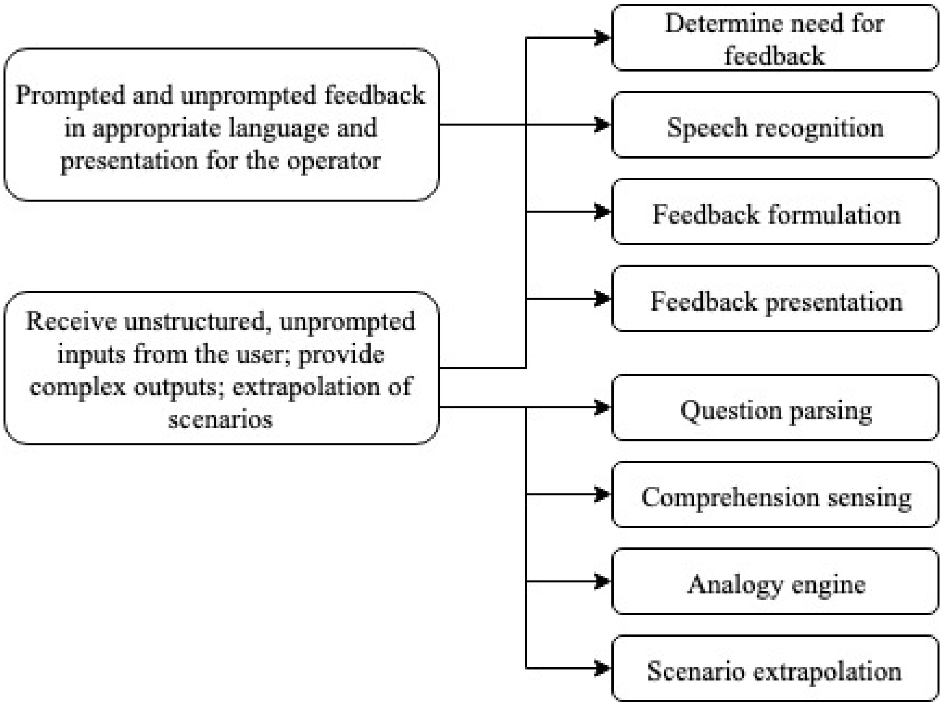
Prompted and unprompted feedback in appropriate language and presentation for the operator
Receive unstructured, unprompted inputs from the user; provide complex outputs; extrapolation of scenarios
Determine need for feedback
Speech recognition
Feedback formulation
Feedback presentation
Question parsing
Comprehension sensing
Analogy engine
Scenario extrapolation

Figure.24. Value-Added Interaction Functions

Finally, the third set of functions aims to ***facilitate better collaboration and knowledge coordination*** within the larger organization through knowledge networking capabilities (Figure 25). Functions to support better human-to-human and human-to-system networking require deeper levels of definition of knowledge networks in a given organization and behaviors regarding knowledge sharing patterns. Moreover, determining deficits and connection strategies help with prediction and anticipation of user knowledge needs. An additional component of collaboration is system support of richer communication modes and methods between humans to facilitate shared awareness of collaborating analysts.

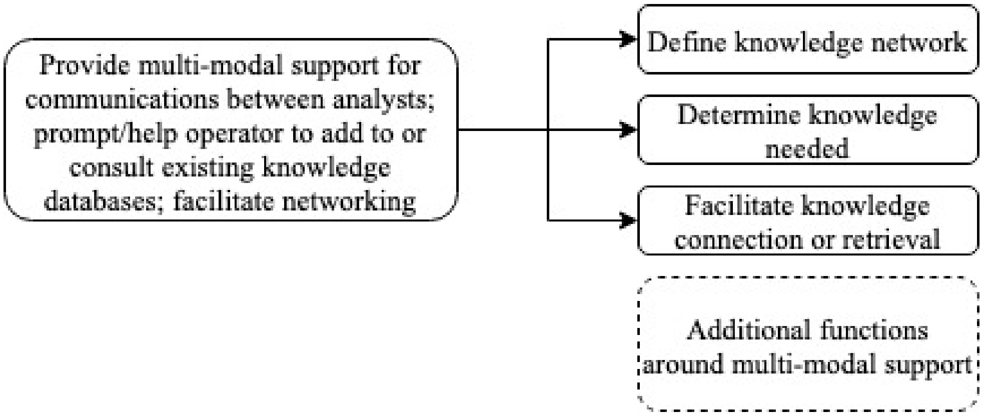
Provide multi-modal support for communications between analysts; prompt/help operator to add to or consult existing knowledge databases; facilitate networking
Define knowledge network
Determine knowledge needed
Facilitate knowledge connection or retrieval
Additional functions around multi-modal support

Figure.25. Collaboration Facilitation Functions

The Functional Analysis presents ideas for conceptual development regarding additional capabilities for SOAR platforms to better meet the needs of CSIRTs. However, in proposing these functions, it is also clear that at the current rate of development and adoption for machine learning and artificial intelligence, not all of the above functions can be developed in a short time frame and deployed to the field. Thus, there is a need to prioritize the functions by feasibility and scope, such that software developers have some options to work on while the state of the art advances to support the other functions. Feasibility is defined here as *the extent to which literature provides evidence that the capability exists or is currently in development, such that the likelihood of realization in industry within the next 5 to 10 years is high.* This definition roughly equates to a Technology Readiness Level (NASA, 2017) of 4 or higher. Based on evidence of current development, these capabilities may be able to be developed and realized within the cyber security industry in a conservative timeframe. The other functions have a much larger estimated scope for research and development, but are ripe opportunities for future activities.

The prioritized function classes from this analysis are *explainability and transparency* and *facilitating collaboration*. Considering the current DARPA focus on XAI, there is a wealth of ongoing research, methods, and tools available to developers to start creating this capability in SOAR. Available resources can also be used to incorporate best practices and cutting edge approaches in order to develop this set of functions. However, facilitating collaboration through knowledge networks will require some additional conceptual development in order to meet the unique needs of this environment.

6.2.2.3. Practical Approaches to Addressing Gaps

With functional needs identified, the next step in answering this research question is to address if there is a practical approach to satisfying those needs. The Functional Analysis narrows the scope of development ideas and prioritizes concepts based on literature support as evidence of feasibility. Additional research and details provided in the operational concept document (Appendix R) support development of a new set of functions for SOAR platforms to help facilitate knowledge sharing and collaboration between humans and between humans and technology. An example of developing one of the corresponding functions is pictured in Figure 26.

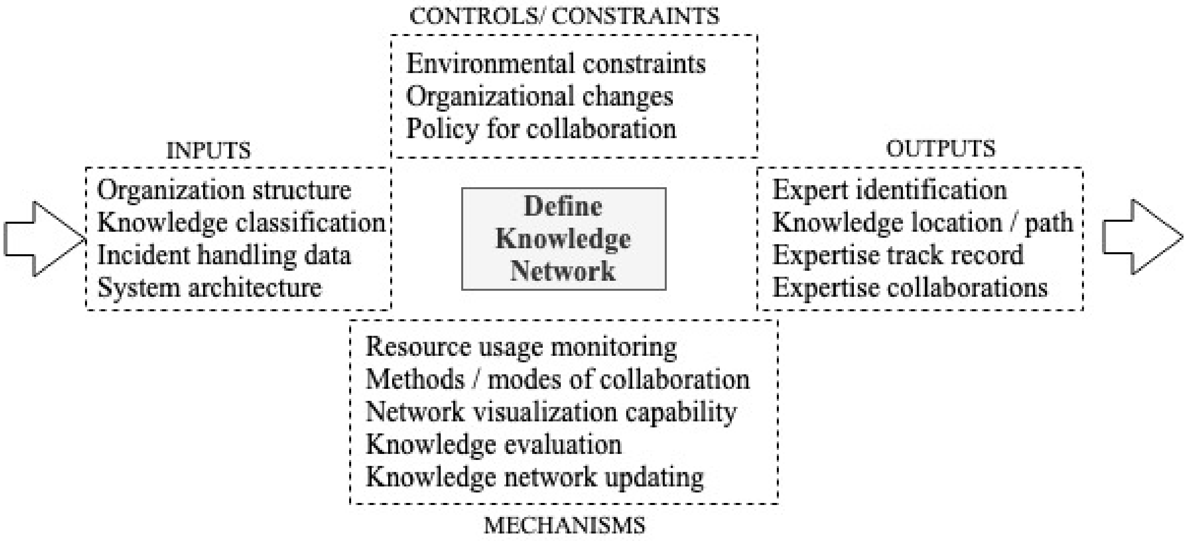

Figure.26. Example of Function to be Developed to Support Knowledge Sharing

The technological opportunities identified support human-automation teaming by increasing capability of automated assistants to build and maintain a schema of knowledge within an organization, as well as facilitate connections between entities to support knowledge sharing. Research supporting knowledge networks and needs of automated assistants suggest that the development of these capabilities in SOAR is indeed feasible.

Adding to the operational concept, a draft of a functional architecture was also developed to support how the new capabilities would work. The functional architecture is comprised of four (4) different models what describe different aspects of the proposed capability. Also included is an integrated data dictionary, which helps define terms and data types and bridge the various models. Materialization or vision (Kossiakoff & Sweet, 2003, p. 122) is supported by architecting the concept from a functional perspective. This allows readers and developers can construct similar mental models and expectations for the added functions.

An example of one of the developed models from the functional architecture is shown in Figure 27. This activity model provides an overview of conceptualized function, specifically regarding process flow and information flow. The activity model can be used to guide further conceptual development, and acts as a road map for other components of the functional architecture. The accompanying models are presented in Appendix R.

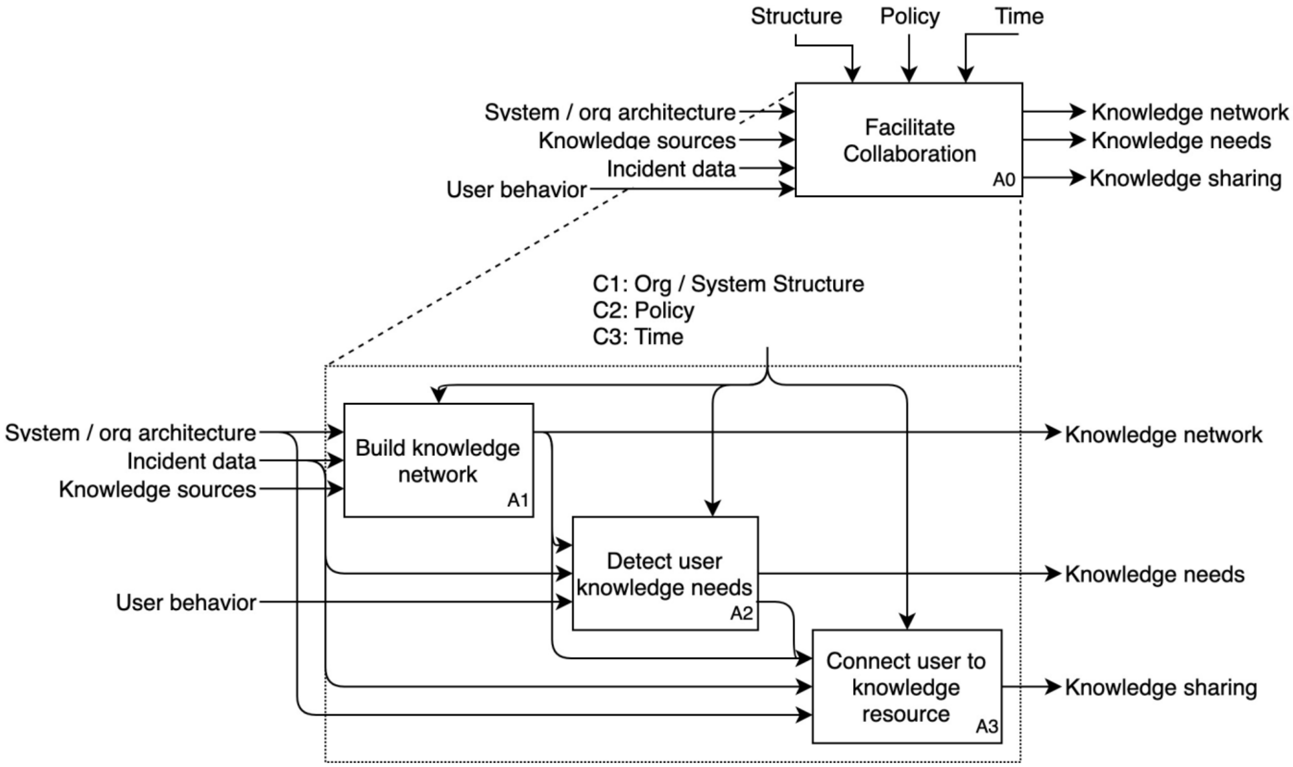

Figure.27. Activity Model of Functional Architecture

Using a human factors approach to collect data from various levels of analysts and experts in CSIR informed a deeper understanding of the larger socio-technical system. Furthermore, human factors methods help synthesize needs of users into ideas for new capabilities in relation to current technology. The conceptual development process (and corresponding appendices) from systems engineering helped translate needs into capabilities through various standardized analyses, and provides a solid foundation for future development of functions to support human-automation teaming in incident response.

CHAPTER 7. DISCUSSION

7.1. Organizational Aspects of CSIR

As the findings in Chapters 4 through 6 proceed through various levels of security organizations, the following sections are organized to transition from the base level of the security analyst, to the team level of CSIRTs (comprised of analysts), to the enterprise level in which CSIRTs exist. The topics within each subsection focus literature and findings to the specific topics explored in the research questions, including information sharing, expertise, and automation. As systems engineering provided much of the lens for analysis across the three studies, themes from various systems perspectives are also discussed in relation to each of these levels.

The base unit of a CSIRT is the security analyst, who conducts knowledge work in order to handle security incidents as quickly and effectively as possible to mitigate threats. Analysts are typically part of a larger team, or CSIRT, which functions as a system to monitor and respond to all incoming threats by coordinating with other analysts and applying expertise where needed. Finally, CSIRTs are often embedded in larger organizations, acting as an internal (or sometimes external) service provider to secure network operations. This hierarchical structure provides the organizational scheme for the following subsections.

7.1.1. The Individual Analyst

As presented in Chapter 1, one of the biggest challenges in CSIR in 2019 is the lack of qualified personnel to fill positions across the various tiers of a security team (HCL Technologies, 2019). The labor shortage has only worsened since 2016 ((ISC)^2, 2018; Bureau of Labor Statistics, 2016), with some reports estimating a threefold increase in demand through 2021 (Morgan, 2017). Literature describes several reasons behind *why* there is a shortage, which include poor education and development pipelines (Assante & Tobey, 2011), dynamically evolving skill requirements (Hoffman et al., 2012), and finding *and keeping* candidates with a balance of personal, technical, and business skills (Cobb, 2016; HCL Technologies, 2019). This dissertation explores the problem through the lens of expertise needed to perform analyst jobs and the automation being developed to augment human analysts.

Research questions pertaining to the analyst level of the hierarchy focused strongly on information sharing processes, pain points from the triage and mitigation level (T1 – T2) perspectives, and expertise requirements to perform those activities. The research findings indicate two major topics worth discussing at the analyst level that relate to the research questions and literature. The first area is about ***developing expertise*** in lower tier analysts, and the second area is about ***two-way exchanges in information sharing*** tasks.

7.1.1.1. Developing Expertise in Analysts

Expertise is a central part of performing security analysis. In particular, researchers and industry experts have developed resources that provide a generic list of KSAs needed to be successful (Newhouse et al., 2017). While the list itself is a valuable reference, one goal of this research was to further explore expertise as a framework of different dimensions (Garrett et al., 2009) in order to better understand what kind of expertise was needed, and at which levels of response. Indeed, literature points out that maintaining relevant expertise, and developing new areas of expertise in the field are both critical to the success of an analyst (Hoffman et al., 2012; Ruefle et al., 2014). As the field expands, the needed areas of expertise grow, and analysts must constantly stay abreast of new trends and developments in order to protect their networks and remain competitive in the security marketplace (Vieane et al., 2016). Despite this fact, the data suggest that this may: 1) depend on the organizational philosophy, and 2) only be expected of higher tier analysts who specialize in particular areas of response. Several teams observed did the opposite of exposing lower tier analysts to new threats or skills by restricting the response space for T1 analysts. Procedures dictated incident types appropriate for T1, and prompted automatic escalation from T1 to T2 for other incident types. One reason given was that these analysts may not be full-time employees, or may not be employees at all (but rather contractors), which poses some level of risk to errors and liability. This scenario was one of the tensions uncovered that was related to growth and development of security analysts.

A recent survey presented similar findings, stating that lack of development is one of the top reasons why analysts leave their roles at a given firm (HCL Technologies, 2019). Through the exploratory study presented in Chapter 4, findings identified one particularly interesting difference regarding growth and development that impacted individual success. *Designing the T1 role to be a development position* allowed one company to invest in the development of each analyst, allowing that person to grow technically and professionally during her time in the SOC. Furthermore, the analyst was not expected to stay within the SOC when her rotation as an analyst concluded. In fact, the person could choose to transfer to a completely different department outside of security or IT to pursue a different path. This investment and added flexibility made the overall program very successful by designing the T1 position to be very educational while also allowing for upward or lateral movement. While this particular program was in a near-constant state of hiring, they also had the high coverage over three shifts (the only team observed to do so) and employed recruiting strategies inside and outside the company to fill the pipeline. The hiring manager mainly focused on character and problem-solving attributes, with less emphasis on certificates or formal degrees in IT. Furthermore, ‘graduates’ of the rotation could apply security knowledge and skills to other areas of the company and spread awareness of security issues to their new departments. Essentially, the employee development approach (Jacobs & Washington, 2003) from this team indicates promise for mitigating retention issues.

This research also found that at the T1 level in the three different settings observed, analysts are often expected to do repetitive tasks in finite areas of response. These tasks often included monitoring particular channels for particular types of threats or alerts, to which they would respond with relatively simple procedures for mitigation and resolution. In literature, these types of activities are connected with fatigue and burnout (Bourget, 2017), and a recent survey of cyber analysts identified these activities as less desirable compared to deeper types of analysis and problem solving ((ISC)^2, 2018). In conjunction with limited development opportunities, fatigue and burnout are common causes for losing qualified candidates at this level. In contrast, the team with T1 development roles expected analysts to do more than described above. These analysts rotated through responsibilities to get exposure to different types of threats. (Another team had rotational duties at the T2 level, but the rotational design was often inhibited by personal preferences and lack of discipline in employees adhering to assigned rotation schedules.) They were given more autonomy than T1 analysts in the other teams, and expected to be active participants during meetings and handoffs by reporting out on their individual investigations.

Research has identified professional development as an important component of building the cyber security workforce (Assante & Tobey, 2011; Burley & Bishop, 2011) Burley and Bishop's report highlights the need for this development in terms of 'roadmaps' for education and recruiting. However, the data indicate that the need for development persists, even after candidates are hired. A good example of observed post-hire professional development was the rotational program and managerial strategy in one team. The combination of both program and strategy allowed them to be successful in recruiting and retaining candidates while also building capacity in CSIR. By embracing the 'entry level' nature of the T1 roles and providing development and growth opportunities, the organization largely considered their program successful and without the same tensions observed in other teams regarding hiring. In summary, the research findings suggest that there are additional and alternative strategies for building capacity in security that include building talent pipelines and investing in experience-based development of analyst expertise, especially at the T1 level.

7.1.1.2. The Role of Information Sharing in Performance and Development

Information sharing is important in several contexts in incident response, such as between individuals, between teams, and between separate firms or organizations (Tetrick et al., 2016). It allows various entities to have awareness of the environment and/or activity at hand, and supports coordination between entities as they work together to handle an incident. In particular, information sharing between individual analysts is critical to conducting effective incident response, such as in the case of an escalation handoff (which transfers ownership from a lower tier analyst to a higher tier analyst). Data from Study 1 (Chapter 4) indicate that, during an escalation, a lower tier analyst has completed as much of an incident investigation as they can, and the incident requires additional expertise or action by someone at a higher tier level. Escalations typically happen within a ticketing system, but are often supported by additional communication between analysts (or analyst groups) via email, instant messaging, or phone calls. This added information flow is typically for the purpose of raising it to someone’s attention or confirming that an incident was received.

Despite the above description of how an escalation typically occurs, and the common belief amongst security analysts that escalation is a one-way transaction, the research findings indicate that the reciprocated response from the receiver is perhaps just as important as the escalating message itself. The feedback loop to the sender acts as a mechanism for awareness as well as learning. Feedback can be immediate, indicating that the incident was received and is being handled by a particular analyst. CSA literature includes feedback as a necessary component at multiple levels of situation awareness (Tadda & Salerno, 2010b) that allows an analyst to more accurately comprehend and project based on past and current data inputs. Extending this concept to shared situation awareness enforces the need for feedback through two-way information sharing to support team coordination (Franke & Brynielsson, 2014), especially due to the distributed nature of cyber settings (Tyworth, Giacobe, Mancuso, & Dancy, 2012).

Feedback can also happen after the receiving analyst reviews the investigative information from the sender. This type of feedback might be more directed at quality of the sender’s investigation, which can support learning and expertise development through socialization and exposure to expertise (Collins & Evans, 2018). Moreover, literature in cyber security has identified self-regulation and metacognition as key components of performance for cyber analysts (Cano et al., 2018; Jøsok, Lugo, Knox, Sütterlin, & Helkala, 2019). These factors also connect with communication effectiveness: a concept central to the six dimensions of expertise construct. Findings in Chapter 5 suggest that self-awareness is an important attribute for successful analysts to have, which directly relates to performance feedback. The research findings determined that self-awareness prompts analysts to seek out and parse the performance feedback such that they know where they can improve. Experts indicated that self-awareness is critical to success for security professionals, as it allows them to remain effective and competitive. Thus, if people with high self-awareness are being recruited for analyst positions, performance feedback is an important mechanism that could be incorporated into information system design and escalation protocols.

The analyst level of CSIRT operations currently dominates the concerns of the labor shortage problem. The dissertation findings support the idea that this level requires cultivation and support through better job design and more complete information flows. Moreover, these strategies will not become obsolete as the field evolves in terms of new threats and technologies. In conclusion, building capacity and creating resilience in cyber security requires attention and careful consideration for the design and operation at the heart of the skills gap: the human analyst.

7.1.2. The CSIR Team

Analysts act as the base unit of a CSIRT, which are often organized by tiers that reflect knowledge and expertise. T1 analysts typically have basic knowledge of networks, systems, hardware, and software that they apply during incident handling. As an analyst progresses to higher levels of response (T2 and higher), his expertise may broaden across a wide range of topics, or may specialize into one or two particular areas. Regardless of this progression, CSIRTs commonly include more than one analyst (Ruefle et al., 2014) such that the range and depth of expertise accommodate the threats experienced by their larger enterprise. From this perspective, ***collaboration*** and ***systems to support knowledge and information sharing*** between analysts on a CSIRT play an important role for fast and effective incident response.

7.1.2.1. Collaboration in CSIRTs

As previously discussed, the nature of CSIRT environments lends itself to organizational, physical, and even technological separation between analysts and other stakeholders working on the same incident, all of which can impact team performance. Incident response often involves multiple individuals with unique expertise and awareness working together in order to address a threat (Ahrend et al., 2016; Rajivan & Cooke, 2018; Tetrick et al., 2016; Werlinger et al., 2010). In this sense, collaborative problem solving is important for this process to occur smoothly. Within the CSIRT itself, collaboration can occur between analysts in the same tier or between analysts in different tiers.

Despite the need for collaboration, team-based literature has found that there are limited collaborative efforts made to bridge gaps between CSIRT members (Champion, Rajivan, Cooke, & Jariwala, 2012), with potential causes being team structure, communication, and information overload. The research findings presented in Chapter 5 expand this list to include social factors between team members; the broader cyber security domain has identified social factors as important for conducting security analysis altogether (Ahrend et al., 2016; Beznosov & Beznosova, 2007). Specifically, related findings in theme 3 of Chapter 5 state that *trust, prior experience, and professional relationship can affect how (multiple dimensions of) expertise can be developed and utilized*. In this sense, decisions regarding collaboration include underlying elements that stem from sociology and psychology. Experts in Study 2 referenced personal reputation as a factor they use to determine with whom they will work or to whom they might escalate a ticket during incident response. Work ethic and past performance, as well as personal trust in an individual, are conceivably important when making a decision that involves risk (Inaba & Takahashi, 2017; Mayer, Davis, Schoorman, Mayer, & Davis, 1995). Experts also indicated that they tended to go to the same people unless some sort of policy or rotational procedure was established to prevent favoritism. However, regarding future technological developments in CSIRTs, it is important to consider and explicitly acknowledge these factors when trying to encourage or design technology to support collaboration or supplement humans performing tasks.

Additionally, the data indicate that *physical distance* and *technological separation* may affect analysts' ability to perform collaborative tasks, such as the tools and environment at the disposal of the analysts. One struggle observed was that, multiple analysts may be involved in an incident, but may not have access to the same tools or information. This disconnect can create alignment issues during collaborative activities. Additionally, information pooling bias is common in teams with distributed information that try to pool unique facts from its members for problem solving (Rajivan & Cooke, 2018); such studies indicate that there are more complex effects of distributed information in collaborative problem solving beyond just physical and technological separation. Another struggle related to the physical layout and separation of team members was interruption in two-way communication flows. For instance, analysts sitting in the same room with no cubicle walls might have a conversation while looking at their own screen; while analysts in separate rooms or offices (or buildings) might reduce interaction to a single one-way message or email.

Related literature in information flows in mission control and online information coupling present some relevant findings to consider when mitigating these issues (Caldwell, 2011, 2015).

7.1.2.2. Systems to Support Knowledge and Information Sharing in CSIRTs

During incident response, teams must balance expertise across team members by seeking out and sharing relevant knowledge about threats and mitigation strategies (Ahrend et al., 2016; Mesmer-Magnus & DeChurch, 2009; Rajivan & Cooke, 2018). This communication can occur over a variety of channels, and may or may not be supported by information systems. From the cybernetics perspective presented in Chapters 3 and 4, viability of the overall organization is supported by a well functioning system dedicated to information sharing such that subsidiaries of the larger organization can efficiently and effectively coordinate activities (Tetrick et al., 2016). Channels should also support knowledge and information sharing in CSIR organizations, such that expertise across diverse areas can be easily and quickly accessed for use in a given incident. However, literature also notes that a technologically centered approach to understanding and addressing collaboration needs may not be effective (Ahrend et al., 2016). Systems designed for information sharing may be circumvented for a number of reasons, including perceived relevance of the information, convenience of using the system, and limitations of the system.

Findings from Study 1 (Chapter 4) indicate weaknesses in organizations around coordination activities, particularly with respect to information systems connecting teams or entities that need to share information during incident response. Separate ticketing systems, cluttered communication channels, and lack of shared network visibility within systems all contributed to poor coordination. Accordingly, teams with disjointed information systems expressed frustration and confusion when performing incident response that required collaboration with entities outside their immediate organizations, which is fairly common in incident response as a whole (Werlinger et al., 2010). This evidence validates the need for stronger information systems (‘System 2’ in VSM) to support effective operations in CSIRTs.

Information systems in System 2 should also facilitate knowledge sharing (Hoverstadt, 2010) through shared knowledge management systems. Study 1 produced empirical evidence of disconnected coordination efforts between cyber subsidiaries that validates the need for more focus on development and maintenance of information systems. These findings support previous literature on CSA regarding information sharing, especially across boundaries that separate different subgroups within the SOC (Tyworth, Giacobe, Mancuso, et al., 2012). Furthermore, the complexity of collaboration in incident response settings may require more from a tool or system than merely providing communication means (Rajivan & Cooke, 2018). Research cautions that it is critical to understand the context of use and user perceptions of IT systems and tools when investigating system improvements (Ahrend et al., 2016). How analysts use technology may vary depending on their task sets, specializations, and operational focus, and individual customization of tools is extremely common (Ahrend et al., 2016; Werlinger et al., 2010).

Some examples of well-functioning knowledge management systems were observed within the CSIRTs, though they were relatively insular within subsidiaries. Wikis and reference documents that were frequently updated and maintained by team members provided fast answers to incident-related questions in lieu of more interruptive communications between analysts. Factors such as common taxonomies and ontologies could also impact ‘findability’ and usability of knowledge management systems. While not specifically explored in the scope of this dissertation, these factors present potential subjects for investigation in future research about knowledge management in CSIRTs and design of augmenting technology. When not updated, utility of observed knowledge management systems was extremely limited to the point that analysts no longer considered them resources. Analyst perceptions of usefulness of both systems and information, have been identified as potential reasons why information is not shared and why systems are circumvented (Ahrend et al., 2016), further supporting the notion that user perception is an important factor to consider in information system design.

The team level of the hierarchy is much more complex than the individual analyst level when investigating retention and performance. Collaboration is essential in performing incident response, but has many underlying factors that support successful collaborative efforts, including those based in teamwork. Furthermore, collaboration in these distributed teams is facilitated by physical layout and technology (information systems) to coordinate tasks and encourage knowledge sharing. Without these social, physical, and technological pillars, teams may face challenges in effective and efficient collaboration during incident response.

7.1.3. The Enterprise and CSIR

CSIRTs provide security services, usually within (or in contract to) larger organizations (Chen et al., 2014). These teams operate within the bounds of the parent firm, following procedures and protocols of the respective organization. Research in the security domain (Beznosov & Beznosova, 2007) has identified organizational factors as having a direct impact on how a security team operates, and thus how individuals interact with each other. Research findings in Chapters 4 and 5 support this fact, indicating that ***mission*** and ***culture*** drive interactions at the team level.

Though limited, some security literature has identified the effects of enterprise-level factors on team structure and performance (Beznosov & Beznosova, 2007); some of these factors include organizational purpose (or mission) and hierarchical structure (Boudreau, Loch, Robey, & Straub, 1998). However, (Beznosov & Beznosova, 2007) also cite aspects of organizational culture (Coleman, 1990; Handy, 1995) as key factors in achieving competitive advantage in security. Findings of this research support this discussion by identifying organizational mission and culture as factors that affect collaboration practices at the team level. Themes 2 and 10 in Chapter 4 both implicate downward influence of managerial (or higher) policies and practices on communication and agency of CSIRTs. Furthermore, theme 2 in Chapter 5 states that *expertise sharing and development is enabled by collaboration practices and environment.*

Recalling the discussion about feedback in the 7.1.1, data suggest that performance feedback is a critical aspect of on-the-job training for analysts. However, findings in Study 1 and Study 2 indicated that performance feedback, as well as other types of feedback notifications, can be disrupted or inhibited by the cultural, physical, social, or technological environment. Technological impacts of the environment were discussed in the previous section. ***Cultural impacts*** may include managerial protocols mandating how performance feedback was given, or a general belief that ‘no news is good news’. For instance, observed negative performance feedback was typically funneled through management or team leads to avoid confrontational and potentially awkward one-on-one conversations between analysts. Additionally, it was unclear that T1 analysts received performance feedback of any kind in certain teams, especially if they are not considered full-time employees. (Note that two of the three teams hired contractors for T1).

The only team that openly discussed overall performance feedback was the team that designed their SOC around a development program in which performance reviews were a formal and continuous part of their processes. Within this team, informal positive and negative feedback was common, especially from the team leads and manager, regarding general performance. Yet even in this team, inter-tier performance feedback specific to incident handling was limited. Higher tier participants in this team acknowledged opportunities to spend more time mentoring T1 analysts, but seemed to reference good attitude and aptitude in the T1 analyst as prerequisites for time investment. As the sample size was relatively small, more investigation is warranted to explore phenomena in non-procedural inter-tier interaction in relation to teamwork, retention, and performance.

Another enterprise-level factor that had an impact on CSIRT operations was the ***mission*** and overall security posture of the firm. Without support from the highest level of management, security tasks may be viewed as inhibiting operations, and thus a nuisance to the rest of the organization. Furthermore, the organizational structure at the top (e.g. CIO, CISO) could impact this posture via conflict of interest. For instance, if the security officer reports through IT operations, there could be conflicting missions at the top of the organization between supporting security (mitigating threats) and supporting operations (maintaining uptime). A subordinate relationship of security reporting up to operations can create tension, as the CIO is ultimately responsible for operations before security. Without support in structure and policy, security operations may have reduced autonomy or independence to take action when security threats are detected. This was observed in the university setting and the state government setting, potentially due to the ‘conglomerate’ nature of their respective structures. Central security acted as a filter and notification leader in these settings, and often handed off incidents to affiliate IT groups for handling and mitigation. Some level of frustration was observed from analysts in these cases, as they had little authority to act on incidents, and even poor visibility regarding the real risks in the networks associated with them.

In summary, security operations can be highly impacted by enterprise mission and culture, and traditional organizational structures have been recognized as a significant obstacle in cyber defense (Staples & Sullivan, 2018). Referencing VSM, these 'System 5' level policies drive the purpose and operational procedures of the CSIRT. If the CSIRT exists in an enterprise that does not value security, the team may have limited agency to act upon threats. Eventually, this could lead to a drop in morale and overall effectiveness, as the team cannot execute upon the tasks they were trained to perform. The research findings link the role of upper management to CSIRT operations, suggesting that CSIRTs should be investigated within the larger context in which they exist.

7.2. Automation in CSIR

7.2.1. How SOCs View the Purpose of 'Automation'

Before addressing some of the discussion topics within the area of automation, this section expands details from Chapter 4 regarding how automation was viewed by the observed CSIRTs. One of the themes generated during analysis in Study 1 (Chapter 4) states that *automation is seen as a potential solution for low-level tasks and coordination, but considered out of reach for teams who don't have the support resources.* Some organizations viewed automation tools as a solution to addressing noise and inconveniences in the process, such as looking up who else worked on a similar issue in the past. (Onken, 2003) refers to this as 'conventional automation', or the idea that automation is a technical resource; the author differentiates this from 'cognitive automation', in which the technology has comprehensive knowledge and scrutinizing capability at a higher level of cognition than conventional automation. This dichotomy nicely summarizes some of the differences in perception about what automation can do in cyber security, which was dominated by the conventional definition.

Accordingly, automation was seen as helpful for overcoming some technical limitations in other programs, such as scripts generated by users. However, Study 1 found that if a team did not have strong and plentiful automation support resources, the general sentiment was that they were better off without it, despite the potential benefits. Supporting findings in Gutzwiller, Fugate, Sawyer &

Hancock (2015a), data indicated that automated scripts require constant adjustment as connected software tools get updated, creating a waterfall effect for needed updates from connected scripts and systems.

An alternative view on automation was that it is a pathway to elevate the maturity of the organization by relieving T1 analysts from doing mostly repetitive work. Managers interviewed in Study 1 expressed hopes that automation would allow the analyst operators to conduct deeper analysis of incident data and remove the need for humans to do low level filtering and routine incident response. This sentiment is especially important to acknowledge as the discussion evolves regarding recent technological advancement and future development in incident response.

7.2.2. Automation and the Evolution of the Analyst

In order to reduce workload, analysts currently rely on tools with some low level of automation that can help filter and monitor vast network systems, such as intrusion detection systems and security event and information management systems. Observations in Study 1 confirmed that many T1 tasks are already being automated at some level in three very diverse CSIRTs, even in the organizations that were considered ‘low maturity’ (NCSC-NL, 2015; Tetrick et al., 2016). While this level of automation is effective in reducing noise and vigilance tasks, it does not entirely mitigate the workload on T1 and T2 analysts. Pressures from the skills shortage, increased incident rates, and data influx have created a market environment that is technology-centric. Primarily, cyber security software companies are rapidly developing automation tools and platforms to help alleviate the pain caused by these pressures. Automation platforms aim to replace some amount of human analysts at the T1 level by automating repetitive *response* activities and guiding human response decisions with predetermined ‘playbooks’ (Bedell, 2019).

During Study 2 (Chapter 5), experts verified that much of the incident response process within the standard T1 level has been automated (though not necessarily instated in firms) through security orchestration, automation, and response (SOAR) platforms. SOAR platforms, in addition to other types of lower-level automation tools, advertise that they can assist improve situation awareness (Albanese et al., 2014), time for decision-making (Neiva et al., 2017), and integration of tools to reduce pivoting and manually piecing together information from different sources (Neiva et al., 2017). Experts indicated that the adoption of this type of automation platform is likely more near-term than long-term, which is supported by recent business trends (Bhargava, 2018; Oltsik, 2018b). With the SOAR platform configured and deployed, humans are expected to be supervisors over the automated T1 tasks and machine learning of the platform while performing T2 tasks themselves. Long-term, experts expect that T2 will be automated with supervisors over both T1 and T2, with the human making only critical decisions.

Though automation is currently viewed as a tool and not a teammate, the analyst level of the hierarchy discussed may evolve from human using automated tools to humans and automation becoming joint actors in CSIR. Evidence that this is already being pursued shows promise for this research path (Bunch et al., 2012). Considering technological trends in automating incident response tasks and replacing human analysts, this dissertation proposes that additional design considerations in human-automation teaming will become increasingly critical as the role of automation becomes more active in SOCs. Moreover the information sharing, collaboration, and knowledge sharing will still be relevant topics as the relationship evolves between higher-level human analysts and automation-based ‘teammates’.

Recalling literature presented in Chapter 2, dynamic function allocation (DFA) and adaptive automation are two relevant areas for investigating human-automation teaming in cyber security. More specifically, the research findings indicate that current assumptions about automation in security are founded in separation of tasks between human and machine. However, consideration of the current dependency of automation on humans when executing more complex tasks (i.e. supervised machine learning) prompts questions about how humans and machines will work together more closely and dynamically, and how automation can adapt to the human user (Gutzwiller et al., 2015b). Outputs from Study 3 provide good starting points for specific functions and functional needs for consideration in automation development in CSIR; some these functions directly support research for automated teammates in cyberspace (Lathrop, 2017).

7.2.3. Automation in CSIRTs

The installation and deployment of higher levels of automation, such as SOAR platforms, may also impact the operation of CSIRTs. Continuing the discussion regarding the evolution of the analyst becoming both human and machine, there could be additional effects at the CSIRT level of the hierarchy. Changes in the role of the human at the T1 level may inhibit the natural growth and progression currently seen in CSIRTs. Moreover, concepts from human-automation interaction (HAI) research can elevate to the team level, especially as the dynamics of the CSIRT change due to T1 automation. These concepts also connect to existing literature on supervisory control (Sheridan, 1992) and supervisory coordination (Caldwell, 2002).

One concern from participants in both Study 1 and Study 2 was the continued development of human analysts as policies and procedures become embedded in automatic or guided playbooks. Participants in Study 1 had conflicting viewpoints regarding how this would affect organizational maturity, as lower tier analysts would no longer have the benefit of a high rate of low-level incidents with which to practice and learn. Some higher tier analysts expressed concern that automating some of these functions would reduce the knowledge gained by training analysts, negatively impacting long-term critical thinking skills and overall growth. Essentially, the key question to balance automation development at T1 is, “Could ‘dumbing down’ T1 make the problem worse?” By decreasing the overall demand for skilled workers, the workers who still fill the seats of T1 analysts may not develop the needed understanding and expertise to progress to the next level of analyst. The subsequent concerns regarding fully automated T1 activities adds to this critical question. If the pipeline for higher-level analysts (mainly T1 analysts as they gain experience) is diminished or eliminated, the larger skills shortage is made worse at higher, more skilled levels of response in the future.

As human roles are overtaken by automation at the T1 level, there are also opportunities to explore HAI at the team level of incident response (Maymí & Thomson, 2018). One useful framework that could be applied to HAI in cyber teams is presented in Cuevas et al. (2007), which includes a broad set of topics within team cognition. Human supervision of automated T1 tasks is imminent, and warrants investigation regarding supervised machine learning and validation of decisions (especially in a dynamic threat setting) across multiple human entities. Automation is quickly making strides in ingesting security data to make decisions based on machine learning algorithms (Staples & Sullivan, 2018). However, it is unclear how much the learning processes will be supervised by humans, and thus how automation will *actually* affect the overall workload of the analysts. Interviewed analysts and managers had not discussed or mentioned these implications during Study 1, as the overall goals of installing automation seemed to be at the forefront of automation-related dialogues. However, considerations stemming from the VSM analysis highlight a need for increased discussion about how new tasks and processes will be managed by human counterparts at each level of incident response, and how current subsidiary roles will change.

In reference to findings and previous discussion about the current state of CSIRTs regarding collaboration and information sharing, an increase in the role of automation prompts questions about what information sharing will look like in a given organization, and how that will impact the overall dynamics and viability. Again, participants had not expressly addressed implications of higher levels of automation on the team's dynamics. However, after installation of technologies like SOAR, security organizations will likely reflect on the larger effects of how the technology prompts change elsewhere within the team. Topics in this area may include: responsibility and accountability of automation trust in decisions made by automation and design of information sharing between human and automation team members.

Automation across other research areas is expected to change the role and education of the human operator as well as team-level operations (Barnes & Jentsch, 2010; Best, 2018; Shively, Lachter, Koteskey, & Brandt, 2018). Conceivably similar effects will be seen in security as automation overtakes T1 tasks, though one might detect some amount of hyperbole in recent articles about the promise of automation in cyberspace. Recent reports highlight challenges in implementing automation due to the necessary resources required, especially in relation to human expertise (Filkins, 2019; Ponemon Institute, 2019). Data collected in Study 1 indicate that role and education of human teammates may not be fully in the scope of what organizations are considering when they think about automation deployment in CSIRTs, indicating opportunities to further explore these effects and the role of human expertise in security teams.

7.2.4. Automation and Security Organizations

Continuing the discussion of mission and culture impacts on CSIRTs, these factors also have a direct impact on the adoption and use of automation within security organizations. Findings in Study 1 indicated that some managers view automation as a way to increase the maturity of an organization with respect to security. The goal in this case is to improve the overall security position of the company or firm by creating consistency of coverage and execution of tasks at the 'front line'. In turn, this reduces the risk of breaches. Experts in Study 2 believed this is the desired direction of the field: to reduce risk through less vulnerability and more consistency. Yet, the ability of an organization to act on these goals is directly impacted by the agency given to the security group by higher levels of management up the to 'System 5' level in VSM terms. These findings frame a discussion about how the priorities of a given organization impact decisions to purchase, configure, and deploy automation, especially in relation to their current human teams.

Organizations observed in Study 1 operated under different missions, which had a direct effect on their CSIRTs. One firm had the mission of improving their overall security position, as it was vital to their viability and that of their products (and thus, their customers). Within this team, discussion about automation was in direct alignment with this mission, and focused on improving the maturity of the overall security group. Within the other two teams, the organizational missions were more focused on maintaining stable operations; the interpretation of this at the security team level was to minimize disruption. The teams in these two organizations were significantly smaller with only one shift of coverage, and both had less emphasis on technological advancement at the team level. When discussing automation with these teams, they largely believed that these technologies were out of scope for their role within the organization and their respective budgets. Note that security spending over the last several years has increased (ISACA Cybersecurity Nexus, 2017; Scale Venture Partners, 2017; Vieane et al., 2016), but was a significant barrier for many security teams to expand and improve operations. Overall the research findings support the idea that the upper echelons of an organization will drive the adoption of automation within CSIRTs.

7.3. Research Contributions

The objectives of this dissertation were to study expertise in the context of CSIRTs, specifically in how it relates to information sharing, and to develop subsequent functional requirements for automation deployed in this domain. In meeting these objectives, this dissertation has made several theoretical, methodological, and practical contributions. The below subsections describe these contributions in terms of research (theory and methods) and application.

7.3.1. Theoretical and Methodological Contributions

One major issue currently being studied and addressed in this field is the labor shortage (Bureau of Labor Statistics, 2016), which is being addressed in several ways, including education, training, and technological development (Assante & Tobey, 2011; Chen et al., 2014; Cobb, 2016; Hoffman et al., 2012; Peusquens, 2017; G. White & Granado, 2009). However, trends in business literature suggest that these methods have been largely ineffective at closing the skills gap (HCL Technologies, 2019; Morgan, 2017; Oltsik, 2019). Research has studied knowledge requirements by way of KSAs (Chen et al., 2014; Cobb, 2016; Newhouse et al., 2017), which has resulted in guideline documents (National Initiative for Cybersecurity Careers and Studies, 2017) that are extremely broad and difficult to apply through recruitment, especially as human resources has already been identified as weak for these positions (Cobb, 2016). One contribution of this dissertation is that it holistically explored expertise by using the six dimensions of expertise construct as a categorization schema and focusing on the lowest level of analyst. As suggested by the creators of the construct, there is need for expanding the study of generalized expertise in different context areas (Garrett et al., 2009). The expertise framework developed in this dissertation delivered new potential areas of expertise for researchers to consider within CSIRTs. This framework can also be applied to research in machine learning and automation development, especially as the field expands into expert systems (Gamal, Hassan, & Hegazy, 2011; Neiva et al., 2017; Staples & Sullivan, 2018).

In addition to the contribution to the cyber security research domain, this dissertation also makes a contribution to the research literature on expertise. The discussion of findings expanded the original six dimensions of expertise construct by developing an expertise framework that includes other aspects of expertise in CSIR. In so doing, this dissertation increases the ability of the six dimensions of expertise construct to reach new research communities and application domains.

Research findings in this dissertation identified new expertise factors within CSIR specifically by using qualitative methods traditionally founded in building theory; these factors present viable options for future research in expertise and CSIRT recruitment and operations. By analyzing expertise for specific tasks in T1 incident response, context-specific expertise factors, such as policy and self-awareness, were identified. These additions highlight the complexity of the environment as well as the knowledge required to conduct incident response activities. Furthermore, in evaluating the inter-rater reliability, findings identified some discrepancies in how the raters applied components of the original construct. More specifically, there was considerable disagreement in delineating subject matter expertise versus situational context expertise. This disagreement was perhaps due to the fact that both dimensions are required for the task being analyzed, and additional research could be conducted to study correlations between dimensions for particular contexts.

The cyber security research domain is still relatively immature, especially regarding empirical studies of CSIRTs. Many studies included in the literature review approach the study of CSIRTs from a macro level perspective, often focusing on mapping general processes and identifying common issues across these types of teams (Chen et al., 2014; Reed, Abbott, Anderson, Nauer, & Forsythe, 2014; Ruefle et al., 2014; Steinke et al., 2015; Tetrick et al., 2016). Indeed, some of these studies provide useful outputs for practitioners to consider for CSIRT operations (Chen et al., 2014; Steinke et al., 2015; Tetrick et al., 2016). However, this level of analysis does not provide enough contextual information to help researchers consider environmental differences between teams (Ahrend et al., 2016) or to determine robust design requirements for technological development. In essence, the extant CSIRT operations literature paints a very generalized picture for a very nuanced setting. Another contribution of this dissertation is the micro level study of teams in different operational settings, which included vertical views of the respective organizations identified as needed in literature (Beznosov & Beznosova, 2007). By sampling teams in different sectors, key differences could be captured in various types of operations, especially in regards to mission, structure, and autonomy. This sampling approach combined with a multi-method design helps advance the CSIRT literature base by capturing and connecting new factors to operational context, which can be included in future research.

Literature addressing CSIR functions is generally grouped into two categories: those that employ human-centered approaches, and those that are technologically focused. The latter category has been known to dominate the literature base (Beznosov & Beznosova, 2007; Vieane et al., 2016). In response, recent publications have called for more research focused on human factors (Gutzwiller et al., 2015b; Vieane et al., 2016), and social and psychological approaches (Dawson & Thomson, 2018; National Academies of Sciences Engineering and Medicine, 2017). The research conceptual framework applied in this dissertation successfully connected different research traditions, effectively executing a multi-method approach to study CSIRTs. Human factors and systems engineering were two research traditions applied to better understand how human teams share information in this setting, and how technology plays a role. These traditions are well suited for capturing the complexity of the environment, and well outfitted with frameworks and methodologies for applying to a given problem. Indeed, other literature in these domains could be useful in studying CSIRTs by applying contextual design methods (Holtzblatt & Jones, 1993) and using system-of systems approaches to include hierarchical identification of resources, operations, policy, and economics (DeLaurentis, 2005; Guariniello et al., 2016) in security operations. In summary, through applying a multi-tradition, multi-method approach, this dissertation illustrated that numerous perspectives and approaches are valid when studying problems in cyber security incident response.

Furthermore, the multi-tradition, multi-method design used in this dissertation shows that, within cyber security settings, results from a qualitative methodology for data collection can be successfully translated into actionable items for automation developers who design and implement functional solutions and researchers who perform quantitative research. The methodology, which included needed ethnographically informed methods (Ahrend et al., 2016), employed was able to effectively produce these results in a reasonable timeframe with limited resources; researchers could consider something similar when exploring similar problem spaces that present challenges with access, resource, and time limitations. Moreover, the methodology can be replicated in future studies to expand upon findings in different teams.

A major and recent area of development within CSIR is the SOAR platform, which aims to automate and ‘orchestrate’ incident response with predetermined procedures and some learned responses. However, research literature on SOAR is extremely limited. In order to provide better understanding of potential shortcomings of this new technology, Study 3 included a gap analysis, identifying areas in which the platforms may not meet the needs of the analysts. This contribution adds value to studying SOAR platforms beyond basic usability measures, focusing on functionality instead. Additionally, literature on SOAR focuses more on developing new automated capabilities to decrease workload with the intent of creating more time for individual, in-depth analysis (Bedell, 2019; Lathrop, 2017; Neiva et al., 2017); these approaches do not necessarily focus on improving effectiveness of human teams. As information sharing between human entities was a key part of this dissertation, the results offer a contribution to defining SOAR needs at a team coordination level. Finally, the functions identified as part of Study 3 advance potential areas of human factors research on HAI. While human factors literature has identified HAI as a relevant topic in cyber security (Gutzwiller et al., 2015b), the research findings in this dissertation help focus functions for study and evaluation, especially early in the design process. Furthermore, these results address previously identified gaps regarding how technology can play a role in supporting incident responders (Werlinger et al., 2010).

7.3.2. Practical Contributions

The practical contributions of this dissertation are divided into two categories. First, contributions aimed at management in current CSIRT operations help by providing specific areas of potential investigation or improvement that can be immediately applied to a given team. Second, contributions for developers provide direction for future automation development, specifically regarding functional requirements that could help CSIRTs as they currently operate, and how they might operate in the future.

7.3.3.1. For CSIRT Management

The management-oriented contribution is summarized as identified trends and variance in CSIRT organizations. This dissertation did some level of comparison across CSIRTs in different types of organizations, revealing potential considerations for team design and management based on the operational setting. Analysis revealed that teams in different sectors might vary greatly in terms of needs and constraints, which should be factors for consideration when trying to improve performance. Within the teams themselves, this research identified more specific insights regarding communication (including meetings, shift handoffs, and escalations) that could be taken back to an organization for reflection and internal study. These insights also draw attention to organizational maturity and how strategic a given team can be based on their organizational structure, culture, and team capability. All of these factors are potential discussion topics for CSIRT management within a given organization.

Some specific recommendations for management are listed below. These suggestions address only a portion of the important revelations of this dissertation research. Additional insights may be extracted from themes in Studies 1 and 2, with additional considerations related to automation in Study 3. Recommendations include:

- Consider how the mission of the overall organization supports the security organization, and how this operational focus affects the autonomy of the SOC.
- Consider how the organizational structure supports autonomy of security executives and their respective reporting structures.
- Consider what feedback is given to the lowest level analyst; feedback should support growth and learning.
- Identify clear hands-on learning opportunities for analysts, and a clear path for development from the lowest level of analyst to higher or more specialized positions.
- Identify where accountability is created between different tiers of incident response, including through peer feedback and shift handoff meetings.
- Identify how often documentation is updated, and by whom it is updated, to ensure validity and usefulness of reference documents.

- Identify how knowledge is captured within the organization: documentation, wikis, etc.; Incorporate access behaviors of users, usefulness and relevance of the information with respect to time and user base, and usability of medium.
- Consider how teamwork and team processes affect the stability of the team itself; Evaluate collaboration practices, and ensure environmental structure supports team interactions.
- Consider perceptions of roles by other analysts: determine perceptions of who is part of the team and who is not. Consider perceived value, and the effect on turnover and retention.
- Identify the goals of installing future automation to assist with security operations; Identify how this affects your current staffing, development, and effectiveness; Test automation, or observe in other organizations, before full deployment.
- Identify how communication is supported in your organization during incident response; Ensure communication is appropriate and adequate in the eyes of your analysts.

7.3.3.2. For Technology Developers

For developers, the applied contributions of this dissertation are mainly the discussion and concepts of new design considerations for future automation development and implementation in CSIRTs. Reiterating the problem statement in this dissertation, the labor shortage (and thus, expertise shortage) cannot be completely addressed with software solutions or workforce training. One way of changing the approaches is to consider human expertise and automated teammates in tandem, marrying the respective research traditions to address the problem in a new way.

One goal of current SOAR technologies is to embed subject matter expertise into the technology itself. In so doing, developers aim to alleviate some of the shortage of said expertise, especially at lower tiers of security operations. Indeed, addressing the fact that there are not enough people to meet minimum coverage is a high priority to keep businesses running securely. Yet, when extrapolated to the near future, this particular approach may actually make the problem worse. Recent reports indicate that the shortage of security analysts spans all levels of the hierarchy, from expert analysts to CISOs (VIB & Demisto, 2018). Thus, the labor shortage problem persists in the entire pipeline. Augmenting the first line of defense (or even second line of defense) may help businesses maintain operability, but may actually prevent those individuals from gaining expertise in the traditional way: through repeated experience and thinking through real problems.

Literature supporting cyber education emphasizes the need for live training of cyber analysts (Urias, Leeuwen, Stout, & Lin, 2017), which could feasibly take place while conducting live incident response (on-the-job experience). Embedding expertise in the system creates requirements about establishing transparency and effective transfer of knowledge to the human using the system so that their knowledge and expertise continue to grow and develop (Buchanan et al., 2018). Continuous learning and development of analysts is critical to maintaining vital skills and increasing knowledge of new threats (Oltsik, 2019); a recent survey of security analysts suggests that analysts themselves rate improving their skillsets with high importance (VIB & Demisto, 2018). However, these statements do not align with current SOAR goals of automating at the source level of skilled analysts, creating some concern that SOAR technologies may not successfully address the larger problem.

Other aspects of CSIRT operation at the team level are also pertinent discussions that involve technology developer stakeholders. Communication in particular has facets unacknowledged by technology companies in reference to tools and systems installed to help facilitate communication. Study 3 findings indicated that software companies are addressing "communication" issues by creating better chat features in programs. Indeed, some practical issues regarding physical team separation and segregation of tiers may be overcome using these features. Yet, deeper aspects of communication in Themes 2 and 3 within Study 2 indicate that better understanding of these factors and how they affect decisions could be beneficial developing future AI in incident response. Trust is already a prominent topic in AI literature; it is reasonable to connect trust to how humans interact (and not just human-automation interactions), and that trust is still central to communication and collaboration, with or without technology. While future development of tools may help analysts facilitate new information sharing patterns and overcome personal bias, it may be beneficial to connect scientific research of trust and collaboration between humans to AI development such that the sensors and feedback of the technology adequately meet the needs of the human components.

7.4. Research Limitations

7.4.1. Practical Limitations

The practical limitations of this research were mainly those imposed by access to participants, participants' time and financial resources. Time constraints of the participants provided limits to the length of interviews and observations, which necessarily limited the richness of the data collected. Within the scope of Study 1, mitigation strategies included revisiting participants for short verification questions. In Study 2, data were audio recorded when permitted by the participants, allowing for thorough data capture of responses. I was not able to secure additional observers or interviewers with available resources, which were limited to my own financial income and time.

7.4.2. Data and Methodology Limitations

In addition to the practical limitations above, the main limitation affecting the data and methodology was accessibility to teams and participants. Though this was introduced in Chapter 4, recruitment of individuals in teams was extremely difficult. This was exacerbated by the on-site nature of the data collection for Study 1, which required multiple conversations with upper level management at each firm to ensure confidentiality. Furthermore, once on-site, data recording was not permitted in any way; in some cases, active data collection was not permitted in the incident response room. To address the potential limitation of overall access, relevant members of the team were included in a meeting to review play-by-play reenactment of past incidents for the purpose of data collection. However, removing participants from the context of their work introduces another limitation of ecological validity for the data (since this data would be based solely on what participants could recall), and time for the participants to spend away from their workstations.

7.4.3. Analysis & Discussion Limitations

The three studies in this dissertation revealed many interesting findings, some of which were beyond the scope of the research questions. Due to time and length constraints, and in order to focus the contributions of the research, the scope of discussion is limited to issues directly relating to the research questions and considerations for improvement in security and future technological design of automation for CSIRTs.

Additionally, the findings presented in Study 3 are concepts based on the data collected in this dissertation and available literature. Thus, the ideas are inherently limited by this scope, but can be overcome with further research to verify and validate design elements, especially with a wide set of stakeholders. Future elaboration of “form” to complement the functions presented will help complete the concepts for product development.

CHAPTER 8. CONCLUSIONS

8.1. Designing Automation for CSIRTs

CSIRTs have been previously identified as complex and diverse teams in time-intensive, stressful environments. This dissertation research has corroborated these statements in context, and it has added another layer of detail to the extent of complexity and diversity seen in different teams in different operational settings. At present, automation does play a minor role in how CSIRTs perform incident response, and trends in the domain indicate an imminent increase in the amount of automation and the self-sufficiency of machines at higher levels of response. In fact, the current marketplace is largely depending on this increase to alleviate pressures of the skills shortage. However, there are many factors to consider when designing automation as it increases in both autonomy and interaction with the human analyst, and with human teams. The research findings emphasize that these factors are not always directly related to automation itself, but rather the environmental and organizational context in which it is deployed. Moreover, any negative impact from factors such as mission, culture, and team structure cannot necessarily be mitigated by the installation of automated tools. This dissertation reveals that, though automation is heralded as an ultimate solution to problems in incident response, it cannot fix everything.

Current automated tools are able to quickly and reliably execute low-level monitoring and filtering tasks. SOAR platforms promise more, and the tools claim to be able to do actual response activities and correlations across threats: tasks normally assigned to a human. However, the quality and effectiveness of deployed SOAR platforms has not been validated with research studies, and the adoption rate, while expected to rise, is still low. Traditionally within CSIR, tool usability (or lack thereof) is a topic of discussion in human factors and other related design disciplines. Indeed, usability of tools in CSIR is notoriously poor. Discussions about the future of automation and orchestration platforms in CSIR transcend screen-level topics of usability, expanding into human-automation teaming. As automation and orchestration gain ground in scope and capability, trust and communication factors become more important in establishing fluid team information sharing and decision making. Designing a teammate goes beyond mere task performance and extends into team human factors, as they dynamically collaborate and coordinate during incident response.

Research on human-automation teaming in CSIR is relatively limited, and studies do not have a clear connection to the design of new automation platforms currently on the market. As new automation is developed and deployed, new opportunities arise to study how teams adapt to this technology. Furthermore, this dissertation revealed some disconnects between needs of human analysts and SOAR capabilities as advertised, indicating that there might be a gap between research and design in this domain.

In order to bridge the gap in automation design, this dissertation has also proposed functional requirements that could meet analyst needs. These proposed requirements, founded in empirical research in CSIRTs, focus on collecting and sharing information across different tools and team members. Future CSIR operations processes and protocols should accommodate multiple inputs, outputs, and feedback to support incident response activities, especially as teams evolve to include multiple humans *and* multiple automation entities. The proposed requirements represent new opportunities to inform design of human-automation teaming, such that the human component is incorporated into the development of the platform.

8.2. Broader Implications

This dissertation draws clear connections between CSIRT operations and automation design, and acts as a focus for opportunities to study and advance research in cyber security, specifically for human factors. Calls to investigate this domain (Gutzwiller et al., 2015b; Vieane et al., 2016) have expressed not only the opportunity, but also the *need* for human factors in this critical area. As the threat landscape grows in scope and complexity, the need for human defenders and well-designed technology becomes crucial to securing systems and networks worldwide. Findings in this dissertation highlight specific areas for focus within incident response in which human factors can be applied, particularly in supporting information sharing and coordination.

The methodology employed in this dissertation is a testament to the breadth of approaches that can be used to study complex environments, particularly in cyber security. Expanding the set of qualitative tools and techniques in cyber security research augments traditional, technologically focused efforts, complementing development with design considerations. In this sense, research can smoothly shift from academic ventures to applied projects. Within CSIR, the potential benefits of research transfer that bridges human and machine elements can be directly observed in current development projects for SOAR platforms.

Finally, while the problem space for this dissertation focused on CSIR, the methodology and findings might also be applicable in other domains that include human-automation teaming. Results highlight opportunities to better understand this team-based relationship in dynamic and complex environments, spurring questions about how teammates could be designed versus how technology should be designed. The paradigm shift from automation-as-a-tool to automation-as-a-teammate requires new perspectives regarding context, communication, and culture, and how these impact team operations and the role of technology.

8.3. Future Research

In addition to the contributions presented in Chapter 7, this dissertation has identified several thrusts of future research that will build upon the work in this dissertation. First, research findings of this dissertation reiterate the need for social science research the in cyber security domain, citing specific areas for investigation identified from the findings of Studies 1 and 2. Second, ideas are presented for expanding the use of the six dimensions of expertise framework, including measurement and assessment. Lastly, this section concludes with thoughts on how to build additional capability in security automation, using the outputs of Study 3 as a springboard for developing new automation functions and how they can fit into CSIRTs.

8.3.1. Bridging Social Science Research in Cyber Security

The need for social science research in security (National Academies of Sciences Engineering and Medicine, 2017) was a key inspiration for this dissertation. While this dissertation makes a contribution by using social science methodology, the security domain presents a wide range of opportunities and issues to explore from a social science perspective (Ahrend et al., 2016; Beznosov & Beznosova, 2007; T. R. Chen et al., 2014; Werlinger et al., 2010), including organizational and social factors as they relate to CSIRT operations and performance. This dissertation has also identified specific areas within the area of collaboration that could benefit from these perspectives.

One potential area of investigation using social science methods is intra-tier and inter-tier collaboration practices and patterns between analysts. Examining what drives individuals in distributed teams to work together (or not), and specific outcomes of this collaboration would aid in understanding more potential needs of technology development as it relates to expertise transfer. Additionally, different dynamics and specific cases for inter-tier collaboration might provide additional insights regarding where automation tools can assist in the investigation and response process. This dissertation has identified three such examples of collaboration for potential study.

The first case expands upon creating shared operational pictures between analysts, and how they might co-monitor an incident before deciding what actions to take. Shared awareness was identified in Study 2 as an information sharing need, and included in Study 3 as a concept for future development. Research questions might explore what that monitoring process looks like from the human perspective, how the analysts each utilize the technology available to them (on shared or separate screens), and how the collaboration efforts affect knowledge transfer between the expert and novice. The second case is when two higher tier analysts with different but overlapping expertise profiles collaborate on an incident and decide together how to proceed. Research questions might explore the transaction level of how they apply knowledge and collaboratively problem solve, especially when their opinions conflict about the correct course of action. The last case includes a security analyst collaborating with an external customer or client (such as in a managed security service provider, or MSSP) to resolve an incident. As the information sharing crosses organizational boundaries, questions might explore what information is shared and in what way (i.e. mode, style, urgency); illustrating this type of collaboration in comparison to intra-organization might reveal opportunities for technology to bridge boundary-related gaps.

8.3.2. Developing Research Tools Using the Six Dimensions of Expertise Framework

Another area of future research is in expanding methodologies and research tools to apply the six dimensions of expertise construct across different application areas. The six dimensions of expertise proved to be a useful construct for expanding discussion on understanding and transfer of expertise in cyber security. It conceivably could provide value in other complex domains that require more flexible frameworks of expertise and how it is applied in dynamic settings. Future research in this area could focus on quantification or assessment tools for measurement of different dimensions by connecting relevant theories and frameworks, such as situation awareness (as it relates to situational context expertise) and transactive memory (as it relates to expert identification expertise). Efforts to complement the existing construct with assessment methods may help advance empirical research regarding dimensional expertise.

Developing measurement instruments would involve several major components to ensure content and construct validity. Mainly, further development in joining related constructs is needed before developing a measurement. Comparison of related constructs, what tools have already been developed, and how these might be combined are some suggested activities.

Future research might also aim to identify connections or correlations between different dimensions, especially within a given domain, to help increase ecological validity of research. Furthermore, investigating specific relationships between dimensions would help overcome certain data analysis struggles noted in this dissertation, specifically in relation to multiple codes per segment and overlap between dimensions noted in Study 2. These connections between dimensions may help determine which design requirements are related or dependent when translating expertise requirements to technological development.

8.3.3. Developing New Concepts for Security Automation

The current technological trends and business literature focus on automation as a necessary development area in security in order to alleviate labor shortages, increase security coverage, and increase efficiency and consistency of security analysis. Though Study 3 proposes future directions of SOAR capabilities, there are myriad paths for future research around automation development. This section presents a select few based on findings and challenges from Study 3.

First, research can explore additional development of a baseline methodology within cyber security specifically to help develop new concepts while also overcoming application-specific issues. Using the SE methodology in Study 3 presented challenges in the validation process; it required more quantitative outputs and extensive access to hard-to-reach populations, including multiple levels of analysts and management. Concept development methodologies derived from user experience and human factors domains might help overcome those challenges. These methodologies are also designed to accommodate qualitative data (collected from human users) and translate into product design requirements. The methodology employed in this dissertation builds the foundation for more formalized design and development methodologies in cyber security, especially as the field grows in applied research and tool development.

Second, with the expected increases in deployment of automation, research should study the effects of automation on current CSIRT team effectiveness. Though there is much hype around the *potential* of automation in these settings, the larger system effects and indirect impacts have not been identified. Studies to establish baselines before installation and collect data from after deployment can assess changes and improvements in CSIRT processes and performance. Moreover, studies similar to Study 1 can investigate changes in team dynamics before and after deployment. Furthermore, practical studies focused on usability of security tools could contribute valuable insights for interface design. While SOAR platforms currently boast 'easy-to-use' interfaces, validation of those claims in different settings could help establish standards for security tool design.

Finally, to expand upon the outputs of this dissertation, research should focus on technological development of "automated assistants" beyond SOAR platforms. Taking into account perceptions and expectations of analysts at different levels will be helpful in identifying potential adaptation needs of the assistant with changes in human roles and expertise levels. Literature on the evolution of expert systems (Buchanan et al., 2018) and the development of artificial intelligence may be useful in helping guide SOAR development. Additionally, publications with applied focus on intelligent assistants can provide useful insights regarding how humans in other contexts interact think about how automation plays a role in their daily activities (Budiu & Whitenton, 2018).

REFERENCES

(ISC)^2. (2018). *Cybersecurity Professionals Focus on Developing New Skills as Workforce Gap Widens: (ISC)^2 Cybersecurity Workforce Study, 2018*.

Abbass, H. A., Petraki, E., Merrick, K., Harvey, J., & Barlow, M. (2016). Trusted Autonomy and Cognitive Cyber Symbiosis: Open Challenges. *Cognitive Computation*, *8*(3), 385–408.

Abdul, A., Vermeulen, J., Wang, D., Lim, B. Y., & Kankanhalli, M. (2018). Trends and Trajectories for Explainable , Accountable and Intelligible Systems : An HCI Research Agenda. In *Conference on Human Factors in Computing Systems*. Montreal, QC, Canada: ACM Press.

Adler, M. (1988). An Algebra for Data Flow Diagram Process Decomposition. *IEEE Transactions on Software Engineering*, *14*(2), 169–183.

Ahrend, J. M., Jirotka, M., & Jones, K. (2016). On the collaborative practices of cyber threat intelligence analysts to develop and utilize tacit Threat and Defence Knowledge. In *Cyber Situational Awareness, Data Analytics And Assessment (CyberSA), 2016 International Conference On* (pp. 1–10). IEEE.

AIAA Guide to the Preparation of Operational Concept Documents (ANSI/AIAA G-043A-2012). (2012). Reston, VA, USA.

Albanese, M., Cam, H., & Jajodia, S. (2014). Automated Cyber Situation Awareness Tools and Models for Improving Analyst Performance. In R. E. Pino, A. Kott, & M. Shevenell (Eds.), *Cybersecurity Systems for Human Cognition Augmentation* (pp. 47–60). Cham: Springer International Publishing.

Annett, J., & Stanton, N. A. (2000). *Task Analysis*. London, UNITED STATES: CRC Press.

Armistead, C., & Meakins, M. (2002). A Framework for Practising Knowledge Management. *Long Range Planning*, *35*(1), 49–71.

Asher-Dotan, L. (2015, February). Metrics that Matter: New Measurements of IT-Security Program Effectiveness. *Cybereason*.

Assante, M. J., & Tobey, D. H. (2011). Enhancing the cybersecurity workforce. *IT Professional*, *13*(1), 12–15.

Auerbach, C., & Silverstein, L. B. (2003). *Qualitative data: An introduction to coding and analysis*. New York University Press.

Bada, M., Creese, S., Goldsmith, M., Mitchell, C., & Phillips, E. (2014). Computer Security Incident Response Teams (CSIRTs): An Overview. *Global Cyber Security Capacity Centre*, (May), 1–23.

Baker, S. E., & Edwards, R. (2012). *How many qualitative interviews is enough? Expert voices and early career reflections on sampling and cases in qualitative research*. Southampton, UK.

Barnes, M., & Jentsch, F. (2010). *Human-Robot Interactions in Future Military Operations*.

Bedell, C. (2019). *Definitive Guide to SOAR*. (S. Shuttleworth, Ed.). Annapolis, MD: CyberEdge Press.

Beer, S. (1984). The Viable System Model: its provenance, development, methodology and pathology. *The Journal of the Operational Research Society*, *35*(1), 7–25.

Beer, S. (1995). *Diagnosing the System for Organizations*. Wiley.

Beitzel, S., Dykstra, J., Toliver, P., & Youzwak, J. (2018). Exploring 3D Cybersecurity Visualization with the Microsoft HoloLens. In D. Nicholson (Ed.), *Advances in Human Factors in Cybersecurity: Proceedings of the AHFE 2017 International Conference on Human Factors in Cybersecurity, July 17−21, 2017, The Westin Bonaventure Hotel, Los Angeles, California, USA* (pp. 197–207). Cham: Springer International Publishing.

Bernard, H. R. (2006). *Research Methods in Anthropology: Qualitative and Quantitative Approaches* (4th ed.). Oxford, England: AltaMira Press.

Best, J. (2018, November). Robots and the NHS: How automation will change surgery and patient care. *ZDNet*.

Beznosov, K., & Beznosova, O. (2007). On the imbalance of the security problem space and its expected consequences. *Information Management & Computer Security*, *15*(5), 420–431.

Bhargava, R. (2018). Gartner SOAR Adoption Rate Prediction: From 1% to 15% by 2020 - Why Should You Care? *InfoSec Island*. Wired Business Media.

Billings, C. E., Lauber, J. K., Funkhouser, H., Lyman, E. G., & Huff, E. M. (1976). *NASA aviation safety reporting system*. Moffett Field, CA.

Birks, Y., Harrison, R., & Bosanquet, K. (2014). An exploration of the implementation fo open disclosure of adverse events in the UK: a scoping review and qualitative exploration. *Health Services and Delivery Research*, *2*(20). https://doi.org/10.3310/hsdr02200

Bishop, M., Burley, D., Buck, S., Ekstrom, J. J., Futcher, L., Gibson, D., … Parrish, A. (2017). Cybersecurity Curricular Guidelines. In M. Bishop, L. Futcher, N. Miloslavskaya, & M. Theocharidou (Eds.), *Information Security Education for a Global Digital Society: 10th IFIP WG 11.8 World Conference, WISE 10, Rome, Italy, May 29-31, 2017, Proceedings* (pp. 3–13). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-58553-6_1

Blanchard, B. S., & Fabrycky, W. J. (2006). *Systems engineering and analysis* (4th Ed.). Prentice Hall, Inc., Upper Saddle River, NJ (USA).

Boden, P. (2016). The Emerging Era of Cyber Defense [Web log post]. Retrieved January 1, 2017, from https://blogs.microsoft.com/microsoftsecure/2016/01/27/the-emerging-era-of-cyber-defense-and-cybercrime/

Bonacina, M. P. (2017). *Automated Reasoning for Explainable Artificial Intelligence*.

Borghetti, B., Funke, G., Pastel, R., & Gutzwiller, R. (2017). Cyber Human Research from the Cyber Operator’s View. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *61*(1), 350–350.

Boroomand, F., Fereidunian, A., Zamani, M. A., Amozegar, M., Jamalabadi, H. R., Nasrollahi, H., … Lucas, C. (2010). Cyber security for Smart Grid: A human-automation interaction framework. In *Innovative Smart Grid Technologies Conference Europe (ISGT Europe), 2010 IEEE PES* (pp. 1–6). IEEE Publishing.

Boudreau, M., Loch, K., Robey, D., & Straub, D. (1998). *Going Global: Using Information Technology to Advance the Competitiveness Of the Virtual Transnational Organization*. *Academy of Management Perspectives* (Vol. 12).

Bourget, J. (2017). *Addressing Analyst Fatigue in the SOC*. Arlington, VA.

Bowen, B. M., Devarajan, R., & Stolfo, S. (2011). Measuring the human factor of cyber security. In *2011 IEEE International Conference on Technologies for Homeland Security (HST)* (pp. 230–235).

Bradshaw, C. (2015, June 10). The rise of the centaurs. *The Drum*.

Breed, B. (2003). The reflective abilities of expert and novice learners in computer programming. In *British Educational Research Association Annual Conference*. Edinburgh, UK.

Brooks, C. (2018). Security Orchestration , Automation and Response (SOAR) - The Pinnacle For Cognitive Cybersecurity. *Security Essentials*. AlienVault.

Brooks, F. A. (1960). Operational Sequence Diagrams *. *IRE Transactions on Human Factors in Electronics*, *1*(March), 33–34.

Brown, P., Christensen, K., & Schuster, D. (2016). An Investigation of Trust in a Cyber Security Tool. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *60*(1), 1454–1458.

Brundage, M., Clark, J., Allen, G. C., Flynn, C., Farquhar, S., Crootof, R., & Bryson, J. (2018). *The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation*.

Buchanan, B. G., Davis, R., Smith, R. G., & Feigenbaum, E. A. (2018). Expert systems: A perspective from computer science. In K. A. Ericsson, R. R. Hoffman, & A. Kozbelt (Eds.), *Cambridge Handbook of Expertise and Expert Performance* (2nd ed.). Cambridge, UK: Cambridge University Press.

Budiu, R., & Whitenton, K. (2018). What Could an Intelligent Assistant Do for You? A Diary Study of User Needs. Nielsen Norman Group.

Buford, J. F., Lewis, L., & Jakobson, G. (2008). Insider threat detection using situation-aware MAS. *Proceedings of the 11th International Conference on Information Fusion, FUSION 2008*.

Bunch, L., Bradshaw, J. M., Carvalho, M., Eskridge, T., Feltovich, P. J., Lott, J., & Uszok, A. (2012). Human-Agent Teamwork in Cyber Operations: Supporting Co-evolution of Tasks and Artifacts with Luna BT - Multiagent System Technologies. In I. J. Timm & C. Guttmann (Eds.) (pp. 53–67). Berlin, Heidelberg: Springer Berlin Heidelberg.

Bureau of Labor Statistics. (2016). Information Security Analysts. Retrieved October 11, 2017, from https://www.bls.gov/ooh/computer-and-information-technology/information-security-analysts.htm

Burke, C. S. (2005). Team Task Analysis. In N. Stanton, A. Hedge, K. Brookhuis, E. Salas, & H. W. Hendrick (Eds.), *Handbook of Human Factors and Ergonomics Methods* (pp. 1–8). Boca Raton, FL: CRC Press.

Burley, D. L., & Bishop, M. (2011). *Summit on Education in Secure Software Final Report*. *C-Suite Challenge 2019: Survey*. (2019). New York, NY, USA.

Cain, A. A., & Schuster, D. (2014). Measurement of situation awareness among diverse agents in cyber security. In *Cognitive Methods in Situation Awareness and Decision Support (CogSIMA), 2014 IEEE International Inter-Disciplinary Conference on* (pp. 124–129). IEEE.

Caldwell, B. S. (2002). Developing Tools to Support Knowledge Synchronization in Distributed Supervisory Coordination. In G. Luczak, H. , Çakir, A. E., Çakir (Ed.), *6th International Scientic Conference on Work with Display Units: WWDU 2002 ± World Wide Work (Berlin: ERGONOMIC)* (pp. 554–556). BERLIN.

Caldwell, B. S. (2008). Knowledge sharing and expertise coordination of event response in organizations. *Applied Ergonomics*, *39*(4), 427–438.

Caldwell, B. S. (2009). Perspectives on Systems Engineering and Impacts on SE Education. In *Proceedings of the 2009 Industrial Engineering Research Conference*.

Caldwell, B. S. (2011). Connection, Coupling, and Persistence in Online Social Networks. In D. Haftor & A. Mirijamdotter (Eds.), *Information and Communication Technologies, Society and Human Beings: Theory and Framework (Festschrift in honor of Gunilla Bradley)* (pp. 346–354). Hershey, PA: IGI Global.

Caldwell, B. S. (2015). Framing, Information Alignment, and Resilience in Distributed Human Coordination of Critical Infrastructure Event Response. *Procedia Manufacturing*, *3*, 5095–5101.

Cano, J., Hernandez, R., Pastor, R., Ros, S., Tobarra, L., & Robles-Gomez, A. (2018). Developing Metacognitive Skills for Training on Information Security BT - Online Engineering & Internet of Things. In M. E. Auer & D. G. Zutin (Eds.) (pp. 708–720). Cham: Springer International Publishing.

Carson, R. S., & Sheeley, B. J. (2013). 2.5.1 Functional Architecture as the Core of Model-Based Systems Engineering. *INCOSE International Symposium*, *23*(1), 29–45.

Cellucci, T. (2008). *Developing Operational Requirements: A Guide to the Cost-Effective and Efficient Communication of Needs*. U.S. Department of Homeland Security.

Chakraborti, T., Fadnis, K. P., Talamadupula, K., Dholakia, M., Srivastava, B., Kephart, J. O., & Bellamy, R. K. E. (2018). Visualizations for an Explainable Planning Agent. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, {IJCAI-18}* (pp. 5820–5822). International Joint Conferences on Artificial Intelligence Organization.

Challenger, R., Clegg, C. W., & Shepherd, C. (2013). Function allocation in complex systems: Reframing an old problem. *Ergonomics*, *56*(7), 1051–1069.

Champion, M., Rajivan, P., Cooke, N. J., & Jariwala, S. (2012). Team-based cyber defense analysis. *Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*, 218–221.

Checkland, P. (1985). From Optimizing to Learning: A Development of Systems Thinking for the 1990s. *The Journal of the Operational Research Society*, *36*(9), 757–767.

Checkland, P., & Scholes, J. (1990). *Soft Systems Methodology in Action*. New York, NY, USA: John Wiley & Sons, Inc.

Chen, H., & Lynch, K. J. (1992). Automatic construction of networks of concepts characterizing document databases. *IEEE Transactions on Systems, Man, and Cybernetics*, *22*(5), 885–902.

Chen, T. R., Shore, D. B., Zaccaro, S. J., Dalal, R. S., Tetrick, L. E., & Gorab, A. K. (2014). An organizational psychology perspective to examining computer security incident response teams. *IEEE Security & Privacy*, *12*(5), 61–67.

Chuvakin, A., & Barros, A. (2018). *Preparing Your Security Operations for Orchestration and Automation Tools*.

Cisco Systems. (2017). *CISCO 2017 Midyear Cybersecurity Report*. San Jose, CA.

Cisco Systems. (2018). *2018 Annual Cybersecurity Report*. San Jose, CA.

Clark, D., Berson, T., & Lin, H. S. (2014). *At the Nexus of Cybersecurity: Some Basic Concepts and Issues*. Washington, DC, USA: The National Academies Press.

Coats, D. R. (2017). *Worldwide Threat Assessment of the US Intelligence Community*. Washington, DC, USA.

Cobb, S. (2016). Mind This Gap: Criminal Hacking and the Global Cybersecurity Skills Shortage, A Critical Analysis. In *Virus Bulletin Conference*. Virus Bulletin.

Cohen, J. (1960). A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, *20*.

Cohen, M. S., Freeman, J. T., & Wolf, S. (1996). Meta-recognition in time-stressed decision making: {Recognizing}, critiquing, and correcting. *Human Factors*, *38*(2), 206–219.

Coleman, J. (1990). *Foundations of Social Theory* (1st ed.). Cambridge, MA, USA: Harvard University Press.

Collins, H., & Evans, R. (2002). The Third Wave of Science Studies: Studies of Expertise and Experience. *Social Studies of Science*, *32*(2), 235–296.

Collins, H., & Evans, R. (2018). A sociological/philosophical perspective on expertise: The acquisition of expertise through socialization. In K. A. Ericsson, R. R. Hoffman, & A. Kozbelt (Eds.), *Cambridge Handbook of Expertise and Expert Performance* (2nd ed.). Cambridge, UK: Cambridge University Press.

CompTIA. (2018). *Building a Culture of Cybersecurity: A Guide for Corporate Executives and Board Members* (Vol. April). Downers Grove, IL.

Cooke, N. (2004). Measuring Team Knowledge. In *Handbook of Human Factors and Ergonomics Methods* (pp. 46–49). CRC Press.

Core, M. G., Lane, H. C., Lent, M. Van, Gomboc, D., Solomon, S., Rosenberg, M., & Rey, M. (2006). Building Explainable Artificial Intelligence Systems. In *IAAI'06 Proceedings of the 18th conference on Innovative applications of artificial intelligence* (pp. 1766–1773). Boston, MA: AAAI Press, Inc.

Crandall, B., Klein, G., & Hoffman, R. R. (2006). *Working Minds: A Practitioner's Guide to Cognitive Task Analysis*. Cambridge, MA: MIT Press.

Cuevas, H. M., Fiore, S. M., Caldwell, B. S., & Strater, L. (2007). Augmenting team cognition in human-automation teams performing in complex operational environments. *Aviation Space and Environmental Medicine*, *78*(5, Section II), 63–70.

Cuevas, H. M., Fiore, S. M., Salas, E., & Bowers, C. A. (2002). A Macroergonomic Approach to Distributed Team Performance. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *46*(15), 1335–1339.

Curry, S. (2019, February). Mirror Chess is Not Good For Cyber. *Forbes*.

Cusick, J. J., & Ma, G. (2010). Creating an ITIL inspired Incident Management approach: Roots, response, and results. In *2010 IEEE/IFIP Network Operations and Management Symposium Workshops* (pp. 142–148).

Cypress, B. S. (2017). Rigor or Reliability and Validity in Qualitative Research: Perspectives, Strategies, Reconceptualization, and Recommendations. *Dimensions of Critical Care Nursing*, *36*(4), 253–263.

D'Amico, A., O'Brien, B., Whitley, K., Tesone, D., & Roth, E. (2005). Achieving Cyber Defense Situational Awareness: a Cognitive Task Analysis of Information Assurance Analysts. In *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting—2005* (pp. 229–233).

D3 Security. (2019). *Product Guide: Cyber Incident Response*. Vancouver, BC, Canada.

Damelio, R. (2011). *The Basics of Process Mapping* (2nd ed.). Boca Raton, FL: CRC Press.

Darktrace. (2018). *Catching the Silent Attacker, and the Next Phase of Cyber AI*. San Francisco, CA, USA.

Dautenhahn, K., Ogden, B., & Quick, T. (2002). From embodied to socially embedded agents – Implications for interaction-aware robots. *Cognitive Systems Research*, *3*(3), 397–428.

Dawson, J., & Thomson, R. (2018). The future cybersecurity workforce: Going beyond technical skills for successful cyber performance. *Frontiers in Psychology*, *9*.

DCRO Cyber Risk Governance Council. (2018). *GUIDING PRINCIPLES FOR CYBER RISK GOVERNANCE : Principles for Directors in Overseeing Cybersecurity*.

de Visser, E., & Parasuraman, R. (2011). Adaptive Aiding of Human-Robot Teaming: Effects of Imperfect Automation on Performance, Trust, and Workload. *Journal of Cognitive Engineering and Decision Making*, *5*(2), 209–231.

DeGreene, K. (Ed.). (1970). *Systems Psychology*. New York, NY, USA: McGraw-Hill, Inc.

DeLaurentis, D. (2005). Understanding transportation as a system-of-systems design problem. *43rd AIAA Aerospace Sciences Meeting and Exhibit*, (January), 1–14.

Demir, M., McNeese, N. J., & Cooke, N. J. (2016). Team communication behaviors of the human-automation teaming. In *2016 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)* (pp. 28–34).

Demisto. (2018). *Demisto Enterprise for Incident Management*. Cupertino, CA.

Dhingra, K. R., Elms, A., & Hobgood, C. (2010). Reducing error in the emergency department: a call for standardization of the sign-out process. *Annals of Emergency Medicine*, *56*(6), 637–642.

Domínguez, I. X., Goodwin, P. R., Roberts, D. L., & Amant, R. St. (2017). Human Subtlety Proofs: Using Computer Games to Model Cognitive Processes for Cybersecurity. *International Journal of Human–Computer Interaction*, *33*(1), 44–54.

Doran, D., Schulz, S., & Besold, T. (2017). *What Does Explainable AI Really Mean? A New Conceptualization of Perspectives*.

Dunne, E. (2018, October 10). The US military's cybersecurity is extraordinarily weak. *Washington Examiner*.

Eldardiry, O., & Caldwell, B. (2015). Improving Information and Task Coordination in Cyber Security Operation Centers. *IIE Annual Conference. Proceedings*, 1224–1233.

Emes, M. R., Bryant, P. A., Wilkinson, M. K., King, P., James, A. M., & Arnold, S. (2012). Interpreting "systems architecting." *Systems Engineering*, *15*(4), 369–395.

Endsley, M. R. (2018). Expertise and Situation Awareness. In K. A. Ericsson, R. R. Hoffman, & A. Kozbelt (Eds.), *Cambridge Handbook of Expertise and Expert Performance* (2nd ed.). Cambridge, UK: Cambridge University Press.

Endsley, M. R., & Connors, E. S. (2014). Cyber Defense and Situational Awareness. In A. Kott, C. Wang, & R. F. Erbacher (Eds.), *Cyber Defense and Situational Awareness* (Vol. 62, pp. 7–27). Springer International Publishing.

Engelbrecht, S. (2018, July 27). The Evolution of SOAR Platforms. *SecurityWeek by Wired Business Media*.

Ericsson, K. A., & Smith, J. (1991). Prospects and limits to the empirical study of expertise: an introduction. In K. A. Ericsson & J. Smith (Eds.), *Toward a General Theory of Expertise* (1st ed., pp. 1–38). Cambridge: Cambridge University Pres.

Ertmer, P., & Newby, T. (1996). The expert learner: Strategic, self-regulated, and reflective. *Instructional Science*, *24*(1), 1–24.

Espejo, R., & Reyes, A. (2011). The Viable System Model: Effective Strategies to Manage Complexity BT - Organizational Systems: Managing Complexity with the Viable System Model. In R. Espejo & A. Reyes (Eds.) (pp. 91–112). Berlin, Heidelberg: Springer Berlin Heidelberg.

EY. (2018). *Is cybersecurity about more than protection?: EY Global Information Security Survey 2018-2019*.

Faysel, M. A., & Haque, S. S. (2010). Towards cyber defense: research in intrusion detection and intrusion prevention systems. *IJCSNS International Journal of Computer Science and Network Security*, *10*(7), 316–325.

Feigh, K. M., Dorneich, M. C., & Hayes, C. C. (2012). Toward a Characterization of Adaptive Systems. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *54*(6), 1008–1024.

Feigh, K. M., & Pritchett, A. R. (2014). Requirements for Effective Function Allocation. *Journal of Cognitive Engineering and Decision Making*, *8*(1), 23–32.

FEMA. (n.d.). The Four Phases of Emergency Management. In *Animals in Disasters* (p. A-3-1-A-3-19).

Filkins, B. (2019). *2019 SANS Automation & Integration Survey*.

FireEye. (2019). *M-Trends 2019: Special Report*. Malpitas, CA, USA.

Fitts, P. M. (Ed.). (1951). *Human engineering for an effective air-navigation and traffic-control system. Human engineering for an effective air-navigation and traffic-control system.* Oxford, England: National Research Council.

Flood, R. L., & Jackson, M. C. (1991). *Creative problem solving : total systems intervention*. Chichester ; New York: Wiley.

Foroushani, V. (n.d.). *Saving Valuable Time with Incident Response Management*. Vancouver, BC, Canada.

Franke, U., & Brynielsson, J. (2014). Cyber situational awareness - A systematic review of the literature. *Computers & Security*, *46*.

Freed, S. E. (2014). *Examination of Personality Characteristics among Cybersecurity and Information Technology Professionals*. University of Tennesee at Chattanooga.

Frey, B. B. (2016). *There's a stat for that! : what to do & when to do it*. SAGE Publications, Inc.

Frost & Sullivan. (2017). *2017 Global Information Security Workforce Study: Benchmarking Workforce Capacity and Response to Cyber Risk*.

Fusch, P. I., & Ness, L. R. (2015). Are We There Yet? Data Saturation in Qualitative Research. *The Qualitative Report*, *20*(9), 1408–1416.

Gaggiotti, H., Kostera, M., & Krzyworzeka, P. (2017). More than a method? Organisational ethnography as a way of imagining the social. *Culture and Organization*, *23*(5), 325–340.

Gamal, M. M., Hassan, B., & Hegazy, A. F. (2011). A Security Analysis Framework Powered by an Expert System. *International Journal of Computer Science and Security (IJCSS)*, *4*(6).

Garrett, S. K., & Caldwell, B. S. (2002). Mission Control Knowledge Synchronization: Operations to Reference Performance Cycles. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *46*(15), 1345–1349.

Garrett, S. K., & Caldwell, B. S. (2011). Describing functional requirements for knowledge sharing communities. *Behavior and Information Technology*, *21*(5), 359–364.

Garrett, S. K., Caldwell, B. S., Harris, E. C., & Gonzalez, M. C. (2009). Six dimensions of expertise: a more comprehensive definition of cognitive expertise for team coordination. *Theoretical Issues in Ergonomics Science*, *10*(2), 93–105.

Gartner. (2018). Gartner Forecasts Worldwide Information Security Spending to Exceed $124 Billion in 2019. Sydney, AUS: Gartner, Inc.

Gilbreth, F., & Gilbreth, L. (1922). Process charts and their place in management. *Mechanical Engineering*, *44*.

Gillham, B. (2008). *Observation techniques : structured to unstructured*. London ; New York: Continuum International Pub.

Glaser, R., & Chi, M. T. H. (1988). Overview. In M. T. H. Chi, R. Glaser, & M. J. Farr (Eds.), *The nature of expertise*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Gokhale, G. B., & Banks, D. A. (2004). Organisational Information Security: A Viable System Perspective. In *2nd Australian Information Security Management Conference* (pp. 178–184). Perth, Australia: School of Computer & Information Science Edith Cowan University.

Goldstein, A. (2016). Components of a multi-perspective modeling method for designing and managing IT security systems. *Information Systems & E-Business Management*, *14*(1), 101–141.

Gomboc, D., Solomon, S., Core, M. G., Lane, H. C., & Lent, M. Van. (2005). Design recommendations to support automated explanation and tutoring. In *In Proceedings of the 14th Conference on Behavior Representation in Modeling and Simulation (BRIMS05), Universal*.

Gonzalez, C., Ben-Asher, N., Oltramari, A., & Lebiere, C. (2014). Cognition and Technology. In A. Kott, C. Wang, & R. F. Erbacher (Eds.), *Cyber Defense and Situational Awareness* (pp. 93–117). Cham: Springer International Publishing.

Goodall, J. R., D'Amico, A., & Kopylec, J. K. (2009). Camus: Automatically mapping cyber assets to missions and users. *Proceedings - IEEE Military Communications Conference MILCOM*, 1–7.

Goodell, L. S., Stage, V. C., & Cooke, N. K. (2016). Report Practical Qualitative Research Strategies : Training Interviewers and Coders. *Journal of Nutrition Education and Behavior*, *48*(8), 578–585.

Goulding, C. (2002). Grounded Theory. London: SAGE Publications Ltd.

Graham, B. B. (2004). *Detail process charting : speaking the language of process*. Hoboken, NJ: John Wiley & Sons. Inc.

Graves, J. (2019, February). Reactive vs. Proactive Cybersecurity: 5 Reasons Why Traditional Security No Longer Works. *Fortinet*.

Green, H. E. (2014). Use of theoretical and conceptual frameworks in qualitative research.(Report), *21*(6), 34.

Grose, T. (2007). "Patch and Pray." *ASEE Prism*, *17*(3), 27–32.

Growing the Security Analyst: Hiring, Training, and Retention. (2014). (No. 4AA5- 3982ENN).

Guariniello, C., O'Neill, W., Ukai, T., Dumbacher, D., Caldwell, B., & DeLaurentis, D. (2016). Understanding human space exploration. In *Proceedings of the International Astronautical Congress, IAC*. International Astronautical Federation, IAF.

Guinery, J. (2011). Capturing decision input content to support work system design. In *Tenth International Symposium on Human Factors in Organizational Design and Management.* Grahamstown, South Africa: IEA Press.

Gunning, D. (2017). Explainable Artificial Intelligence (XAI). DARPA.

Gutzwiller, R. S., Fugate, S., Sawyer, B. D., & Hancock, P. A. (2015a). The human factors of cyber network defense. *Proceedings of the Human Factors and Ergonomics Society*, *2015-Janua*, 322–326.

Gutzwiller, R. S., Fugate, S., Sawyer, B. D., & Hancock, P. A. (2015b). The Human Factors of Cyber Network Defense. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *59*(1), 322–326.

Hale, R. (2017). Cyber Competencies and the Cybersecurity Officer. In D. Antonucci (Ed.), *The Cyber Risk Handbook: Creating and Measuring Cybersecurity Capabilities* (pp. 359–368). John Wiley & Sons, Inc.

Hallgren, K. A. (2012). Computing Inter-Rater Reliability for Observational Data: An Overview and Tutorial. *Tutor Quant Methods Psychol.*, *8*(1), 23–34.

Hancock, P. A., & Scallen, S. F. (1996). The Future of Function Allocation. *Ergonomics in Design*, *4*(4), 24–29.

Handy, C. (1995). Trust and the Virtual Organization. *Harvard Business Review*.

Hartson, R., & Pyla, P. (2012). *The UX Book: Process and Guidelines for Ensuring a Quality User Experience* (1st ed.). San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.

HCL Technologies. (2019). *State of Cybersecurity 2019: Current Trends in Workforce Development*. Schaumburg, IL.

Healey, C. G., Hao, L., & Hutchinson, S. E. (2014). Visualizations and Analysts. In A. Kott, C. Wang, & R. F. Erbacher (Eds.), *Cyber Defense and Situational Awareness* (pp. 145–165). Cham: Springer International Publishing.

Herring, C., & Kaplan, S. (2000). The viable system model for software. In *4th World Multiconference on Systemics, Cybernetics and Informatics (SCI'2000)*. Orlando, FL, USA.

Hilburn, B. G., Molloy, R., Wong, D., & Parasuraman, R. (1993). Operator versus computer control of adaptive automation. *International Symposium on Aviation Psychology*, (October 2016), 161–166.

Hoffman, L., Burley, D., & Toregas, C. (2012). Holistically Building the Cybersecurity Workforce. *IEEE Security & Privacy*.

Hoffman, R. R. (1987). The ‘ Problem of Extracting the Knowledge of Experts from the Perspective of Experimental Psychology. *AI Magazine*, *8*(2), 53–67.

Hoffman, R. R. (2008). Influencing versus Informing Design, Part 2: Macrocognitive Modeling. *Intelligent Systems, IEEE*, *23*(6).

Hoffman, R. R., & Deal, S. V. (2008). Influencing versus Informing Design, Part 1: A Gap Analysis. *Intelligent Systems, IEEE*, *23*(5).

Hoffman, R. R., Feltovich, P. J., & Ford, K. M. (1997). A general framework for conceiving of expertise and expert systems in context. In P. J. Feltovich, K. M. Ford, & R. R. Hoffman (Eds.), *Expertise in context : human and machine* (pp. 543–580). Cambridge, MA: MIT Press.

Hoffman, R. R., Klein, G., & Mueller, S. T. (2018). Explaining Explanation For " Explainable AI ." In *Proceedings of the Human Factors and Ergonomics Society 2018 Annual Meeting* (pp. 197–201). Los Angeles, CA: SAGE Publications.

Hoffman, R. R., Mueller, S., Klein, G., & Litman, J. (2018). *Metrics for Explainable AI: Challenges and Prospects*.

Hollnagel, E., & Bye, A. (2000). Principles for modelling function allocation. *International Journal of Human-Computer Studies*, *52*(2), 253–265.

Holtel, S. (2015, February). Cognitive tools turn the rules upside down. (Cover story). *KM World*, *24*(2), 1–18.

Holtzblatt, K. (2016). *Contextual design: design for life* (2nd ed.). Elsevier.

Holtzblatt, K., & Jones, S. (1993). Contextual Inquiry: A Participatory Technique for System Design. In D. Schuler & A. Namioka (Eds.), *Participatory design: Principles and practices* (pp. 177–210). Hillsdale, NJ: Lawrence Erlbaum Associates.

Horne, B. (2014). On Computer Security Incident Response Teams. *IEEE Security & Privacy*, *12*(5), 13–15.

Hoverstadt, P. (2010). The Viable System Model. In M. Reynolds & S. Holwell (Eds.), *Systems Approaches to Managing Change: A Practical Guide* (pp. 87–133).

Humphrey, C. M., & Adams, J. A. (2011). Analysis of complex team-based systems: augmentations to goal-directed task analysis and cognitive work analysis. *Theoretical Issues in Ergonomics Science*, *12*(2), 149–175.

IBM. (2017). *Resilient Incident Response Platform*. Somers, NY.

Identity Theft Resource Center. (2019). Multi-Year Breach Summary. Retrieved November 7, 2019, from https://www.idtheftcenter.org/wp-content/uploads/2019/02/Multi-Year-Chart.pdf

IEEE Guide for Information Technology - System Definition - Concept of Operations (ConOps) Document (1362-1998). (1998). *IEEE Guide for Information Technology - System Definition - Concept of Operations (ConOps) Document*. Piscataway, USA: IEEE.

Inaba, M., & Takahashi, N. (2017). The use of reputation in repeated dyadic interactions. *Rationality and Society*, *30*(1), 54–79.

ISACA Cybersecurity Nexus. (2017). *State of Cyber Security 2017 Part 2: Current Trends in the Threat Landscape*. Rolling Meadows, IL, USA.

ISO/IEC/IEEE 29148:2011. (2011). Geneva, Switzerland.

ISO/IEC/IEEE Std 31320-1:2012. (2012). Geneva, Switzerland.

ISO/IEC. (2011). *Systems and software engineering - Life cycle processes - Requirements engineering* (No. ISO/IEC 29148). Geneva, Switzerland.

Jackson, M. C. (2000). *Systems approaches to management*. New York: Kluwer Academic/Plenum.

Jackson, M. C. (2003). *Systems thinking : creative holism for managers*. Chichester: J. Willey.

Jackson, P., & Klobas, J. (2008). Transactive memory systems in organizations: Implications for knowledge directories. *Decision Support Systems*, *44*(2), 409–424.

Jacobs, R., & Washington, C. (2003). Employee development and organizational performance: a review of literature and directions for future research. *Human Resource Development International*, *6*(3), 343–354.

Johnson, M., Bradshaw, J. M., & Feltovich, P. J. (2017). Tomorrow's Human–Machine Design Tools: From Levels of Automation to Interdependencies. *Journal of Cognitive Engineering and Decision Making*, *12*(1), 77–82.

Jøsok, Ø., Lugo, R., Knox, B. J., Sütterlin, S., & Helkala, K. (2019). Self-Regulation and Cognitive Agility in Cyber Operations. *Frontiers in Psychology*, *10*, 875.

Kammüller, F. (2018). Human Centric Security and Privacy for the IoT Using Formal Techniques. In D. Nicholson (Ed.), *Advances in Human Factors in Cybersecurity: Proceedings of the AHFE 2017 International Conference on Human Factors in Cybersecurity, July 17−21, 2017, The Westin Bonaventure Hotel, Los Angeles, California, USA* (pp. 106–116). Cham: Springer International Publishing.

Kerlinger, F. N., & Lee, H. B. (2000). *Foundations of Behavioral Research* (4th ed.). Fort Worth, TX: Harcourt College Publishers.

Killcreece, G., Kossakowski, K.-P., Ruefle, R., & Zajicek, M. (2003). Organizational Models for Computer Security Incident Response Teams (CSIRTs). Pittsburgh, PA: Carnegie Mellon Software Engineering Institute.

Klein, G. A., & Hoffman, R. R. (1993). Seeing the invisible: {Perceptual}/cognitive aspects of expertise. In M. Rabinowitz (Ed.), *Cognitive science foundations of instruction* (pp. 203–226). Hillsdale, NJ: Lawrence Erlbaum Associates.

Klinger, D., & Hahn, B. (2004). Team Decision Requirement Exercise. In *Handbook of Human Factors and Ergonomics Methods* (pp. 52–55). CRC Press.

Knox, W. B., Stone, P., & Breazeal, C. (2013). Training a Robot via Human Feedback: A Case Study. In *Proceedings of the 5th International Conference on Social Robotics - Volume 8239* (pp. 460–470). Berlin, Heidelberg: Springer-Verlag.

Kossiakoff, A., & Sweet, W. N. (2003). *Systems Engineering: Principles and Practice*. Hoboken, NJ: John Wiley & Sons, Inc.

Koulouris, T., Mont, M. C., & Arnell, S. (2017). *SDN4S: Software Defined Networking for Security* (No. HPE-2017-07).

KPMG International. (2015). *Cyber security : a failure of imagination by CEOs Global CEOs*.

Kral, R. (2018, June 11). Governing Cybersecurity : Cybersecurity Committees On The Rise A Growing Trend of Cybersecurity Committees. *Corporate Compliance Insights*.

Krippendorff, K. (1980). *Content analysis : an introduction to its methodology*. Beverly Hills: Sage Publications.

Kumar, V. (2005). Parallel and Distributed Computing for Cybersecurity. *IEEE Distributed Systems Online*, *6*(10), 1–9. https://doi.org/10.1109/MDSO.2005.53

Kurke, M. I. (1961). Operational Sequence Diagrams in System Design. *Human Factors*, *3*(1), 66–73.

Kurowski, S., & Frings, S. (2011). Computational Documentation of IT Incidents as Support for Forensic Operations. In *2011 Sixth International Conference on IT Security Incident Management and IT Forensics* (pp. 37–47).

Landis, J. R., & Koch, G. G. (1977). The Measurement of Observer Agreement for Categorical Data. *Biometrics*, *33*(1), 159–174.

Langford, G., Franck, R. E., Huynh, T., & Lewis, I. A. (2008). Gap Analysis: Rethinking the Conceptual Foundations. In *Fifth Annual Acquisition Research Symposium*. Monterrey, CA: Naval Postgraduate School.

Lathrop, S. D. (2017). Interacting with Synthetic Teammates in Cyberspace. In D. Nicholson (Ed.), *Advances in Human Factors in Cybersecurity: Proceedings of the AHFE 2017 International Conference on Human Factors in Cybersecurity, July 17−21, 2017, The Westin Bonaventure Hotel, Los Angeles, California, USA* (pp. 133–145). Springer International Publishing.

Lee, C.-H., Soong, F. K., & Paliwal, K. K. (Eds.). (1996). *Automatic speech and speaker recognition: advanced topics*. Norwell, MA, USA: Kluwer Academic Publishers.

Lee, Y., & Lee, S.-J. (2004). An Exploratory Investigation fo Factors Affecting Computer Security Incident Response Team Performance. In *AMCIS 2004 Proceedings* (pp. 4402–4406). AIS Electronic Library.

Lehto, M. R., & Buck, J. (2008). Design to Fit Tasks, Processes, and People. In *Introduction to human factors and ergonomics for engineers*. New York: Lawrence Erlbaum.

Leung, L. (2015). Validity, reliability, and generalizability in qualitative research. *Journal of Family Medicine and Primary Care*, *4*(3), 324–327.

Levis, A. H., & Wagenhals, L. W. (2000). C4ISR architectures: I. Developing a process for C4ISR architecture design. *Systems Engineering*, *3*(4), 225–247.

Lin, C., Cao, N., Liu, S. X., Papadimitriou, S., Sun, J., & Yan, X. (2009). SmallBlue: Social Network Analysis for Expertise Search and Collective Intelligence. In *2009 IEEE 25th International Conference on Data Engineering* (pp. 1483–1486).

Lincoln, Y. S., & Guba, E. G. (2011). The Sage handbook of qualitative research. In N. K. Denzin & Y. S. Lincoln (Eds.), *The Sage handbook of qualitative research* (4th ed.., pp. 163–187). Thousand Oaks: Thousand Oaks : Sage.

LogRhythm. (2019). *2018 Cybersecurity: Perceptions & Practices Table of Contents*. Boulder, CO.

London, J., Caldwell, B., & Patsavas, K. (2013). Aligning Learning Outcomes and the Engineering Educational Accreditation Expectations. *IIE Annual Conference. Proceedings*, 512–521.

Lowe, A., Norris, A. C., Farris, A. J., & Babbage, D. R. (2018). Quantifying Thematic Saturation in Qualitative Data Analysis. *Field Methods*, *30*(3), 191–207.

Maier, M. W., & Rechtin, E. (2000). *The Art of Systems Architecting* (2nd ed.). Boca Raton, FL: CRC Press.

Mancuso, V. (2012). *An Interdisciplinary Evaluation of Transactive Memory in Distributed Cyber Teams*.

Mancuso, V., Minotra, D., Giacobe, N., McNeese, M., & Tyworth, M. (2012). idsNETS: An experimental platform to study situation awareness for intrusion detection analysts. In *2012 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support* (pp. 73–79).

Mancuso, V., Staheli, D., Leahy, M. J., & Kalke, M. M. (2016). Cloudbreak: Answering the Challenges of Cyber Command and Control. *Lincoln Laboratory Journal*, *22*(1).

Massey, L., Seker, R., & Nicholson, D. (2018). Feasibility of Leveraging an Adaptive Presentation Layer for Cyber Security Visualizations. In D. Nicholson (Ed.), *Advances in Human Factors in Cybersecurity: Proceedings of the AHFE 2017 International Conference on Human Factors in Cybersecurity, July 17−21, 2017, The Westin Bonaventure Hotel, Los Angeles, California, USA* (pp. 117–129). Cham: Springer International Publishing.

Mayer, R. C., Davis, J. H., Schoorman, F. D., Mayer, R. C., & Davis, J. H. (1995). An Integrative Model of Organizational Trust. *The Academy of Management Review*, *20*(3), 709–734.

Maymí, F. J., & Thomson, R. (2018). Human-Machine Teaming and Cyberspace BT - Augmented Cognition: Intelligent Technologies. In D. D. Schmorrow & C. M. Fidopiastis (Eds.) (pp. 299–315). Cham: Springer International Publishing.

McGrath, J. E. (1984). *Groups: Interaction and Performance*. Englewood Cliffs, NJ: Prentice-Hall, Inc.

Mclaughlin, N., Rodstein, J., Burke, M., & Martin, N. (2014). Demystifying Process Mapping: A Key Step in Neurosurgical Quality Improvement Initiatives. *Neurosurgery*, *75*(2), 99–109.

Meadows, D. (2008). *Thinking in Systems: A Primer*. (D. Wright, Ed.). Sterling: Earthscan.

Memon, M. (2014). Security modeling for service-oriented systems using security pattern refinement approach. *Software & Systems Modeling*, *13*(2), 549–573.

Mercado, J. E., Rupp, M. A., Chen, J. Y. C., Barnes, M. J., Barber, D., & Procci, K. (2016). Intelligent Agent Transparency in Human–Agent Teaming for Multi-UxV Management. *Human Factors*, *58*(3), 401–415.

Merriam, S. B., & Tisdell, E. (2016). *Qualitative research : a guide to design and implementation* (Fourth edi). Jossey-Bass.

Mesmer-Magnus, J. R., & DeChurch, L. A. (2009). Information sharing and team performance: A meta-analysis. *Journal of Applied Psychology*, *94*(2), 535–546.

Mihajlov, M., & Jerman-Blazic, B. (2018). Eye Tracking Graphical Passwords. In D. Nicholson (Ed.), *Advances in Human Factors in Cybersecurity: Proceedings of the AHFE 2017 International Conference on Human Factors in Cybersecurity, July 17−21, 2017, The Westin Bonaventure Hotel, Los Angeles, California, USA* (pp. 37–44). Cham: Springer International Publishing.

Militello, L. G., Hutton, R. J. B., Pliske, R. M., Knight, B. J., Klein, G., & Randel, J. (1997). ACTA Methodology. *Navy Personnel Research and Development Center*, (November), 1–59.

Miller, C. A., & Parasuraman, R. (2007). Designing for Flexible Interaction Between Humans and Automation: Delegation Interfaces for Supervisory Control. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *49*(1), 57–75.

Mineraud, J., Mazhelis, O., Su, X., & Tarkoma, S. (2016). A gap analysis of Internet-of-Things

Morgan, S. (2015). IBM’s CEO On Hackers: “Cyber Crime Is The Greatest Threat To Every Company In The World.” *Forbes*.

Morgan, S. (2017). *Cybersecurity Jobs Report: 2017 Edition*. Toronto, ON, Canada.

Muggler, M., Eshwarappa, R., & Cankaya, E. C. (2018). Cybersecurity Management Through Logging Analytics. In D. Nicholson (Ed.), *Advances in Human Factors in Cybersecurity: Proceedings of the AHFE 2017 International Conference on Human Factors in Cybersecurity, July 17−21, 2017, The Westin Bonaventure Hotel, Los Angeles, California, USA* (pp. 3–15). Cham: Springer International Publishing.

NASA. (2017). Technology Readiness Level. Retrieved from https://www.nasa.gov/directorates/heo/scan/engineering/technology/txt_accordion1.html

NASA Systems Engineering Handbook. (2007). Washington, DC, USA: National Aeronautics and Space Administration.

National Academies of Sciences Engineering and Medicine. (2017). *Foundational Cybersecurity Research*. (L. I. Millett, B. Fischhoff, & P. J. Weinberger, Eds.). Washington, DC, USA: National Academies Press.

National Initiative for Cybersecurity Careers and Studies. (2017). NICE Cybersecurity Workforce Framework. Retrieved January 1, 2017, from https://niccs.us-cert.gov/workforce-development/cyber-security-workforce-framework

NCSC-NL. (2015). *CSIRT Maturity Toolkit*. The Hague, The Netherlands.

Neiva, C., Lawson, C., Bussa, T., & Sadowski, G. (2017). *Innovation Insight for Security Orchestration, Automation and Response (SOAR)* (No. G00338719).

Newhouse, W., Keith, S., Scribner, B., & Witte, G. (2017). *National Initiative for Cybersecurity Education (NICE) Cybersecurity Workforce Framework*.

Newman, L. H. (2018, May). The Bleak State of Federal Government Cybersecurity. *Wired*.

NICCS. (2017). Glossary - NATIONAL INITIATIVE FOR CYBERSECURITY CAREERS AND STUDIES. Washington, DC, USA: U.S. Department of Homeland Security.

Nielsen, J. (1994). Enhancing the Explanatory Power of Usability Heuristics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 152–158). New York, NY, USA: ACM.

Nielsen, J., & Landauer, T. K. (1993). A Mathematical Model of the Finding of Usability Problems. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems* (pp. 206–213). New York, NY, USA: ACM.

Noble, H., & Smith, J. (2015). Issues of validity and reliability in qualitative research. *Evidence Based Nursing*, *18*(2), 34 LP – 35.

Nyre, M. M. (2016). *Developing agent-based simulation models of task performance of cognitively diverse teams*. Purdue University.

Office of the Deputy Under Secretary of Defense for Acquisition and Technology: Systems and Software Engineering. (2008). *Defense Acquisition Program Support (DAPS) Methodology*. Washington, DC, USA.

Oltsik, J. (2017). *The Life and Times of Cybersecurity Professionals: A Cooperative Research Project by ESG and ISSA.*

Oltsik, J. (2018a, May 9). The evolution of security operations, automation and orchestration. *CSO by IDG Communications*.

Oltsik, J. (2018b, May 30). The rise of analyst-centric security operations technologies. *CSO Online*.

Oltsik, J. (2019, January 10). The cybersecurity skills shortage is getting worse. *CSO Online*.

Onken, J. D. (2012). *Modeling Real-Time Coordination of Distrubuted Expertise and Event Response in NASA Mission Control Center Operations*. Purdue University.

Onken, R. (2003). *Cognitive Cooperation for the Sake of the Human-Machine Team Effectiveness*.

Optiv. (2019). *Enterprise Attitudes to Cybersecurity: Tackling the Modern Threat Landscape in the United Kingdom*. Denver, CO.

Oyewole, T. (2016). Application fo Situation Awareness in Incident Response. *ISACA Journal*, *3*, 1–5.

Palazzolo, E. T. (2005). Organizing for Information Retrieval in Transactive Memory Systems. *Communication Research*, *32*(6), 726–761.

Papadopoulous, L. (2017). How Watson AI is helping companies stay ahead of hackers and cybersecurity attacks. Retrieved January 1, 2017, from https://www.ibm.com/blogs/watson/2017/08/how-watson-ai-is-helping-companies-stay-ahead-of-cybersecurity-attacks/

Parasuraman, R, Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*.

Parasuraman, Raja, Barnes, M., & Cosenzo, K. (2007). Adaptive automation for human-robot teaming in future command and control systems. *The International C2 Journal*, *1*(2), 43–68.

Paris, C. R., Salas, E., & Cannon-Bowers, J. A. (2000). Teamwork in multi-person systems: a review and analysis. *Ergonomics*, *43*(8), 1052–1075.

Patton, M. Q. (1985). Quality in qualitative research: Methodological principles and recent developments. In *Invited address to Division J of the American Educational Research Association*. Chicago.

Paulheim, H. (2016). Knowledge graph refinement: A survey of approaches and evaluation methods. *Semantic Web*, *1*, 489–508.

Pearson, G. (n.d.). *Speech Recognition: Defining System Requirements*.

Pelto, P. J. (2016). Applied ethnography: Guidelines for field research. In *Applied ethnography: Guidelines for field research*. Routledge.

Pesanka, D. A., Greenhouse, P. K., Rack, L. L., Delucia, G. A., Perret, R. W., Scholle, C. C., … Janov, C. L. (2009). Ticket to Ride: Reducing Handoff Risk During Hospital Patient Transport. *Journal of Nursing Care Quality*, *24*(2).

Petersen, C., & Lentz, R. (2015). *Surfacing Critical Cyber Threats Through Security Intelligence: A Reference Model for IT Security Practitioners*.

Peusquens, R. (2017). CSP, not 007: Integrated Cybersecurity Skills Training BT - Cyber Security. Simply. Make it Happen.: Leveraging Digitization Through IT Security. In F. Abolhassan (Ed.) (pp. 71–74). Cham: Springer International Publishing.

Piasecki, A. M., Fendley, M. E., & Warren, R. (2017). Improving Anomaly Detection Through Identification of Physiological Signatures of Unconscious Awareness. In I. L. Nunes (Ed.), *AHFE* (pp. 259–269). Cham: Springer International Publishing.

Plenary Session 5: Directors' Perspectives on National Intelligence. (2017). In *Intelligence & National Security Summit*. Washington, DC, USA.

Plumptre, I., Mulki, O., Granados, A., Gayle, C., Ahmed, S., Jenny, N. L., & Fernando, H. (2017). Standardizing bimanual vaginal examination using cognitive task analysis, (August), 114–119.

Polson, P. G., Lewis, C., Rieman, J., & Wharton, C. (1992). Cognitive walkthroughs: a method for theory-based evaluation of user interfaces. *International Journal of Man-Machine Studies*, *36*(5), 741–773.

Pomerleau, M. (2016). Pentagon research chief: AI is powerful but has critical limitations. *Defense Systems*.

Ponemon Institute. (2017). *2017 Cost of Data Breach Study*. Traverse City, MI.

Ponemon Institute. (2019). *Staffing the IT Security Function in the Age of Automation: A Study of Organizations in the United States, United Kingdom and APAC*.

Price, H. E. (1985). The Allocation of Functions in Systems. *Human Factors*, *27*(1), 33–45.

Pritchett, A. R., Kim, S. Y., & Feigh, K. M. (2014). Measuring Human-Automation Function Allocation. *Journal of Cognitive Engineering and Decision Making*, *8*(1), 33–51.

Proctor, R. W. (2015, August). The role of human factors/ergonomics in the science of security: Decision making and action selection in cyberspace. *Human Factors*.

Pujara, J., Miao, H., Getoor, L., & Cohen, W. (2013). Knowledge Graph Identification. In H. Alani, L. Kagal, A. Fokoue, P. Groth, C. Biemann, J. X. Parreira, … K. Janowicz (Eds.), *The Semantic Web – ISWC 2013* (pp. 542–557). Berlin, Heidelberg: Springer Berlin Heidelberg.

Rajivan, P., & Cooke, N. J. (2018). Information-Pooling Bias in Collaborative Security Incident Correlation Analysis. *Human Factors*, *60*(5), 626–639.

Ravindranath, M. (2015, September). Director: DARPA's Cyber Strategy So Far Has Been "Patch and Pray." *Nextgov*.

Read, E. K. (2013). Using Cognitive Task Analysis to Create a Teaching Protocol for Bovine Dystocia, *40*(4), 397–401.

Reagans, R., & McEvily, B. (2003). Network Structure and Knowledge Transfer: The Effects of Cohesion and Range. *Administrative Science Quarterly*, *48*(2), 240–267.

Reed, G., Philip, P. A., Barchowsky, A., Lippert, C. J., & Sparacino, A. R. (2010). Sample survey of smart grid approaches and technology gap analysis.

Reed, T., Abbott, R. G., Anderson, B., Nauer, K., & Forsythe, C. (2014). Simulation of Workflow and Threat Characteristics for Cyber Security Incident Response Teams. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *58*(1), 427–431.

Rhee, H.-S., Ryu, Y., & Kim, C.-T. (2012). Unrealistic optimism on information security management. *Computers & Security*, *31*(2).

Richmond, C., & Lindstrom, P. (2015). *IDC Security Survey: As the Job Churns*.

Roberts, R., Flin, R., & Cleland, J. (2016). Journal of Loss Prevention in the Process Industries How to recognise a kick : A cognitive task analysis of drillers ' situation awareness during well operations. *Journal of Loss Prevention in the Process Industries*, *43*, 503–513.

Romero, D., Bernus, P., Noran, O., Stahre, J., & Fast-berglund, Å. (2016). The Operator 4.0 : Human Cyber-Physical Systems & Adaptive Automation towards Human-Automation Symbiosis Work Systems. In *IFIP international conference on advances in production management systems* (pp. 677–686). Springer.

Roy, S., Ellis, C., Shiva, S., Dasgupta, D., Shandilya, V., & Wu, Q. (2010). A Survey of Game Theory as Applied to Network Security. In *2010 43rd Hawaii International Conference on System Sciences* (pp. 1–10).

Ruefle, R., Dorofee, A., Mundie, D., Householder, A. D., Murray, M., & Perl, S. J. (2014). Computer Security Incident Response Team Development and Evolution. *IEEE Security & Privacy*, *12*(5), 16–26.

Rutkin, A. (2015). Hybrid dream team. *New Scientist*, *227*(3034), 20.

Saldana, J. (2009). *The coding manual for qualitative researchers*. Los Angeles, [Calif.] ; London: SAGE.

Sasaki, M., Abe, S., Takei, S., Kameda, Y., Ichikawa, T., Momoi, Y., … Nagao, T. (2017). *Six Ws on cybersecurity information sharing for enhancing SOC/CSIRT*.

Saunders, B., Sim, J., Kingstone, T., Baker, S., Waterfield, J., & Bartlam, B. (2018). Saturation in qualitative research: exploring its conceptualization and operationalization. *Quality & Quantity*, *52*(4), 1893–1907.

Scale Venture Partners. (2017). *The State of Cybersecurity Priorities and Strategies 2017*.

Scallen, S. F., & Hancock, P. A. (2001). Implementing Adaptive Function Allocation. *The International Journal of Aviation Psychology*, *11*(2), 197–221.

Scardamalia, M., & Bereiter, C. (1991). Literate Expertise. In K. A. Ericsson & J. Smith (Eds.), *Toward a General Theory of Expertise* (pp. 172–194). New York: Cambridge University Pres.

Scharre, P. D. (2003). The Opportunity & Challenge of Autonomous Systems. In A. P. Williams & P. D. Scharre (Eds.), *Autonomous Systems: Issues for Defence Policy Makers* (pp. 3–26). Norfolk, VA: NATO Communications and Information Agency.

Schilberg, D., & Schmitz, S. (2017). Information Model for Intention-Based Robot-Human Interaction BT - Advances in Human Factors and System Interactions. In I. L. Nunes (Ed.), *AHFE* (pp. 129–139). Cham: Springer International Publishing.

Schultz, E. E. (2012). Human Factors and Information Security. In G. Salvendy (Ed.), *Handbook*

Schwartzman, H. B. (1993). *Ethnography in organizations* (Vol. 27). Sage.

SEBoK. (2017). *Functional Architecture (glossary)* (v1.8).

Session 7: Humans and Machines Working Together with Big Data. (2017). In *Challenges in Machine Generation of Analytic Products from Multi-Source Data: Proceedings of a Workshop* (pp. 28–30). Washington, DC, USA: National Academies Press.

Session 8: Use of Machine Learning for Privacy Ethics. (2017). In *Challenges in Machine Generation of Analytic Products from Multi-Source Data: Proceedings of a Workshop* (pp. 31–33). Washington, DC, USA: National Academies Press.

Sheridan, T. B. (1992). *Telerobotics, Automation, and Human Supervisory Control*. Cambridge, MA, USA: MIT Press.

Sheridan, T. B. (2016). Human – Robot Interaction: Status and Challenges. *Human Factors*, *58*(4), 525–532.

Sheridan, T. B., & Ferrell, W. R. (1974). *Man-machine systems; Information, control, and decision models of human performance. Man-machine systems; Information, control, and decision models of human performance.* Cambridge, MA, US: The MIT Press.

Sheridan, T. B., & Verplank, W. L. (1978). *Human and computer control of undersea teleoperators*. Cambridge, MA, US.

Shively, R. J., Lachter, J., Koteskey, R., & Brandt, S. L. (2018). Crew Resource Management for Automated Teammates (CRM-A) BT - Engineering Psychology and Cognitive Ergonomics. In D. Harris (Ed.) (pp. 215–229). Cham: Springer International Publishing.

Shoshitaishvili, Y., Weissbacher, M., Dresel, L., Salls, C., Wang, R., Kruegel, C., & Vigna, G. (2017). Rise of the HaCRS: Augmenting Autonomous Cyber Reasoning Systems with Human Assistance. *CoRR*.

Siemplify. (2018). *The Business Case for SOAR*. New York, NY.

Siemplify. (2019). *Siemplify Product Overview*. New York, NY.

Silva, A., Emmanuel, G., McClain, J. T., Matzen, L., & Forsythe, C. (2015). Measuring Expert and Novice Performance Within Computer Security Incident Response Teams. In D. D. Schmorrow & C. M. Fidopiastis (Eds.), *Foundations of Augmented Cognition: 9th International Conference, AC 2015, Held as Part of HCI International 2015, Los Angeles, CA, USA, August 2-7, 2015, Proceedings* (pp. 144–152). Cham: Springer International Publishing.

Singh, J., & Nene, M. J. (2013). A Survey on Machine Learning Techniques for Intrusion Detection Systems. *International Journal of Advanced Research in Computer and Communication Engineering*, *2*(11), 4349–4355.

Skierka, I., Morgus, R., Hohmann, M., & Maurer, T. (2015). *CSIRT Basics for Policy-Makers: The History, Types & Culture of Computer Security Incident Response Teams* (Transatlantic Dialogues on Security and Freedom in the Digital Age).

Solet, D. J., Norvell, J. M., Rutan, G. H., & Frankel, R. M. (2005). Lost in translation: challenges and opportunities in physician-to-physician communication during patient handoffs. *Academic Medicine : Journal of the Association of American Medical Colleges*, *80 12*, 1094–1099.

Solomon, D. (n.d.). The End of Reactive Security, and the Move to a Doctrine of Cyber Defence?

Sonnentag, S., Niessen, C., & Volmer, J. (2006). Expertise in Software Design. In K. A. Ericsson, N. Charness, P. J. Feltovich, & R. R. Hoffman (Eds.), *The Cambridge Handbook of Expertise and Expert Performance* (1st ed., pp. 373–388). New York, NY, USA: Cambridge University Press.

Sophos. (2019, February 19). Sophos Central Management Platform Now Features All Next-Gen Cybersecurity Protection from Sophos. *Associated Press*.

Sophos Labs. (2019). *Sophoslabs 2019 threat report*. Oxford, UK.

Staples, Z., & Sullivan, M. (2018). *2nd Age of Cyber: Philosophy and Principles to Shift the Advantage to the Cyber Defender*. Austin, TX.

Starmer, A., Spector, N., Srivastava, R., Rosenbluth, G., Dalal, A., Keohane, C., … Landrigan, C. (2014). Changes in medical errors after implementation fo a handoff program. *New England Journal of Medicine*, *371*(November), 1803–1812.

Steinke, J., Bolunmez, B., Fletcher, L., Wang, V., Tomassetti, A. J., Repchick, K. M., … Tetrick, L. E. (2015). Improving Cybersecurity Incident Response Team Effectiveness Using Teams-Based Research. *IEEE Security and Privacy*, *13*(4), 20–29.

Su, C., Huan, M., & Contractor, N. (2010). Understanding the structures, antecedent and outcomes of organisational learning and knowledge transfer: a multi-theoretical and multilevel network analysis. *European Journal of Internatianal Management*, *4*(6), 576–601.

Sultana, F. (2007). *Reflexivity, positionality and participatory ethics: Negotiating fieldwork dilemmas in international research*. *ACME* (Vol. 6).

Sundaramurthy, S. C., McHugh, J., Ou, X., Rajagopalan, S. R., & Wesch, M. (2014). An Anthropological Approach to Studying CSIRTs. *IEEE Security & Privacy*.

Sycara, K., & Lewis, M. (2004). Integrating intelligent agents into human teams. In *Team cognition: Understanding the factors that drive process and performance.* (pp. 203–231). Washington, DC, US: American Psychological Association.

Symantec. (2017). *Internet Security Threat Report*. *Annual Internet Security Threat Report*. Mountain View, CA.

Symantec Corporation. (2019). *Internet Security Threat Report Volume 24 | February 2019* (Vol. 24). Mountain View, CA.

Tadda, G. P., & Salerno, J. S. (2010a). Cyber situational awareness: Issues and research. In S. Jajodia, P. Liu, V. Swarup, & C. Wang (Eds.), *Advances in Information Security* (Vol. 46, pp. 15–36). Springer Science+Business Media, LLC.

Tadda, G. P., & Salerno, J. S. (2010b). Overview of Cyber Situation Awareness. In S. Jajodia (Ed.), *Cyber Situational Awareness* (pp. 15–36). Springer Science+Business Media, LLC.

Tanczer, L. M., Brass, I., & Carr, M. (2018). CSIRTs and Global Cybersecurity: How Technical Experts Support Science Diplomacy. *Global Policy*, *9*(S3), 60–66.

Tannenbaum, S. I., Beard, R. L., & Salas, E. (1992). Team Building and its Influence on Team Effectiveness: An Examination of Conceptual and Empirical Developments. *Advances in Psychology*, *82*(C), 117–153.

Taylor, F. W. (1911). The Principles of Scientific Management New York Harper & Row.

Tetrick, L. E., Zaccaro, S. J., Dalal, R. S., Steinke, J. A., Repchick, K. M., Hargrove, A. K., … Wang, V. (2016). *Improving Social Maturity of Cybersecurity Incident Response Teams*. Fairfax, VA: George Mason University.

The Council of Economic Advisors. (2018). *The Cost of Malicious Cyber Activity to the U.S. Economy*. Washington, DC, USA.

The MITRE Corporation. (n.d.). Operational Needs Assessment. Retrieved March 29, 2019, from https://www.mitre.org/publications/systems-engineering-guide/se-lifecycle-building-blocks/concept-development/operational-needs-assessment

The White House. (2018). *National Cyber Strategy of the United States of America*. Washington, DC, USA.

Tøndel, I. A., Line, M. B., & Jaatun, M. G. (2014). Information security incident management: Current practice as reported in the literature. *Computers & Security*, *45*, 42–57.

Truvé, S. (2017). *Machine Learning in Cyber Security: Age of the Centaurs*. *Recorded Future*. Sommerville, MA.

Tyworth, M., Giacobe, N. A., & Mancuso, V. (2012). Cyber situation awareness as distributed socio-cognitive work (pp. 84080F1-84080F9).

Tyworth, M., Giacobe, N. A., Mancuso, V., & Dancy, C. (2012). The distributed nature of cyber situation awareness. In *2012 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support* (pp. 174–178). IEEE.

United States Government Accountability Office. (2017). *High Risk Series: Progress on Many High-Risk Areas, While Substantial Efforts Needed on Others*. *Report to Congressional Committees*.

Urias, V. E., Leeuwen, B. V, Stout, W. M. S., & Lin, H. W. (2017). Dynamic cybersecurity training environments for an evolving cyber workforce. In *2017 IEEE International Symposium on Technologies for Homeland Security (HST)* (pp. 1–6).

Van der Kleij, R., Kleinhuis, G., & Young, H. (2017). Computer Security Incident Response Team Effectiveness: A Needs Assessment. *Frontiers in Psychology*, *8*, 2179.

Van Maanen, J. (2011). Ethnography as Work: Some Rules of Engagement. *Journal of Management Studies*, *48*(1), 218–234.

van Zadelhoff, M. (2017, May 4). Cybersecurity Has a Serious Talent Shortage. Here's How to Fix It. *Harvard Business Review: Security & Privacy*.

VIB, & Demisto. (2018). *The State of SOAR Report, 2018*.

Vidgen, R. (1998). Cybernetics and business processes: using the viable system model to develop an enterprise process architecture. *Knowledge and Process Management*, *5*(2), 118–131.

Vieane, A. Z., Funke, G. J., Gutzwiller, R. S., Mancuso, V., Sawyer, B. D., & Wickens, C. D. (2016). Addressing human factors gaps in cyber defense. *Proceedings of the Human Factors and Ergonomics Society*, *60*, 770–773.

Vlahos, J. (2015, September 20). Goodbye imaginary friends; hello A.I. dolls. *New York Times Magazine*, p. 44.

Wagenhals, L. W., Shin, I., Kim, D., & Levis, A. H. (2000). C4ISR architectures: II. A structured analysis approach for architecture design. *Systems Engineering*, *3*(4), 248–287.

Walker, J. (1998). The Viable Systems Model: A guide for co-operatives and federations. In *SMSE Strategic Management in the Social Economy training programme* (2.21). Author.

Waltl, B., & Vogl, R. (2018). Explainable Artificial Intelligence – the New Frontier in Legal Informatics. In E. Schweighofer, F. Kummer, A. Saarenpää, & B. Schafer (Eds.), *21st International Legal Informatics Symposium IRIS*.

Warden, P. (2018). Speech Commands: A Dataset for Limited-Vocabulary Speech Recognition. Mountain View, CA: Google Brain.

Werlinger, R., Muldner, K., Hawkey, K., & Beznosov, K. (2010). Preparation, detection, and analysis: the diagnostic work of IT security incident response. *Information Management & Computer Security*, *18*(1), 26–42.

West-Brown, M. J., Stikvoort, D., Kossakowski, K.-P., Killcrece, G., Ruefle, R., & Zajicek, M. (2003). *Handbook for Computer Security Incident Response Teams (CSIRTs)*. *SEI Digital Library* (2nd ed.). Pittsburgh, PA: Carnegie Mellon Software Engineering Institute.

White, A. R. (2019). Human expertise in the interpretation of remote sensing data : A cognitive task analysis of forest disturbance attribution. *Int J Appl Earth Obs Geoinformation*, *74*(March 2018), 37–44.

White, G., & Granado, N. (2009). Developing a Community Cyber Security Incident Response Capability. In *2009 42nd Hawaii International Conference on System Sciences* (pp. 1–9). IEEE.

Williams, B., & Hummelbrunner, R. (2010). Viable System Model. In *Systems Concepts in Action : A Practitioner's Toolkit* (pp. 199–214). Stanford University Press.

Williams, L. C. (2017, September 7). Spy chiefs set sights on AI and cyber. *FCW: The Business of Federal Technology*.

Xiong, W., Wu, L., Alleva, F., Droppo, J., Huang, X., & Stolcke, A. (2018). The Microsoft 2017 conversational speech recognition system. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 5934–5938).

Yates, K., Ed, D., Sullivan, M., Ph, D., Clark, R., & Ed, D. (2012). Integrated studies on the use of cognitive task analysis to capture surgical expertise for central venous catheter placement and open cricothyrotomy. *AJS*, *203*(1), 76–80.

Yen, J., Erbacher, R. F., Zhong, C., & Liu, P. (2014). Cognitive Process. In A. Kott, C. Wang, & R. F. Erbacher (Eds.), *Cyber Defense and Situational Awareness* (pp. 119–144). Cham: Springer International Publishing.

Yu, D., & Deng, L. (2016). *Automatic Speech Recognition*. London, UK: Springer.

Yuan, Y. C., Fulk, J., Monge, P. R., & Contractor, N. (2010). Expertise Directory Development, Shared Task Interdependence, and Strength of Communication Network Ties as Multilevel Predictors of Expertise Exchange in Transactive Memory Work Groups. *Communication Research*, *37*(1), 20–47.

Yufik, Y. (2014). Situational awareness, sensemaking, and situation understanding in cyber warfare. *Advances in Information Security*, *61*, 1–18.

Zachary, W., Techologies, C. M. Z. H., Crandall, B., Miller, T., & Nemeth, C. (2012). “Rapidized ” Cognitive Task Analysis. *IEEE Intelligent Systems*, *27*(2), 61–66.

Zinke, C., Anke, J., Meyer, K., & Schmidt, J. (2018). Modeling, Analysis and Control of Personal Data to Ensure Data Privacy -- A Use Case Driven Approach. In D. Nicholson (Ed.), *Advances in Human Factors in Cybersecurity: Proceedings of the AHFE 2017 International Conference on Human Factors in Cybersecurity, July 17−21, 2017, The Westin Bonaventure Hotel, Los Angeles, California, USA* (pp. 87–96). Cham: Springer International Publishing.

APPENDICES

Appendix A: Orienting questions during Study 1 interviews

1. Please describe the general organizational structure, starting with the SOC and working upward.
2. How are incidents reported within the organization?
3. Please walk me through, step-by-step, a recent incident as an example to frame your incident response process.
    a. Where does an incident “change hands”?
    b. What is the general sequence of events?
    c. Are there protocols or standards for these processes?
    d. Are these internal or external to your organization?
4. How are higher levels of management notified when an incident occurs?
5. How are team members held accountable during these “handoffs”
    a. Is there a verification step? If so, who performs this?
6. What are the top 3 types of incidents seen by your organization?
7. Where do problems tend to arise in your incident response process? (reference Q3)
8. How has the organization tried to address these problems?

Appendix B: Dimensions of expertise questions

1. As a T1 analyst, how would you contact other analysts? What about during an escalation?
2. How important is it regarding how you contact someone?
3. In your opinion, does your relationship with the other analyst affect your decision to contact that person, or *how* you decide to contact them?
4. Do you find yourself changing the style of communication to particular people? How? Why?
5. How do you choose who to escalate an incident to?
6. How do you know to send it to that person?
7. When you have to choose between two people to whom to escalate an incident (who have the same expertise/credentials), how would you choose?
8. What other aspects of the problem/incident are relevant when escalating a ticket, or passing to another analyst?
9. When would you decide NOT to escalate an incident? Why?

Appendix C: Study 1 IRB Package

APPLICATION NARRATIVE FORM

Purdue University, Institutional Review Board

1. Project Title: Determining System Requirements for Human-Machine Integration in Cyber Security Incident Response: Part 1

2. Principal Investigator: Barrett S. Caldwell, Professor, IE, bscaldwell@purdue.edu, +1 765 49-45412

Please address the following points regarding your proposed research:

A. PROPOSED RESEARCH RATIONALE

This observational field study is being conducted to gain better understanding of how work tasks vary across Computer Security Incident Response Teams (CSIRTs). Literature offers only a general view of workflow, making it difficult to capture the context and use it for designing better work solutions. The study will focus on observing tasks in "incident response", especially those in which information is passed between two people or two work teams. In so doing, researchers will be able to identify critical "handoff" points, where the teams currently use automation tools, and the areas in which the teams struggle. The study also includes some semi-structured interviews with knowledgeable team members to gain context and clarity.

The research questions to be addressed in this study include:

- *What is the general flow of work tasks within the team? What are the major task areas (categorized groups of tasks performed) (e.g incident handling, vulnerability handling, forensics investigation)?* (NOTE: these are terms with which the subjects are familiar as it is part of their job)
- *How is the organization structured (e.g. group of specialized teams, ad-hoc team, etc)? How is information shared within the organization?*
- *Are there any procedures or protocols for when information is shared? If so, what?*
- *Where does one major task area share information to another area?*
- *During these information sharing activities, is information flow constant, circumstantial? What dictates this flow of information?*
- *Where do the teams use automation in information sharing activities?*
- *Which information sharing activities are disrupted the easiest (i.e. when time is short, or when errors are made)?*

- *During which information sharing activities do teams struggle with the most? Which seem to be the most critical?*

B. SPECIFIC PROCEDURES TO BE FOLLOWED

With respect to observation participants, no specific procedures are required of subjects for this study. Only participants who consent to being a part of the study will be observed. Observations will be performed of normal daily work activities with no intervention from the investigators.

Before observations start, the observer will solicit an 'externship' with a supervisor or manager. This externship will include asking the participant (manager or supervisor) some semi-structured interview questions regarding the team, its context, process, and struggles, followed by shadowing the supervisor or manager, and introductions with team members. The interview questions will help the observer get a better understanding of the non-observable processes and procedures. The following questions are examples of what may be asked:

- *Please tell me about an incident your team has responded to recently. Start at the beginning, and walk me through it step by step.*
- *Can you generally describe for me the process?*
- *How is performance measured for this process?*
- *How do incidents typically escalate? How are higher levels of management notified?*
- *Where in these steps information is shared within the team, or between teams?*
- *Let's explore one of those information sharing activities in more detail. What is the sequence of events when information passed the next team? Please describe them in general terms for me. (For example, passing an incident to the internal vulnerability group, forensics group, or law enforcement)*
- *Are there protocols or standards that you are required to follow for these activities? (Where do these standards originate?)*
- *How are members of the team held accountable during these activities? Are there verification steps in place? Who performs them?*
- *Which, if any, of these information sharing points is considered the most important? (Top 3? Why?)*
- *Where do problems typically arise during these activities? Please describe one or two of these problems in detail.*
- *How has the organization tried to address these problems?*

Observations will include information such as the mode of information sharing (email, verbal, document, etc), the role of the person sharing it and receiving it (analyst, supervisor, manager), what *type of information is being shared, and characteristics of the information, such as time sensitivity, decision-making information, etc.*

Finally, if available, the researcher may collect existing documents regarding process flow and protocols to supplement observational findings. Many organizations already have a process flow diagram of some sort, and having this document would help reduce the data needed to be collected.

C. SUBJECTS TO BE INCLUDED

For the study described, subjects are a part of a team of computer security analysts working together on an established Computer Security Incident Response Teams (CSIRTs). These teams are formed within their respective organizations, and individuals are hired for expertise in computer security areas. The criteria for this study are not limited to any particular demographics.

The inclusion criterion for the subject population is that the team is an operating CSIRT that performs incident handling functions. If the CSIRT does not perform incident handling (i.e. these functions are outsourced), the team is not eligible for this study. Additionally, if the incident handling group within the CSIRT exceeds 100 members, the team will not be a good fit for this study based on resources available.

The investigators request to observe work activities on four (4) different CSIRTs, which vary in size. The maximum expected number of individuals the investigators expect to interact with is 400. The four teams should give some indication regarding variance and trends in handoff procedures and struggle points.

D. RECRUITMENT OF SUBJECTS AND OBTAINING INFORMED CONSENT

The researchers will by working with Purdue CERIAS to identify parent organizations that are willing to allow investigators access to their CSIRTs. Language regarding the general study will be shared with CERIAS for distribution, after which, official recruitment language will be shared with companies that express interest.

Access will then be secured through documentation that will be submitted to the IRB. Investigators will interact with managers (access granters) to explain the study, its motivations, and its procedures.

The manager or supervisor may also given consent at this point to be part of the study as both the interview participant and observation participant. After access is given, investigators will recruit participants through the manager via email (manager as a conduit, so no personal information is accessible to the researcher), which will describe the study and include a form of consent to the team members. This email will also emphasize that participation is voluntary, that their decision to participate will not affect their job or relationship with their employer, that their performance is not being evaluated, and that no directly or indirectly personally identifiable information will be collected as part of the study (see Appendix C.1). After participants express interest and grant consent, observations with interested parties will commence for approximately 40 work hours (multiple shifts across multiple days) for each CSIRT.

Personal contacts will also be utilized for recruitment.

E. PROCEDURES FOR PAYMENT OF SUBJECTS

Subjects will not be compensated, and the observations do not require time away from their normal work duties.

F. CONFIDENTIALITY

With respect to the observation participants, subjects' personally identifiable information (name, phone number, email, etc) and demographic information will not be recorded for this research.. Generalized roles (analyst 1, analyst 2, supervisor A, etc) will be recorded in notes to assist the investigators with memory of event sequence. These roles will be restricted to investigator notes, and will be destroyed after data analysis and dissertation defense are complete.

With respect to the interview participant (one from each team), audio recordings of the interview will be collected to assist with data capture. These recordings will be transcribed in a timely manner, then destroyed. Any files will be kept on the campus of Purdue University, and securely stored to minimize risk of a breach.

STORAGE:

Investigator notes (with no personally identifiable information) will be kept in their original hand-written form with no copies made, and maintained on Purdue's campus.

Recordings of semi-structured interviews will not include the name of the person being interviewed. Recordings will be transcribed in a timely manner to minimize storage time. Recordings will be kept on Purdue's campus on an external hard drive until transcribed. Transcribed (de-identified) interviews will be kept indefinitely for future reference and research.

MAINTAINING AND/OR DESTRUCTION:

[As above] Investigator notes (with no personally identifiable information) will be kept in their original hand-written form with no copies made, and securely stored on Purdue's campus.

Recordings of semi-structured interviews will not be digitally copied, and will be deleted from storage in a timely manner after transcription. The name of the individual will not be recorded in any way.

Any existing documents collected will be returned to the originating organization with no copies made.

G. POTENTIAL RISKS TO SUBJECTS

With respect to observations, the potential risk to subjects is minimal, as they will be observed performing their normal daily activities in their respective workplaces. No personally identifiable information will be recorded, and observations will include all members of a team. No individual will be singled out. Other levels of management (supervisors and managers) are included in the scope of the research, thus any risk related to manager-operator relations are spread across both subpopulations. Observations are the least intrusive and most minimal risk way of collecting this information. As signed consent forms are required for this portion of the study, and these consent forms (plus contact for obtaining signed forms) are the only records linking participants to the study, the forms will be securely stored on Purdue's campus to minimize risk of a confidentiality breach. Data from observations will be kept indefinitely for future reference and research.

With respect to the interviews, the potential risks to the subject is minimal, as the subject matter of the interviews is limited to normal daily activities. Due to audio recordings of the interview, one risk is a breach of confidentiality, which is being addressed with safeguards, such as secure storage on FileLocker or Purdue University Research Repository, timely transcription, and destruction of audio recordings. Audio recordings will not include name or contact information (subject will be de-identified). Transcribed (de-identified) interviews will be kept indefinitely for future reference and research.

There is no expected need for medical or professional interventions based on the observations or interviews being performed.

The employer will not receive raw data collected from either the observations or interview. Rather, synthesized data in the form of flow charts or general trends will be given to management, and will not include names or positions of employee participants.

H. BENEFITS TO BE GAINED BY THE INDIVIDUAL AND/OR SOCIETY

There are no direct benefits to subjects.

The possible benefits to society include a better understanding of the differences in the team operations of these teams. By collecting valuable contextual insights, the investigators will highlight key differences in operations to help drive customizable solutions (versus the current generalizable solutions) for cyber security incident response. The overall goal is to help make these incident response teams more effective, and improve overall computer security operations.

I. INVESTIGATOR’S EVALUATION OF THE RISK-BENEFIT RATIO

The probability and magnitude of possible harms are minimal for this observational study. The researchers are more interested in the flow of work tasks than the individuals that perform them, and no personally identifiable information will be collected. Furthermore, subjects will be observed performing their normal daily duties, and no interventions will be performed.

J. WRITTEN INFORMED CONSENT FORM *(SEE ATTACHED)*

M. SUPPORTING DOCUMENTS *(check all document that you will be submitting to IRB)*

X Recruitment advertisements, flyers, emails and letters.

___Survey instruments, questionnaires, tests, debriefing information, etc.

X Consent Form, Parental Permission, Assent Form

___Translated consent and recruitment documents

___If the research is a collaboration with another institution, that institution’s IRB or ethical board approval for the research or request for IRB deferral.

___If the research accesses the PSYC 120 Subject pool include the description to be posted on the web-based

recruitment program and the debriefing form to be used.

___Local review approval or affirmation of appropriateness for international research.

___If the research will be conducted in schools, businesses or organizations, include a letter from an

appropriate administrator or official permitting the conduct of the research.

___If the study involves an investigational drug/device, include product information or investigator brochure

___Other (please list)

Appendix C.1:

Solicitation Language, to be distributed through CERIAS network before Letter of Permission is received

Note: this recruitment letter also includes some language for Part 3 of this study, which as not yet been submitted to IRB. However, because the same team will need to accessed for both Part 1 and Part 3, the researchers wanted to ensure understanding of needed access on the part of the hosting organization.

Dear CERIAS Strategic Partner:

My name is Megan Nyre-Yu, and I am a CERIAS student and PhD candidate in Industrial Engineering. I am emailing to gauge interest in participating in my dissertation research, which aims to apply Human Factors methods to better understand information sharing functions (such as handoffs) in the cyber incident response process. The goal is to collect needs of the team (through understanding context and process), and produce insights and functional requirements for potential collaborative automation.

Two of the studies I am conducting involve interacting directly with CSIRTs performing incident response. An observation study will help identify how information flows during the general process, and during handoffs. I would observe different teams in their normal working environments for roughly 1 week (each) to build some context around how and why they do things the way they do. The second study (done at a later time with the same teams) includes a 10-min survey to understand perceived knowledge levels on the team, and how team members navigate that knowledge (do they know who has expertise in what).

As my focus is the process (not individuals), I will not identify participating team members (demographics, names, etc), nor will I identify a company by name. I may note characteristics of a company (CSIRT size, company size, company sector - Finance, Tech, Healthcare, Defense).

I am willing to sign a Non-Disclosure Agreement between myself and any company that is interested, and I have U.S. Citizenship, in case that is a requirement.

If you would like more information about this research, please contact me using the information below.

Sincerely,

Megan Nyre-Yu
PhD Candidate | *GROUPER Lab*
Purdue CERIAS Student
School of Industrial Engineering
Purdue University
mnyre@purdue.edu
mnyre19@gmail.com

Appendix C.2:

Recruitment Language, to be distributed through CERIAS network before Letter of Permission is received

Dear CERIAS Strategic Partner:

Thank you for your interest in my research. This research is split into two separate studies that involve the same CSIRTs. The studies have separate protocols, and will be separated by a matter of months.

The following email is in regards to the first of the two studies, which involves observations CSIRTs during normal operations. To ensure understanding and compliance, I will explain in the below email what is needed from your firm to participate in the study.

1st STUDY NEEDS:

The first study will involve your CSIRT, specifically the part of the team that performs incident handling. I have attached protocol information regarding what involvement is needed on behalf of your employees, expected risks and benefits to participating, as well as what specific questions I might ask. These forms are extensive, and include the motivation of the study, the research questions, and what I'm trying to investigate specifically. The attachments also include a Letter of Informed Consent, which will be presented to all the individuals I might interact with during the site visit. I am available via phone or email if additional information is needed from me regarding the study, site visit, and details thereof.

To reiterate from my original email, I will not identify participating team members (demographics, names, etc). Nor will I identify a company by name. I may note characteristics of a company (CSIRT size, company size, company sector - Finance, Tech, Healthcare, Defense). I am willing to sign a Non-Disclosure Agreement between myself and any company that is interested, and I have U.S. Citizenship, in case that is a requirement.

LETTER OF PERMISSION:

In compliance with Purdue's requirements for ethical human subjects research, the protocols for both studies must be reviewed and approved by the Institutional Review Board. Part of these requirements is a Letter of Permission from a team or site manager, which allows the researcher to visit your facility and interact with your employees.

Before the research can proceed, please provide this letter, along with name and contact information for your site.

{ENCLOSED: Application Narrative Form, Letter of Informed Consent}

Sincerely,

Megan Nyre-Yu
PhD Candidate **|** *GROUPER Lab*
Purdue CERIAS Student
School of Industrial Engineering
Purdue University
mnyre@purdue.edu
mnyre19@gmail.com

Appendix C.3:

Recruitment Language, to be distributed through company conduit to recruit participants

Dear Potential Participant:

I am a PhD candidate currently studying information sharing in CSIRTs, specifically how to improve them. In order to investigate the problem, I would like to observe teams who perform incident response, such as the team you are currently a part of. This email is an invitation to participate in these observation studies.

Basics of the study

Who: Cyber security incident response teams
What: Observations of the team performing incident handling, especially regarding information sharing practices
Why: To understand the process and how it might be improved
How Long: 40 hours, though you will not be present for all 40 hours
What if I don’t want to: Your participation is voluntary, and will not affect your employment

I have attached an information letter regarding what involvement is needed from you, expected risks and benefits to participating, as well as what specific aspects of what I am investigating. If you are interested in participating, please review and sign the Letter of Informed Consent, and email directly to me at the email address below. I am available via phone or email if additional information is needed from me regarding the study and details thereof.

{ENCLOSED: Letter of Informed Consent}

Sincerely,

Megan Nyre-Yu
PhD Candidate **|** *GROUPER Lab*
Purdue CERIAS Student
School of Industrial Engineering
Purdue University
mnyre@purdue.edu

Appendix C.4: Study 1 IRB Approval Letter

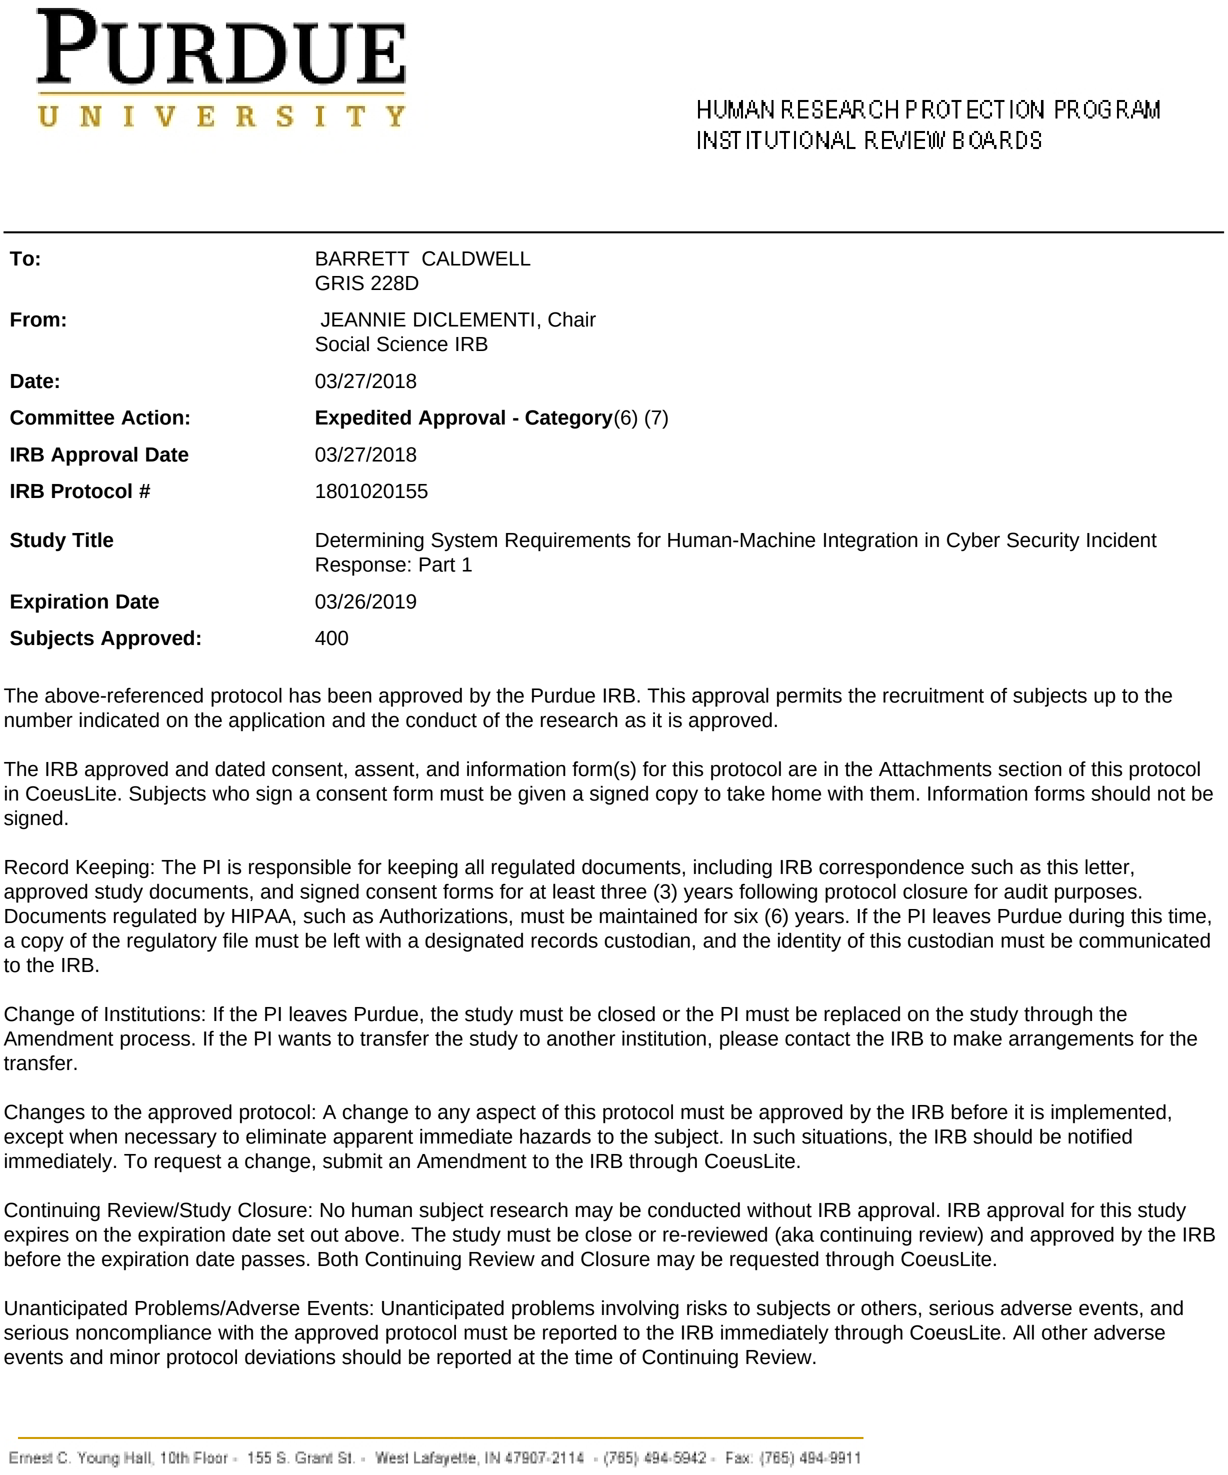

PURDUE
UNIVERSITY

HUMAN RESEARCH PROTECTION PROGRAM
INSTITUTIONAL REVIEW BOARDS

To:	BARRETT CALDWELL GRIS 228D
From:	JEANNIE DICLEMENTI, Chair Social Science IRB
Date:	03/27/2018
Committee Action:	**Expedited Approval - Category**(6) (7)
IRB Approval Date	03/27/2018
IRB Protocol #	1801020155
Study Title	Determining System Requirements for Human-Machine Integration in Cyber Security Incident Response: Part 1
Expiration Date	03/26/2019
Subjects Approved:	400

The above-referenced protocol has been approved by the Purdue IRB. This approval permits the recruitment of subjects up to the number indicated on the application and the conduct of the research as it is approved.

The IRB approved and dated consent, assent, and information form(s) for this protocol are in the Attachments section of this protocol in CoeusLite. Subjects who sign a consent form must be given a signed copy to take home with them. Information forms should not be signed.

Record Keeping: The PI is responsible for keeping all regulated documents, including IRB correspondence such as this letter, approved study documents, and signed consent forms for at least three (3) years following protocol closure for audit purposes. Documents regulated by HIPAA, such as Authorizations, must be maintained for six (6) years. If the PI leaves Purdue during this time, a copy of the regulatory file must be left with a designated records custodian, and the identity of this custodian must be communicated to the IRB.

Change of Institutions: If the PI leaves Purdue, the study must be closed or the PI must be replaced on the study through the Amendment process. If the PI wants to transfer the study to another institution, please contact the IRB to make arrangements for the transfer.

Changes to the approved protocol: A change to any aspect of this protocol must be approved by the IRB before it is implemented, except when necessary to eliminate apparent immediate hazards to the subject. In such situations, the IRB should be notified immediately. To request a change, submit an Amendment to the IRB through CoeusLite.

Continuing Review/Study Closure: No human subject research may be conducted without IRB approval. IRB approval for this study expires on the expiration date set out above. The study must be close or re-reviewed (aka continuing review) and approved by the IRB before the expiration date passes. Both Continuing Review and Closure may be requested through CoeusLite.

Unanticipated Problems/Adverse Events: Unanticipated problems involving risks to subjects or others, serious adverse events, and serious noncompliance with the approved protocol must be reported to the IRB immediately through CoeusLite. All other adverse events and minor protocol deviations should be reported at the time of Continuing Review.

Ernest C. Young Hall, 10th Floor - 155 S. Grant St. - West Lafayette, IN 47907-2114 - (765) 494-5942 - Fax: (765) 494-9911

Appendix C.5: Study 1 Observation Consent Form

Purdue IRB Protocol #: 1801020155 - Expires on: 26-MAR-2019

RESEARCH PARTICIPANT CONSENT FORM
OBSERVATIONS

Observations of Functioning CSIRTs to Determine Workflow and Critical Handoffs
Barrett S. Caldwell, Ph.D
School of Industrial Engineering
Purdue University

What is the purpose of this study?

This study aims to better understand how members of Computer Security Incident Response Teams (CSIRTs) share information during normal work activities, the sequence of work tasks performed, and which points in the process teams struggle with (in terms of errors, time, etc.). More specifically, the researchers are trying to identify opportunities in information sharing activities. Additionally, the researchers would like to understand where automation tools are currently used in those information sharing activities. This study aims to do this though observations.

You are being asked to participate because of your membership in an active Computer Security Incident Response Team (CSIRT).

The study is aiming to collect observation data from four (4) different CSIRTs, with roughly 100 people or less on each team (or sub-team) that performs incident handling functions.

What will I do if I choose to be in this study?

If you are participating in the observation part of this study, there is nothing you will be asked to do outside of your normal daily activities. The researchers may ask some clarification questions to verify what was seen, but interactions will be kept to a minimum to reduce interruption to your job. The researchers will be performing observations, which is watching normal operations within your team, and taking notes on sequence of events. Again, short clarification questions may be asked to verify findings.

Note that these research procedures are non-experimental, meaning there are no treatments being tested in this setting. The researchers are solely interested in the team's normal daily activities to understand the general process and struggle points.

Data collected includes a general order of events, and how normal daily activities happen, in what order, where errors or struggles occur, what kind of information is shared within the team, and how it is shared (email, verbal, shared system). The student researcher will be watching activities for roughly 40 hours over a week, and taking notes.

Please note that your performance is not being evaluated. The researchers are first and foremost here to observe the team (not individuals), and are not here to test you or measure how good you are at your job. They are here to better understand your work environment and how your team shares information.

How long will I be in the study?

The student researcher will be watching activities for roughly 40 hours over the period of one week. Other shifts may be covered during this time, so you may not be present for all 40 hours.

What are the possible risks or discomforts?

The risks for this observation study are minimal. The risks are no greater than you would encounter during routine tasks in your workplace. It may be awkward or uncomfortable to have an outside member of the CSIRT observe operations. However, be assured that the researchers are not interested in individual performance, but rather the general workflow and how information passes between team members.

The observers will be spreading out observation time over several shifts, and will be shifting around the workspace to observe the entire team, including supervisors and managers when interactions involve them.

Please know that all research carries the risk for breach of confidentiality. Safeguards are in place to prevent this risk, which can be found in the confidentiality section of the form.

Are there any potential benefits?

There are no anticipated direct benefits you or your fellow teammates. The general benefits to society may include a better understanding of CSIRTs needs based on insights collected during these observations.

Will information about me and my participation be kept confidential?

The researchers are not collecting personally identifiable information from you during the observations. The researchers will also be retaining a signed copy of this consent form, which link you to this study.

The researchers will be taking notes about how information is passed between roles within in the team, which means that your role will be recorded, but in a way that you cannot be identified (Analyst 1, Analyst 2, Supervisor 1, etc). These notes are solely for assisting the researcher in remembering observations. In addition to omission of personally identifiable details in observation data, and as a safeguard to confidentiality, these notes and the signed copy of this consent form will be securely stored on the campus of Purdue University.

Results will be made into a flow chart or diagram of what happens during normal work activities, which will include what level of the organization (operator, supervisor, manager) receives or sends information. These resulting flow charts or diagrams, which reflect daily operations, will be shared with your company.

The project's research records may be reviewed by departments at Purdue University responsible for regulatory and research oversight.

What are my rights if I take part in this study?

Your participation in this study is voluntary. You may choose not to participate or, if you agree to participate, you can withdraw your participation at any time without penalty or loss of benefits to which you are otherwise entitled.

If at any time you would like to withdraw from the study, please approach the observer and verbalize this fact. The observer will stop observations of you. However, observations of the rest of the team will continue. Your decision to participate or not in the research will have no effect on the your relationship with your employer.

Who can I contact if I have questions about the study?

If you have questions, comments or concerns about this research project, you can talk to one of the researchers. Please contact the PI for the project, Dr. Barrett S. Caldwell (bscaldwell@purdue.edu, +1 765 494 5412) or the student investigator and primary point of contact, Megan Nyre-Yu (mnyre@purdue.edu , +1 224 622 9765)

If you have questions about your rights while taking part in the study or have concerns about the treatment of research participants, please call the Human Research Protection Program at (765) 494-5942, email (irb@purdue.edu)or write to:

Human Research Protection Program - Purdue University
Ernest C. Young Hall, Room 1032
155 S. Grant St.,
West Lafayette, IN 47907-2114

Documentation of Informed Consent

I have had the opportunity to read this consent form and have the research study explained. I have had the opportunity to ask questions about the research study, and my questions have been answered. I am prepared to participate in the research study described above. I will be offered a copy of this consent form after I sign it.

____ ________________________
Participant’s Signature Date

____
Participant’s Name

____ __________________________
Researcher’s Signature Date

Appendix C.6: Study 1 Interview Consent Form

Purdue IRB Protocol #: 1801020155 - Expires on: 26-MAR-2019

RESEARCH PARTICIPANT CONSENT FORM
SEMI-STRUCTURED INTERVIEW
Observations of Functioning CSIRTs to Determine Workflow and Critical Handoffs
Barrett S. Caldwell, Ph.D
School of Industrial Engineering
Purdue University

What is the purpose of this study?

This study aims to better understand how members of Computer Security Incident Response Teams (CSIRTs) share information during normal work activities, the sequence of work tasks performed, and which points in the process teams struggle with (in terms of errors, time, etc.). More specifically, the researchers are trying to identify opportunities in information sharing activities. Additionally, the researchers would like to understand where automation tools are currently used in those information sharing activities. This study aims to do this though observations of CSIRTs and a semi-structured interview with one knowledgeable person (such as a supervisor), from the team.

You are being asked to participate because of your membership in an active Computer Security Incident Response Team (CSIRT), or your position as a knowledgeable team member or supervisor.

The study is aiming to collect observation data from four (4) different CSIRTs, with roughly 100 people or less on each team (or sub-team) that performs incident handling functions.

What will I do if I choose to be in this study?

If you are participating in the interview part of the study, the researcher will need an hour or so of your time to interview you as an experienced member of the team, which will help the researcher get an introduction to your team and its processes. The researcher will be asking questions related to the team's structure, roles, and main responsibilities. This part of the interview will be recorded.

Following the recorded interview, the researcher will begin observations, starting with shadowing your activities as they relate to the team (as allowed), and asking clarifications regarding the activities the team is performing, such as "how does person A send information to person B? How do they know the information was received? How do they know it was correct?".

Note that these research procedures are non-experimental, meaning there are no treatments being tested in this setting. The researchers are solely interested in the team's normal daily activities to understand the general process and struggle points. The semi-structured interview is aimed at asking a series of questions to get a better understanding of these daily activities from someone with experience with the team and its processes. The researcher is interested in a general order of events, and how normal daily activities happen, in what order, where errors or struggles occur, what kind of information is shared within the team, and how it is shared (email, verbal, shared system).

Please note that your performance is not being evaluated. The researchers are first and foremost here to observe the team (not individuals), and are not here to test you or measure how good you are at your job. They are here to better understand your work environment and how your team shares information.

How long will I be in the study?

The student researcher will need roughly 1 hour of your time away from the team for the interview. Following the interview, the researcher will begin observations by shadowing your activities as they relate to the team (as allowed) for another hour, after which the researcher will transition to observations of the rest of the team. The researcher will be watching activities for roughly 40 hours over the period of one week. Other shifts may be covered during this time, so you may not be present for all 40 hours.

What are the possible risks or discomforts?

The risks for this interview are minimal. As the questions are related to standard work processes that occur on a regular basis, the risks are no greater than you would encounter during routine tasks in your workplace. It may be awkward or uncomfortable to have an outside member of the CSIRT shadow you and observe operations. However, be assured that the researchers are not interested in individual performance, but rather the general workflow and how information passes between team members. The shadowing is a way to start the observations with a team member that is familiar with the researcher, and has already shared dialogue regarding the process.

The observers will be spreading out the remaining observation time over several shifts, and will be shifting around the workspace to observe the entire team, including supervisors and managers when interactions involve them.

Please know that all research carries the risk for breach of confidentiality. Safeguards are in place to prevent this risk, and can be found in the confidentiality section of this form.

Are there any potential benefits?

There are no anticipated direct benefits you or your fellow teammates. The general benefits to society may include a better understanding of CSIRTs needs based on insights collected during the interview and observations.

Will information about me and my participation be kept confidential?

As the interview portion will be recorded, the researchers will be collecting personally identifiable information from you (voice recording). The researchers will also be retaining a signed copy of this consent form, which link you to this study.

Please note that all recordings and research records will be securely stored on Purdue's campus. Recordings will be transcribed in a timely manner, and then promptly destroyed to help reduce risk of a confidentiality breach. The signed consent form will be the only remaining record of your participation, which will be securely stored on Purdue University's campus. Your participation will remain anonymous for the data analysis and publication aspects of the study.

Transcribed interviews will be de-identified, meaning there will be no connection between you and the script. These de-identified transcripts may be maintained for future reference or research.

Results will be made into a flow charts or diagrams of what happens during normal work activities, which will include what level of the organization (operator, supervisor, manager) receives or sends information.

Your company will receive this analyzed form of the data (flow charts and insights regarding information flow within the team).

The project's research records may be reviewed by departments at Purdue University responsible for regulatory and research oversight.

What are my rights if I take part in this study?

Your participation in this study is voluntary. You may choose not to participate or, if you agree to participate, you can withdraw your participation at any time without penalty or loss of benefits to which you are otherwise entitled.

If at any time you would like to withdraw from the study, please inform the interviewer. Your decision to participate or not in the research will have no effect on the your relationship with your employer.

Who can I contact if I have questions about the study?

If you have questions, comments or concerns about this research project, you can talk to one of the researchers. Please contact the PI for the project, Dr. Barrett S. Caldwell (bscaldwell@purdue.edu, +1 765 494 5412) or the student investigator and primary point of contact, Megan Nyre-Yu (mnyre@purdue.edu , +1 224 622 9765)

If you have questions about your rights while taking part in the study or have concerns about the treatment of research participants, please call the Human Research Protection Program at (765) 494-5942, email (irb@purdue.edu)or write to:

Human Research Protection Program - Purdue University
Ernest C. Young Hall, Room 1032
155 S. Grant St.,
West Lafayette, IN 47907-2114

Documentation of Informed Consent

I have had the opportunity to read this consent form and have the research study explained. I have had the opportunity to ask questions about the research study, and my questions have been answered. I am prepared to participate in the research study described above. I will be offered a copy of this consent form after I sign it.

_____________________________________ _____________________
Participant's Signature Date

____
Participant's Name

_____________________________________ _____________________
Researcher's Signature Date

Appendix D: Study 1 Themes

Rank	Freq.	Theme
1	38	Communication, feedback and accountability are necessary for IR, awareness, and learning; If lacking within or between levels of org, issues arise
2	36	Organizational alignment on security priorities and awareness of IR issues is important for "full-cycle" IR process
3	35	Continuity of awareness and documentation around incidents is important
4	23	IR requires a wide range of skills and flexibility; Workforce may not be able to maintain if not designed to do so
5	22	IR requires a wide range of activities, including filtering and decision making; These can be split based on expertise or authority of analysts
6	20	Automation is seen as a potential solution for low-level tasks and coordination, but considered out of reach for teams who don't have the support resources
7	14	Knowledge sharing (in a repository, in person, or through other channels) may be important for learning and shared awareness
8	12	Formal and informal roles emerge to meet an organizational need for management, communication, and decision making
9	11	Identity and culture of the team affect communications and responsibilities
10	10	Handoffs are varied in terms of procedure, formality, and documentation; In whatever form, they are important for continuity in several contexts
11	8	Incident handling methods may be indicators of organizational maturity; Maturity as a focus may drive incident handling methods

Appendix E: Study 1 Workflow Diagrams

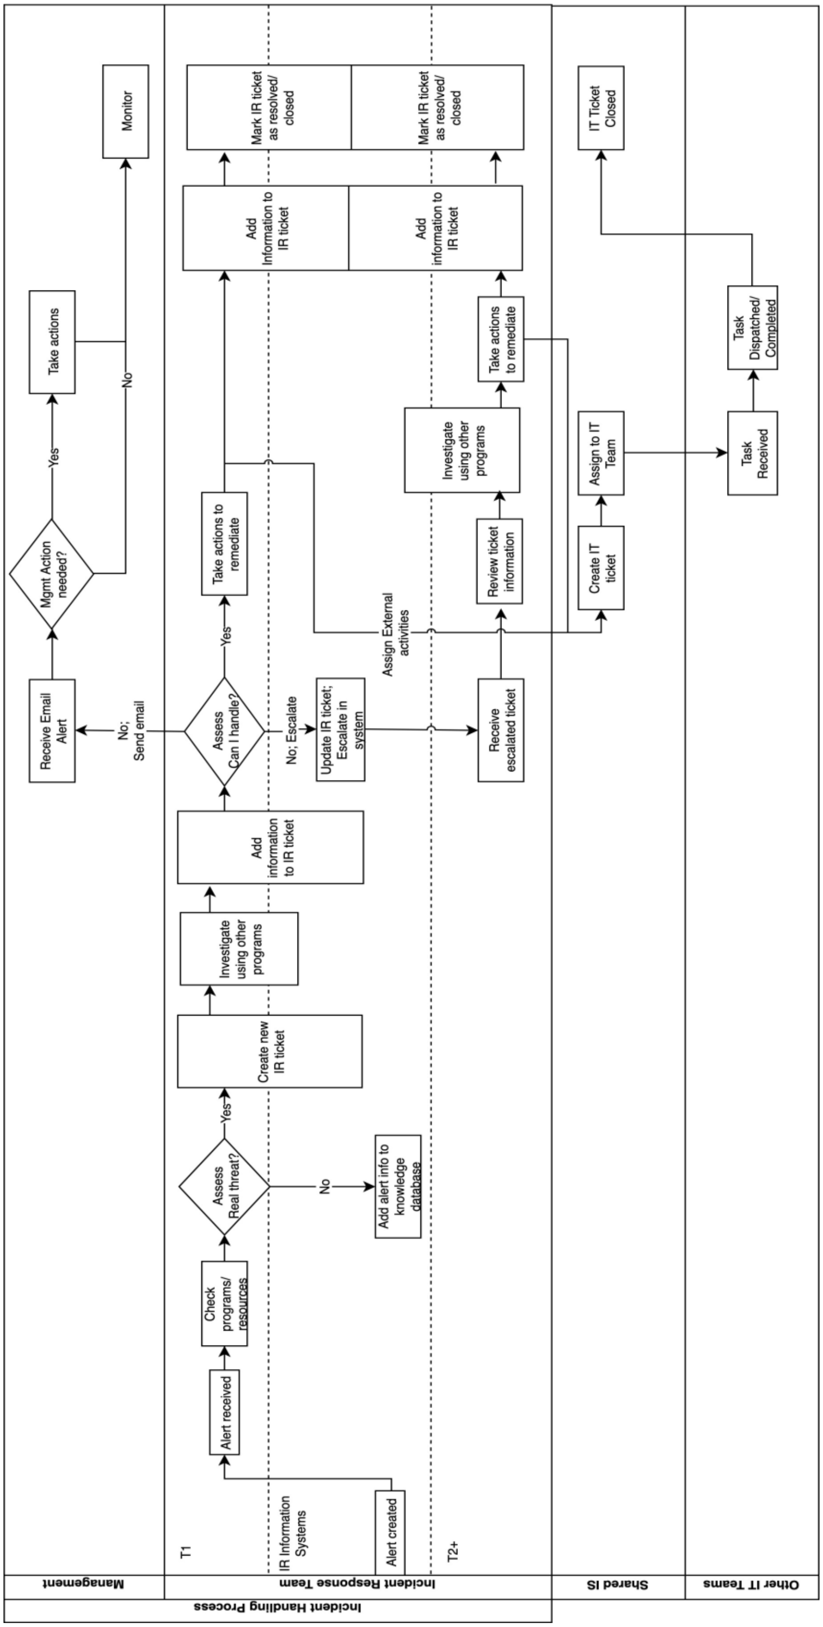
Incident Handling Process
Management
Incident Response Team
Shared IS
Other IT Teams
T1
IR Information Systems
T2+
Alert created
Alert received
Check programs/ resources
Assess Real threat?
No
Add alert info to knowledge database
Yes
Create new IR ticket
Investigate using other programs
Add information to IR ticket
Assess Can I handle?
No; Send email
Receive Email Alert
Mgmt Action needed?
Yes
Take actions
No
Monitor
Yes
Take actions to remediate
Add Information to IR ticket
Mark IR ticket as resolved/ closed
No; Escalate
Update IR ticket; Escalate in system
Receive escalated ticket
Assign External activities
Review ticket information
Investigate using other programs
Take actions to remediate
Add information to IR ticket
Mark IR ticket as resolved/ closed
Create IT ticket
Assign to IT Team
Task Received
Task Dispatched/ Completed
IT Ticket Closed

Figure.D.1. Workflow Diagram from Team 1 (For-Profit Company)

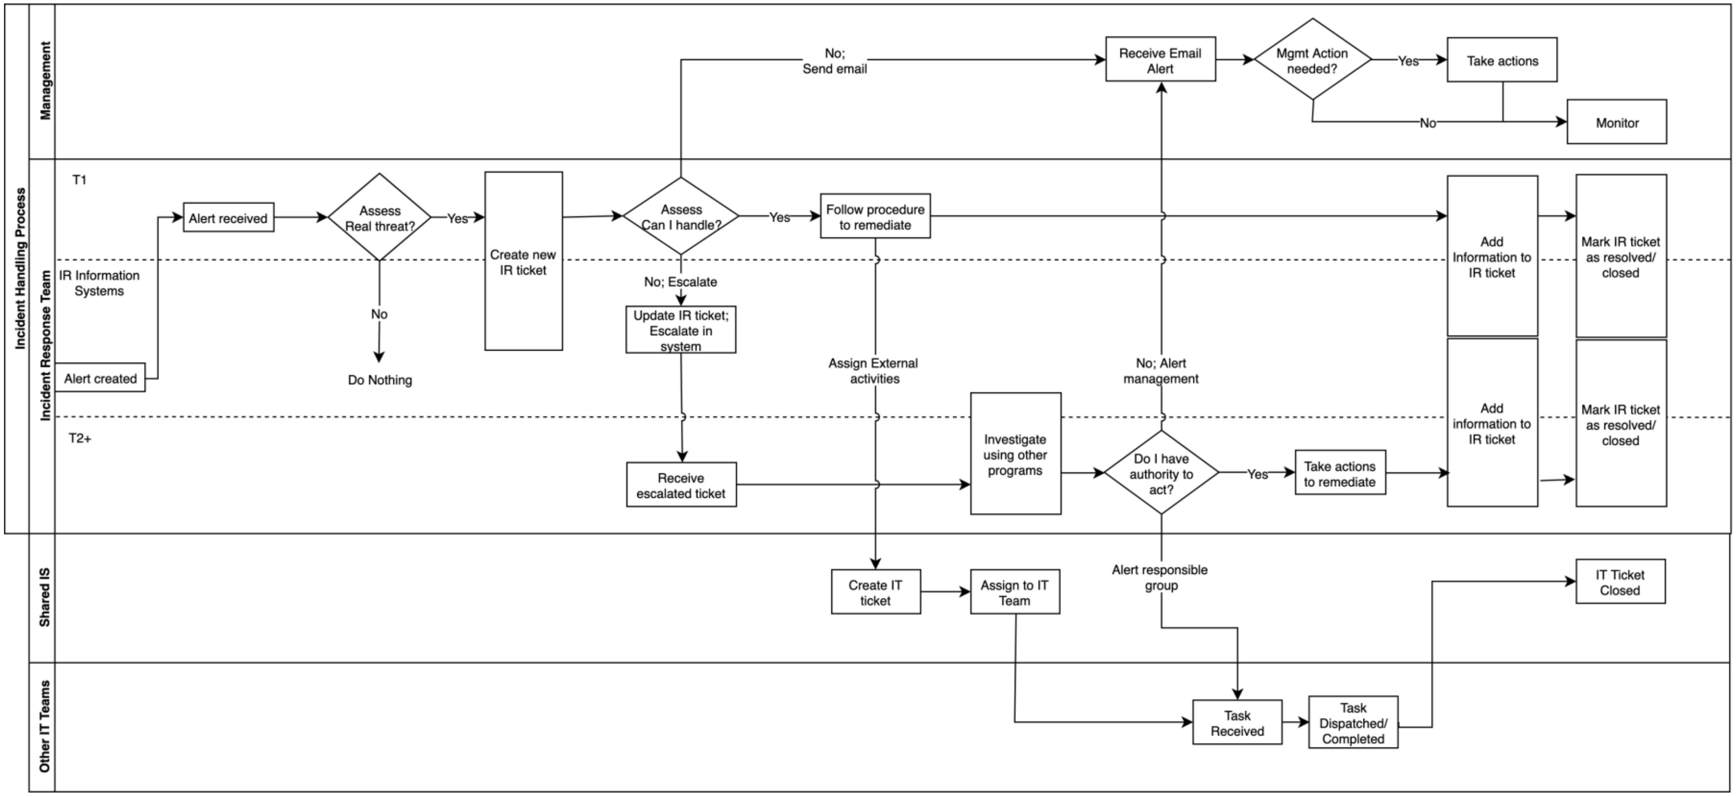

Figure.D.2. Workflow Diagram from Team 2 (Public University)

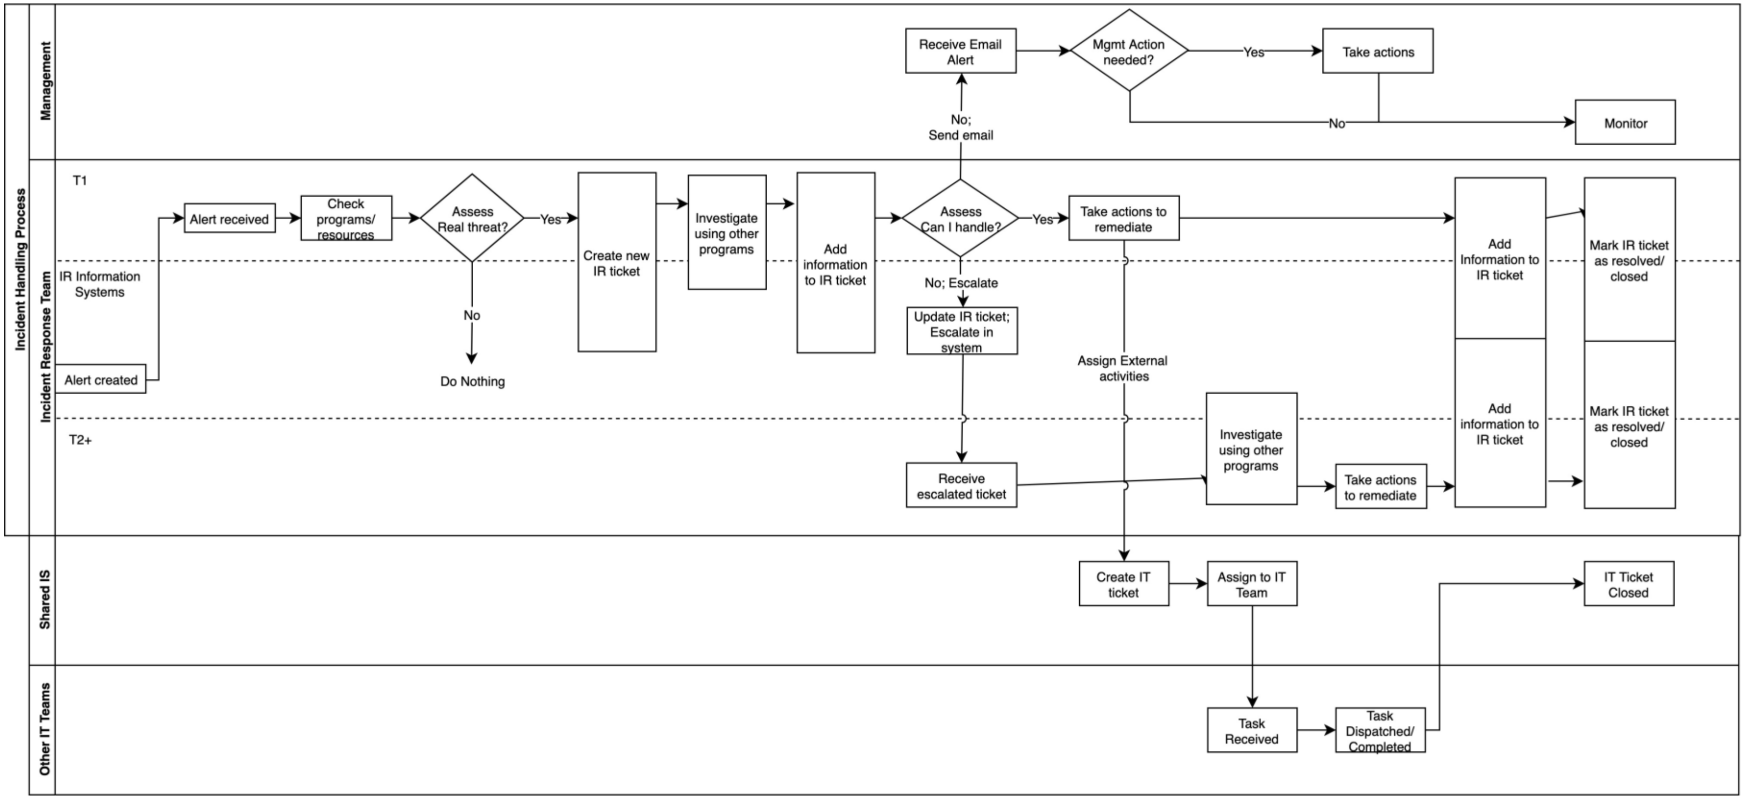
Incident Handling Process
Management
Incident Response Team
Shared IS
Other IT Teams
T1
IR Information Systems
T2+
Alert created
Alert received
Check programs/ resources
Assess Real threat?
Yes
No
Do Nothing
Create new IR ticket
Investigate using other programs
Add information to IR ticket
Assess Can I handle?
No; Send email
Receive Email Alert
Mgmt Action needed?
Yes
Take actions
No
Monitor
Yes
Take actions to remediate
No; Escalate
Update IR ticket; Escalate in system
Assign External activities
Receive escalated ticket
Investigate using other programs
Take actions to remediate
Add Information to IR ticket
Add information to IR ticket
Mark IR ticket as resolved/ closed
Mark IR ticket as resolved/ closed
Create IT ticket
Assign to IT Team
IT Ticket Closed
Task Received
Task Dispatched/ Completed

Figure.D.3. Workflow Diagram from Team 3 (State Government)

Appendix F: Study 2 Task Diagrams

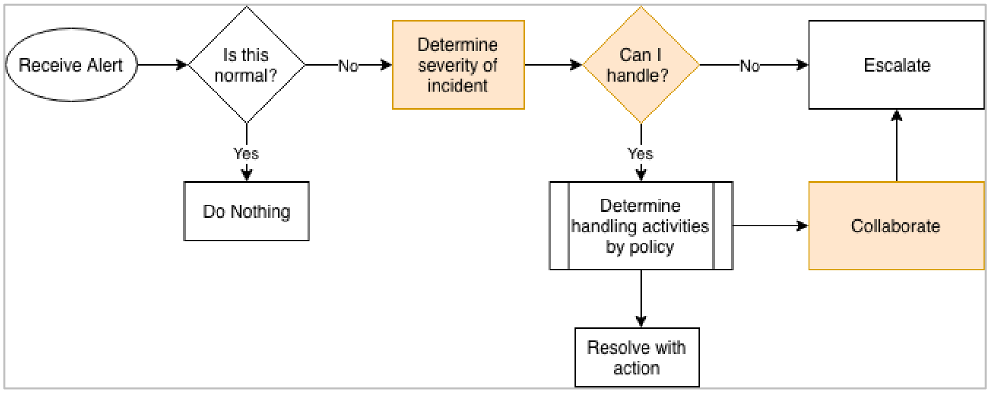

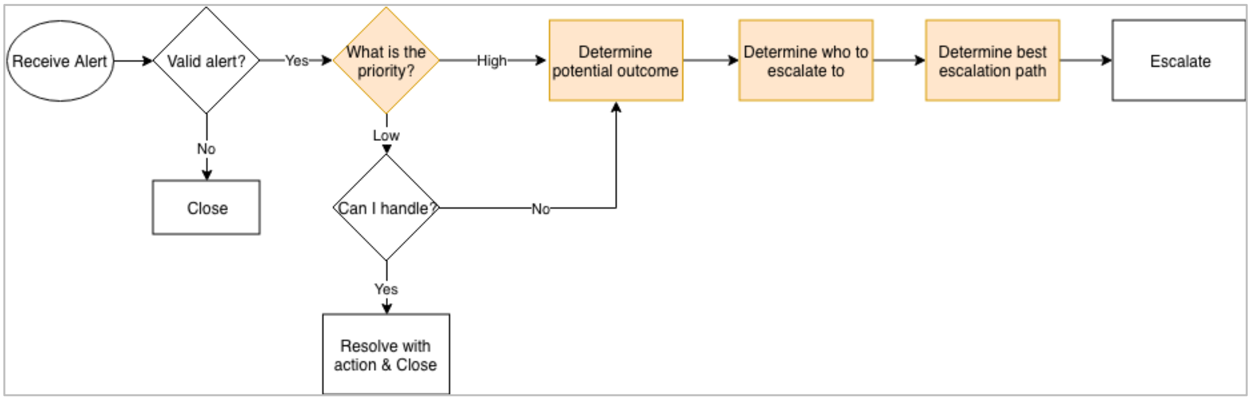
Receive Alert
Valid alert?
Yes
What is the priority?
High
Determine potential outcome
Determine who to escalate to
Determine best escalation path
Escalate
No
Close
Low
Can I handle?
No
Yes
Resolve with action & Close

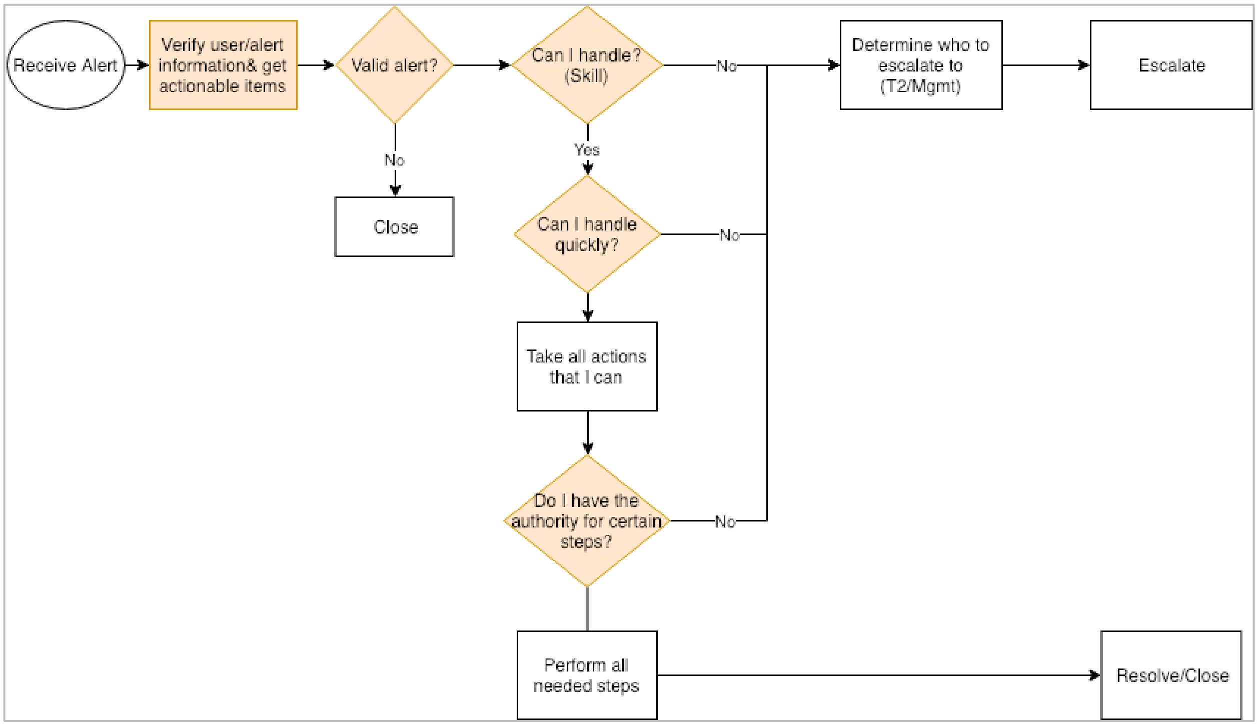
Receive Alert
Verify user/alert information& get actionable items
Valid alert?
Can I handle? (Skill)
No
Determine who to escalate to (T2/Mgmt)
Escalate
No
Close
Yes
Can I handle quickly?
No
Take all actions that I can
Do I have the authority for certain steps?
No
Perform all needed steps
Resolve/Close

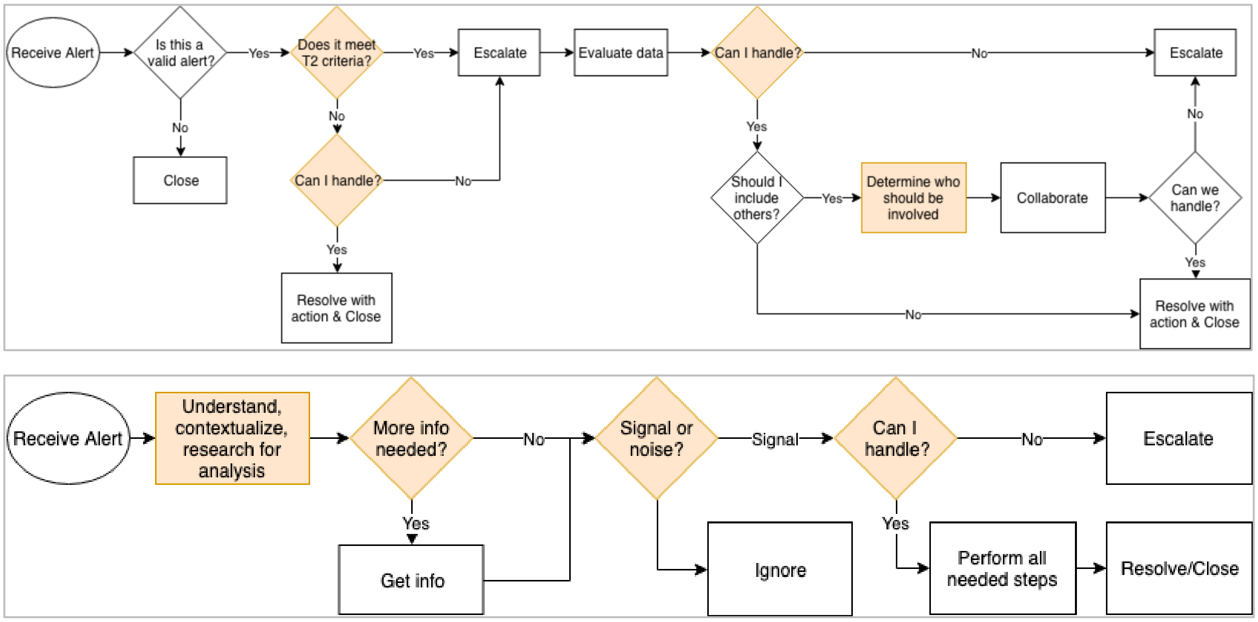

Figures: Participant Task Diagrams of Escalations

Appendix G: VSM Diagnostic

System 1 Diagnosis

System 1 Overview

In order to start the system diagnosis, it was first critical to ask some basic questions about the system of interest. Mainly, what creates value for external "customers?" (Hoverstadt, 2010) While the larger organizations (company, academic institution, and state government) have a broader scope of what provides value, this analysis starts with the value provided by IT/security operations in relation to primary and support functions of the larger whole. Thus, for the following analysis, System 1 is comprised of subsidiaries relating to IT/security operations. Within these organizations, IT was the main function of providing stable operations, while the support function was monitoring or maintaining the security of those operations.

With IT/security operations, several complexity drivers (Hoverstadt, 2010) were identified. Technology is a complexity driver, as there are different groups doing different things, or specialized teams, within the respective organizations. These teams work on different systems to do different activities to support IT and security functions. For instance, there could be different teams to support each security appliance, manage hardware, or conduct different kinds of network analysis. Geography was also a complexity driver. All three teams had evidence of geographic complexity in the physical distribution of System 1 subsidiaries. Lastly, time is a complexity driver due to the fact that many incidents extend beyond the staying power of one individual team (Hoverstadt, 2010). With tiered organizations and shift schedules, it is not unusual for an incident to be passed up through tiers, passed between shifts, or tasked as particular actions to other subsidiaries in System 1.

As previously described, an incident may involve one or more tiers within a SOC subsidiary before being passed to another (separate) subsidiary for additional action. The act of passing an incident from a lower tier to a higher tier is called *escalation*. Escalation can be described as a cybernetic process in which information flows between tiers regarding incident details with the goal of communicating and controlling the overall incident handling process. However, collected data in this study indicated that escalation is often treated as a one-way passing of incident data that has no reciprocating information flow or feedback. Thus, from a systems perspective, the incomplete cybernetic loop can prevent learning and process improvement at lower levels of incident response.

Table.F.1. System 1 Comparison

SYSTEM 1	**Location 1**	**Location 2**	**Location 3**
Environment	Mostly collocated teams of Tiered responders (some work from home); Large room with many desks; Usually quiet, but conversations not discouraged; Shared screen for monitoring incoming incidents; High percentage of team present;	Mostly collocated groups of Tiered responders (some work from home); Small quiet room with cubicles, one office; Very quiet, hushed conversations between individuals	Separate groups of tiered responders (all T2 and above work in a different city); Small quiet room with many desks; Only 1-3 people present; Non-operational shared screen; Quiet but conversations not discouraged
Operations	Incidents come in, are handled by T1, ticket is created, information is put in to describe the problem and handling steps; If needed, incident is escalated to T2; If other activities are required by other teams, the analyst assigns to another team	Incidents come in, are handled by T1, ticket is created, information is put in to describe the problem and handling steps; If T1 does not recognize the problem, incident is escalated to T2	Incidents come in, are handled by T1, ticket is created, information is put in to describe the problem and handling steps; If T1 does not recognize the problem, lead takes a look; If not familiar to lead, incident is escalated by lead to T2; If other activities are required by other teams, the lead assigns to another team
Management	Shift lead, Deputy manager (metrics, quality), Manager (staffing, performance)	(informal) T1 lead, Interim Manager (covered on all incidents)	T1 Lead, Manager (covered on all incidents)
Other parts of System 1 that are not included above	T2-4, ENGR		T2-4, ENGR
Constraints	Work 8 hour shift; Respond to incidents quickly; System access restricted by level	Only handle certain incidents (pass up if not in that group); System access restricted by level; Always cover manager	Only handle certain incidents (pass up if not in that group); System access restricted by level; Always cover manager
Accountability and Measures of Performance	Lead/Deputy check quality of ticketing (depth); Verbal reporting at shift change Time to analyze	Accountability not really exercised; No measures of performance were indicated	Lead check quality of ticketing (completeness)

System 1 Environment

The physical environments of Locations 1 – 3 were similar with respect to noise level, independent versus interdependent taskwork, and the fact that similar tier analysts were collocated, often in the same room. All teams had some off-site collaborators, usually higher-tier analysts or engineers that work from home several days a week. The authors note that, on two of the three teams, these other analysts/engineers reported to a different manager, and would be considered a different System 1 component (not studied here).

Some of the key differences between team environments included spatial layout, shared resources, and general analyst presence. In terms of space, two teams had no barriers between desks, allowing for direct verbal communication between analysts if needed. The last team had cubicles, which forced analysts to get up from their desk to directly interact with another analyst, or use electronic forms of communication, even though they were located in the same room. The general ambience of the environment was different between sites with physical barriers between analysts and no barriers. That is, more conversation and collaboration was evident in teams with no physical barriers between analysts. Next, some teams had shared screens for monitoring incoming incidents, which allowed for shared awareness of the current incident pool. Two of the three teams either had no shared screen, or the screen was not operational. For smaller teams (such as Location 3), a shared screen may not be warranted). Lastly, the number of people present in the security operations room varied greatly between teams, with one teams having at most 3 analysts, another having on average 5-8, and the last having up to 12 analysts at a time.

System 1 Operations

All teams had the tiered structure for operations, and included some form of support subsidiary for the general System 1 group. Other types of activities (preventative, maintenance) were also part of System 1. The tier levels effectively sort incidents and handling activities by the steps in the incident handling process. For instance, Tier 1 (T1) includes monitoring, filtering, triage, and preliminary investigation. T1 provides low-level response and coverage for routine incident types. Tier 2 (T2) conducts additional investigation and remediation activities, if needed. Tiers above T2 (usually Tier 3 or 4, depending on the company) provide deep analysis, reverse engineering, forensics investigations and more. Support activities seen in all organizations included engineering and development, which were responsible for maintaining appliances, creating new scripts/code, configuring of software, and more. Lastly, some teams had additional groups or individuals who conducted other incident response related activities, such as penetration testing (Steinke et al., 2015).

Operations within each organization were very similar in terms of procedure. Incoming incidents were investigated by T1 analysts, who look into the problem, create a ticket to document, and decide what happens next (close/task another group, escalate). The major difference between the organizations was the policy to direct how/when/to whom an incident is escalated. One team encouraged analysts to go as far as they could before escalating, and had a wide range of incident types that a T1 analyst could work on. Even if escalation were imminent, the analyst would do preliminary investigative steps and include all information in the ticket before escalating. The other two teams had strict policy regarding what a T1 analyst should and should not handle. Some incidents went straight to escalation after a skeleton ticket was created (with basic information). Within the third team, a lead who was capable of both T1 and some T2 would view incoming tickets in tandem with the T1 analyst who claimed the ticket. Sometimes this person would "pull up" an incident (reverse escalation) to investigate further before performing a full escalation to another group/team in System 1.

In general, operations were dispersed geographically for all locations between System 1 groups/teams. For the most part, Tier 1 activities were collocated within a single room. However, Tier 2 analysts were not collocated (in the same building, or even the same city), but were somehow accessible through shared information systems. Due to some of this geographic complexity, information sharing in operations was not always consistent in terms of channel and synchronicity of communication.

Management of System 1 Subsidiaries

The management structures looked very different between the three locations within System 1. In two of the three teams, the higher tier/support personnel reported to different managers (different parts of System 1. In the remaining team, all people reported to the same manager. Additionally, All three teams had some form of lead, whether formal or informal, to help manage T1 responders. Leads were responsible for everything from helping individuals with incident handling tasks to mediating issues within the smaller group. Leads tended to have more experience than average T1 analyst, or more capability to perform the job (which was not always measured in tenure). The lead in two of the three teams was responsible for shift handoff procedures, and did some form of quality check on the tickets completed during their shift. Policy regarding coverage of managers on communications was also different between the three teams. In Location 1, managers were only covered on high-impact incidents that might need to be communicated upwards to System 3 level management. The other two teams had a cover-all list of people to include in escalations or email chains, which included the manager.

Management-Imposed Constraints

All teams had relative time constraints with respect to security operations. There were a certain number of hours in a shift in which analysts were assigned to work. In general, all T1 analysts had to work in one location on a strict schedule. T1 analysts were not allowed to work from home or from another location. There was also the expectation that, within a given shift, high-priority incidents were identified, handled, and escalated accordingly. There was generally a constraint of access. For example, if an analyst did not have access to a particular tool or system, the constraint imposed was that they could not use that tool/system. In some locations, higher-level credentials or clearances drove this constraint. In general, the constraint was that the tool/information/system would give the lower tier (less experience, skilled) person too much power, or potential to cause damage/harm.

One management-imposed constraint in information sharing was visibility. In two of the three organizations, management created a policy to always cover management via email on escalated incidents or incidents that were handed off to other System 1 subsidiaries. Though more of a policy than a constraint, management did impose protocols to drive immediate escalation of incidents of certain types in some teams, such that Tier 1 analysts would not do anything more than create a ticket with basic information. The alternative scenario allowed the T1 analyst some freedom to work on it as much as they were able, and then escalate the incident. From a cybernetics perspective, the creation of the ticket itself creates no cybernetic feedback. However, the imposed email policy drives algedonic signals upward to Systems 3 – 5, and could potentially overload higher order systems with information that has not been appropriately filtered. Additionally, after the ticket is escalated, the T1 analyst may or may not receive “sensory” feedback from systems or higher tier analysts that would allow them to learn, adapt, or adjust their response for future incidents.

Accountability Practices and Measures of Performance of System 1 Subsidiaries

Cybernetic feedback loops can facilitate accountability, which was perhaps the most varied point between these common System 1 components. In some teams, the T1 analysts were not even included in meetings with management, but yet did not have the autonomy to make decisions regarding incident escalation. In two of the three teams, the T1 analysts were hired on a contractual basis or something similar, working part time with no benefits and no path to development or tenure. This precedes the analysis to give some context regarding the organizational priorities (and perhaps maturity) around the Viable (or not) System of each location.

In one location, accountability of the T1 analysts was created through verbal shift handoffs, in which each analyst would present incidents that he or she worked on for that day. Other analysts could ask questions that the analyst would then have to answer regarding how the incident was handled. This direct cybernetic feedback creates accountability within the team, as well as a learning environment for newer analysts to get peer feedback on incident handling. The same location had more rigorous “quality control” checks of tickets created by shift. The shift lead would review these tickets for thoroughness (“completeness” was not always an indicator of quality, according to the managers), and a low-level manager would then audit the tickets for quality. This direct cybernetic feedback created some level of accountability at two levels of supervision to complete quality analysis of incidents at the T1 level. Other groups in this System 1 also had some similar practices, even some instances of parallel investigations by a novice and an expert, in which the expert would compare progress and outputs to their own, and mentor the novice accordingly.

The other two locations did not have robust accountability practices that could be observed, which was perhaps due to the lack of cybernetic feedback in the system. One location had leads that did audit tickets for completeness (checking that all fields filled out) but did not otherwise evaluate the quality of the ticket. Cybernetic feedback in this case was limited to omission errors in tickets, which reinforces the completeness metric as a primary measure for performance. This was possibly due to the fact that T1 analysts were only allowed to work on certain kinds of high-volume, routine tickets that did not usually require in depth analysis or even non-routine decision making.

Measures of performance were different in each group as well. At the first location, time was important. Time was the main measure of performance for T1 analysts to quickly conduct an investigation and perform remediation actions on incoming incidents. Quality was also an indicator of performance, but not measured or tracked formally. At the second location, there were not any explicitly described measures of performance or activities to create accountability, indicating a lack of cybernetic flow (by way of performance measures) to System 3, and thus no feedback to adjust the team accordingly. At the third location, the metrics that were tracked were actually to monitor the security performance of the larger organization and revolved around cost. That is, information included in the ticket was tied to cost indicators that became billable to each constituent organization. This included time of analysts as well as other remediation activities, such as reimaging a machine or dispatching a technician. Thus, analyst measures of performance did not really exist.

System 2 Diagnosis

System 2 is often described as the coordination mechanisms between System 1 subsidiaries. It reduces inter-operation disruption as they carry out a number of operational activities. The number of activities performed between System 1 subsidiaries was very high for each organization, and often resulted in some sort of electronic handoff. Enter System 2. System 2 is represented by the information systems shared by System 1 subsidiaries, and includes communication systems, ticketing systems, and any shared access of software or programs needed to perform respective tasks. Many of these teams have some degree of interdependence in order to fully execute the incident response process. For instance, the security operations subsidiary may detect, analyze, and decide what to do with an anomaly, then task IT with mitigation activities, such as reimaging a machine, retrieving a machine for forensic analysis, and so on. These coordinated tasks are assigned in a shared ticketing system. The security tasks (detect, analyze, decide) are managed in a separate system that the IT subsidiary does not access.

I note that within the security operations subsidiary are subsystems (tiered analysis groups) that also have coordination needs through information systems. Some of the weaknesses identified through observations around these coordination activities are described below. Furthermore, System 2 should also facilitate knowledge sharing (Hoverstadt, 2010) through shared knowledge management systems. This is a type of coordination between System 1 subsidiaries (and subsystems within them) that was greatly lacking or disconnected in all three teams.

Table.F.2. System 2 Comparison

SYSTEM 2	**Location 1**	**Location 2**	**Location 3**
Possible sources of oscillation or conflict between System 1 components	How task should be sent To whom to send task/incident Quality of T1 document Visibility between T1 and T2/other teams Different Ticketing systems	Who has the authority to act on an incident (system 1 component) Inconsistent visibility of larger network Different ticketing systems Recognition of other System 1 components as the same system	Who has the authority to act on an incident (system 1 component) Inconsistent visibility of larger network Different ticketing systems Recognition of other System 1 components as the same system
Elements of System 2 (Harmonizing or Damping)	Communication systems (H) Shared network monitoring (H) Wiki (H) Separate Ticketing Systems (D) Limited Access to other Support systems (D)	Communication systems (H) Separate Ticketing Systems (D) Limited Access to other Support systems (D)	Communication systems (H) Separate Ticketing Systems (D) Limited Access to other Support systems (D)
How System 2 is perceived (threatening or facilitating)	Facilitating for all (H) items Threatening for Access (D)	Facilitating for all (H) items Threatening for all else	Facilitating for all (H) items Threatening for all else

Oscillations and Conflict between System 1 Subsidiaries

There were multiple points of oscillation and/or conflict in each of the teams. The purpose of System 2 is to help manage these oscillations or conflict between System 1 subsidiaries. Interesting findings in Locations 2 and 3 indicate some potential issues affecting a viable system based on System 2 components.

One team's oscillations and conflicts were mainly around staying abreast of the same information, and ensuring the quality of that information. Shared visibility of incident investigations and relevant software was limited by some system separation, and was noted as a necessity in some organizations. Split ticketing systems between System 1 components were identified as a pain point, but again, had organizational justifications for separating them. The quality of incident handling activities was sometimes a source of conflict, but quickly handled within or between System 1 subsidiary management. I noted that there was sometimes conflict regarding to whom to send a task or incident two, which was mitigated by self-selection practices and protocols to determine how tasks should be passed between components.

The other two teams had much more substantial conflicts between the System 1 components. This included "who has the authority to act on an incident", as the number of System 1 subsidiaries was very high for the second and third organizations, and these groups reported through separate management streams. The system component in focus in the previous section was titled "central security operations", however, the centrality was in title only, and gives more insight into this particular conflict. Thus, the recognition by other System 1 subsidiaries that they were in the same System 1 as the central security operations team was absent.

Though visibility was listed as an issue for the first location, it was even worse in the other two, as they had little to no visibility into sub-networks managed by other System 1 subsidiaries. This led to inconsistent visibility of the whole network, and limited effectiveness in handling incidents within that. Finally, as seen in the industry team, there were different ticketing systems for different subsidiaries, leading to some visibility and information sharing issues at the System 2 level.

System 2 Elements

Harmonizing elements across the different locations were similar from one aspect: communication systems. Email, instant message, phone, and the ticketing system itself acted as information systems to support communication between System 1 components. This was present on all three teams. The industry team had and additional harmonizing element. This was shared network monitoring of operations. That is, all system components could view the same network information regarding incoming incidents, and the higher-tiered analysts could access the same programs as the lower-tiered incidents.

All three teams had a damping element of System 2: separate ticketing systems. As previously discussed, the need for separate ticketing systems was due to differences in credentials, but also differences in types of information needed by different teams. The security operations ticketing system at Tier 1 level needed a place to deposit routine information collected from multiple sources, sometimes with differing fields for different kinds of incidents. The IT ticketing system may have been too rigid for this kind of tracking. Additionally, security ticketing systems are designed specifically for security activities, making an IT ticketing solution improper for incident handling. One team had a completely separate ticketing system for T2 and higher, which was justified by a need for even more flexibility in terms of data, attachments, and search capabilities needed to manage the information. Lastly, limited access to all information systems was observed on all teams, due to the reasons listed above (credentials, clearance, and power). This acted as a damping element, as it created a funnel point between System 1 subsidiaries.

System 1 perceptions of System 2

Though not explicitly investigated as part of this research study, I was able to gain basic insights regarding the perception of System 2 on System 1. Jackson (M. C. Jackson, 2000) references the fact that these perceptions can often affect how much power System 2 has to assert itself for coordination purposes. In general, System 2 was seen as facilitating for all harmonizing elements, but threatening for all damping elements. Yet in relation to Jackson's cybernetics pathology, System 2 was not necessarily able to assert itself in all forms. From a human factors perspective, this assertiveness could be either salience (controlling how apparent an issue is to the user) or actual control (performing some activities related to incident response and coordination).

Email was the most widely used information system across teams. Within the tiered organization of Security Operations, email was an effective tool for communicating and coordinate between analysts. However, in the university and state government setting, email created an asynchronous communication pattern between Security Operations and local IT groups. Email is a common communication method across business in general making it a very crowded channel. Most locations used email filtering to sort inboxes into different categories, such as sender or subject line.

On interesting considering is that, automated organization of email content and perception of the receiver of the sender and alert can affect the overall salience of the incident, effectively transferring the perceptions of System 1 subsidiaries to the System 2 channel. A hypothetical example is if the System 1A subsidiary views the System 1B subsidiary in a negative light. When an email is received from System 1B, and automatically filtered into a specific folder for 1B issues, System 1A might be less likely to check that inbox, respond to the email quickly, or follow up with the System 1B subsidiary.

In many scenarios observed, System 2 had little control or assertiveness over the incident response process beyond pre-programmed salience of alerts. Due to the fact that System 2 capabilities currently do not include intelligence, most System 2 controls are managed by a system administrator or engineering group and seen as somewhat burdensome to upkeep in terms of time and resources. However, several participants referred to new tools they would like to see implemented that include machine intelligence to drive coordination activities in incident response. This technology, called *orchestration*, would make some decisions regarding potential actions to take for a given incident, or who should be assigned particular tasks. Orchestration is also able to amalgamate data from other information systems in System 2, effectively integrating sources for the System 1 users. Yet orchestration still requires some level of human administration, and may not be considered useful in small teams with limited resources (university team). The view or perceptions by analysts (System 1) on potential automation (System 2) varied by location, and could be dependent upon team maturity and access to relevant resources.

System 3 Diagnosis

System 3 can be described as line management (Hoverstadt, 2010) that oversees subsidiaries in System 1. System 3 is responsible for managing performance and resource allocation within each subsidiary, as well as creating synergy within System 1. Two of the three teams had less than synergistic relationships with other System 1 subsidiaries. They were not always helpful to each other, and did not have good communication practices or coordination protocols between them. Furthermore, the perception of the role of each subsidiary was distorted. Some subsidiaries seemed to view themselves as above others, which was not in the control of System 3 to mitigate, as the same problems seemed to persist amongst System 3 components and their affiliates, and even at the System 4 level. Based on Jackson’s (2000) pathology, the second team displayed evidence of a weak System 4 (considering that the System 4 level should be occupied by the CISO, but was not). Participants indicated that the current System 4 commonly prioritized politics over security risks, and was perhaps not able (or even willing) to incorporate external intelligence into strategic decisions in security. The CISO acting in a System 3 capacity is further evidence that the higher-level systems may be collapsing down to an operational level.

Performance is measured in the three locations differently and for different reasons. Resources are vastly different across them in terms of analyst staffing and information management systems. Locations 2 and 3 noted difficulties with staffing, especially at the T1 Lead/T2 level. Location 1 expressed a need for better information management systems, but no staffing issues. Some of these resource perceptions may be indicators of organizational maturity.

Table.F.3. System 3 Comparison

SYSTEM 3	**Location 1**	**Location 2**	**Location 3**
Components	Director of Security Operations Sector Information Security Officers (Liaisons)	CISO (acts as day-to-day manager sometimes)	CISO (acts as day-to-day manager sometimes)
How exercises authority	Performance expectations, through System 1 managers, accountability with Liaison Information Security Officers (by business segment)	Takes lead on issues; assumes analyst role with a higher level analyst	Takes lead on issues; assumes analyst role with a higher level analyst
How the resource bargaining with System 1 is carried out	Indications of needed resources would be noted and actions would be taken to alleviate pressure or stoppage by providing resources; In return for good performance and based on performance goals	None observed; Resources inherently limited – perception that resources were not prioritized	Staffing managed at System 3 level through contracting service; Other resources limited by expressed need from the larger organization; Performance not tracked as much as work completed
Who is responsible for the performance of the parts of System 1	Respective managers of each of the components are responsible	Respective managers of each of the components are responsible	Respective managers of each of the components are responsible
What audit inquiries (3) are conducted by System 3*	Ticket quality and overall performance tracked via quality audit and metrics tracking	None observed	Work completed acted as financial audit function to feed back to organization (bill services rendered)
How relationship between System 1 and System 3 perceived (autocratic or democratic) *How much freedom?*	Democratic – System 1 Management is given freedom to manage how they wish, as long as the job gets done; Freedom to interpret policy and enforce as needed.	Autocratic – very little freedom	Autocratic – very little freedom

Components of System 3

Based on the above description of what System 3 represents, a major difference was noted amongst the different System 3 components of each of the locations observed. One location had multiple components of System 3 that acted as control functions with security and IT operations. There were the respective directors over those aspects of operations, and additional liaisons in each of the business sectors to represent both sector and security respectively. This liaison position worked directly with the other directors to ensure that operations were steady and that policy was being communicated and followed.

The other two locations had similar setups, in which the true System 3 component was the Chief Information Security Officer (in title) that in reality held a managerial role over the respective operations. In both of these teams, the CISO reported to the CIO, who was ultimately responsible for operations, security thereof, and strategic direction of the organization. The conclusion that this position represents System 3 is supported by the fact that the CISO is covered on almost all incidents for these teams, and often steps in to help mediate issues between Systems 1 subsidiaries who don’t always consider themselves parts of the same system.

How System 3 exercises authority

Another major difference between the observed teams is in regards to authority. System 3 exercises authority differently between the first team and the other two. In the first team, the director meets regularly with System 1 managers to review aspects of performance, but generally allows these managers to be autonomous in daily operations. In the other two teams, the System 3 component actually takes the lead on some issues, assuming somewhat of a System 1 role. Authority is exercised through constraints of coverage on all events. This aspect was not thoroughly investigated, but presents some interesting indicators for future cybernetic work in System 3 and above.

Resource Bargaining

System 3 allocates resources to System 1 components in return for certain outputs. For instance, System 3 may make decisions regarding how many people should work in Security Operations, if the overall operation requires shift coverage, and which information systems (System 2) are necessary for adequate functioning and coordination. This resource allocation acts as an accountability function for System 1. Essentially, System 3 needs a certain level of performance from System 1, and determines the resources needed to do that. Should System 1 not perform effectively, resources are considered as a potential reason for that. In other settings, System 3 management might ask: Do they have the right tools? Do they have the right training? Do they have the right access? Do we have enough of the right people? Within the setting of incident response, the same questions could be asked for resource bargaining between management and System 1.

Resource bargaining was inconsistent across the three locations. In one location, System 3 closely tracked the performance of System 1, and in return was able to effectively provide all necessary resources (people, tools, training, development). In other locations, System 3 had a more distant relationship (both figuratively and literally) with System 1, and no resource bargaining was observed. That is, System 1 performance and accountability was not observed through the System 3 channel in the university, and there were no pending resources to be allocated in the state government (unless operations were to be expanded).

System 3 inputs*

Internal audit functions exist in most institutions as a way to create accountability within the larger organization. In the state government location, ticket completeness was expressed as a key indicator for ticket quality. Participants indicated that this was due to the fact that costs were tracked and invoiced throughout the larger organization. Thus, cost factors of incident handling needed to be closely tracked and billed to internal customers. Other locations did not talk about internal audit (or related functions).

Internal audit can be very financially driven, but is not the only channel by which performance can be tracked and managed. In fact, this function has been identified as potentially insufficient for performance accountability in cyber security (Kral, 2018). Though the function might exist in the other two locations, it may not be appropriate or adequate in identifying resource needs and performance shortcomings.

Nature of System 3 Relationship with System 1

The nature of the relationships between Systems 3 and Systems 1 at the three locations was also different. At the first location, the relationship was democratic. System 1 respective management was given freedom to manage the group how they deemed appropriate, as long as the job continued to get done. They had freedom in interpret policy and enforce when needed. System 3 did not force close oversight into System 1 operations. The other two locations were very autocratic. That is, System 3 components assumed most of the responsibility over these System 1 teams, which was evident in communication protocols and handling processes.

Systems 4 and 5 in SOCs

Access to Systems 4 and 5 was a notable issue in the data collection for this study. Thus, these systems were not included in the overall analysis. However, some aspects of these systems were observable such that I could identify Systems 4 and 5 in each organization. I did not employ the diagnostic questions for analyzing higher-level systems.

In general, the System 4 components across all locations were chief officers in some capacity. The major difference in this role was whether or not the CISO reported to the CIO, or was considered a peer in the organizational structure. Conflicts were observed at this level of the system, as the CIO could be responsible for both operations and security, which sometimes create tension from both an operational perspective (some security measures will limit operational capability) and from a cultural perspective (the CIO has very important internal customers who may not agree with changes in security practices). Issues regarding CISO strategic objectives (Hale, 2017) and importance in security organizations (Oltsik, 2017) has been identified in literature as factors affecting the current state of cyber security.

As the analysis extended beyond System 3, I noted that the priorities of the organization became clearer, as did tensions around those priorities and the manifestations of tensions at different levels of the organization. For instance, a CISO reporting to a CIO that always prioritizes operations over security may not be able to fulfill the security organization’s goal of creating a more secure network. In the future, the under-performing security organization may prompt funding to be cut, resulting in overwhelmed and overworked teams of skeleton crews and lower overall morale of the security team. These system dynamics offer some insight into other ways in which security organizations struggle in today’s landscape.

Jackson (2000) offers diagnostic questions to study Systems 4 and 5, as well as a pathology for weak systems. While not all of these questions could be answered within the scope of this study, some key characteristics of the teams studied are manifestations of these items. For instance, System 4 should be embodied in the CISO, who has deep understanding of the security landscape and experience in making strategic decisions for security operations. However, adjacent Systems

3 and 5 do not always recognize this intelligence function, resulting in a System 4 that does not (or cannot) provide environmental knowledge to System 5, and potentially becomes more involved in daily operations of System 3. This is evident in the representations of the second and third teams in which System 3 is identified as the CISO, and is directly contacted for incident response activities in System 1.

System 5 was difficult to capture from this study, and would need further validation from the participant organizations. Formally speaking, the contributors to policy making in the respective organizations included the Board of Directors, the VP executives over different functional areas of the company, or entire branches of government (Executive and Legislative). Though I offer no insight at this level of the system, it should be noted that, should System 4 suppress the internal needs of the systems below it or external threats from the environment, then System 5 may never get the opportunity to create policy addressing those needs and threats.

Appendix H: Study 2 IRB Package

APPLICATION NARRATIVE FORM

Purdue University, Institutional Review Board

1. Project Title: Determining System Requirements for Human-Machine Integration in Cyber Security Incident Response: Part 2

2. Principal Investigator: Barrett S. Caldwell, Professor, IE, bscaldwell@purdue.edu, +1 765 49-45412

Please address the following points regarding your proposed research:

A. PROPOSED RESEARCH RATIONALE

This expert interview study is being conducted to gain better understanding of expertise required to perform critical work tasks in Computer Security Incident Response (CSIR). Literature offers only a general view of knowledge, skills, and abilities needed by position, but not by process or task. The study will include interviewing experts who have experience with CSIR, focusing on information sharing between two people or two work teams. Researchers will be able to identify expertise needed to perform critical "handoffs", which can be used to improve current CSIR operations.

The research questions to be addressed in this study include:

- *What expertise is needed to perform specific information sharing tasks (e.g. incident handling team to law enforcement, incident handling team to forensics team, intrusion detection to incident handling team, etc.)?*
- *What kinds of errors are novices likely to make in these tasks?*
- *What does "a good day" look like with respect to these tasks?*
- *What does "a bad day" look like with respect to these tasks?*
- *What cues do experts identify and use when performing these tasks?*
- *What strategies to experts use to navigate specific scenarios during these tasks?*
- *Where do the teams use automation in information sharing activities?*

B. SPECIFIC PROCEDURES TO BE FOLLOWED

The expert interview consists of a 3-part methodology, called Applied Cognitive Task Analysis (Militello & Hutton, 1997). Appendix H.3 of this document includes the job aids from the methodology that will be used.

Part 1 – Task Diagram: *The researcher will ask the expert participant to break down a specific task (such as a handoff between the CSIR team and the forensics team) into 3-6 steps, and identify which of those steps requires knowledge or expertise to complete. Each identified step will then be further broken down into 3-6 sub-steps, iterating the same process as the first round. The result is very specific subtasks that require expertise to perform well.*

Part 2 – Knowledge Audit: *For each of the sub-steps identified in Part 1, the researcher will ask for details regarding what kind of expertise is needed to perform the task, specifically through examples from their own experiences. The audit includes probing questions that include:*

1. *Is there a time when you walked into the middle of a situation (regarding sub-step A) and knew exactly how things got there and where they were headed?*
2. *Can you give me an example of what is important about the Big Picture for this task? What are the major elements you have to know and keep track of?*
3. *Have you had experiences where part of a situation just 'popped' out at you; where you noticed things going on that others didn't catch? What is an example?*
4. *When you do this task, are there ways of working smart or accomplishing more with less – that you have found especially useful?*
5. *Can you think of an example when you have improvised in this task or noticed an opportunity to do something better?*
6. *Can you think of a time when you realized that you would need to change the way you were performing in order to get the job done?*

Part 3 – Simulation Interview: *The last part of this procedure is to present a specific scenario (e.g. a specific kind of incident, a given outcome, and a handoff about to be performed with a given entity) to the expert and have the person walk through, step-by-step, what he or she might do in that situation. The goal is to identify their problem-solving processes.*

C. SUBJECTS TO BE INCLUDED

For the study described, subjects are experienced individuals in Computer Security Incident Response (CSIR). The criteria for this study are not limited to any particular demographics.

The inclusion criteria for the subject population are:

- *The expert has more than 5 years of experience in Security Operations*
- *The expert has experience working in incident response, or experience managing it*
- *The expert does not need to be associated with the teams in Parts 1 and 3 of the research program*

The maximum expected number of individuals for this study is 20.

D. RECRUITMENT OF SUBJECTS AND OBTAINING INFORMED CONSENT

The researchers will by working with Purdue CERIAS to identify experts who are willing to participate. Language regarding the general study will be shared with CERIAS for distribution, after which, official recruitment language will be shared with individuals or companies that express interest. Personal contacts will also be utilized for recruitment.

The attached recruitment letters will be distributed (version depends on where the interview will take place: company site, Purdue, or other) through the CERIAS network. Participants can read and return signed consent forms directly to the researcher, or simply express interest via email to the researcher and fill out the consent form just before the in-person interview. As the interviews are solicited to individuals at companies who are interested, the risk of undue influence is mitigated, as participation is not tied to employment. This fact is stated in the Consent Form.

Some interviews may occur on the Purdue West Lafayette campus during the CERIAS Symposium, 3-4 April, 2018.

Should on-site access be needed, this will then be secured through documentation that will be submitted to the IRB. Investigators will present a form of consent (waiver requested for documenting) to the experts via email and in-person (if needed). After signed consent is obtained, the interview will commence, lasting about 2 hours per expert.

E. PROCEDURES FOR PAYMENT OF SUBJECTS

Subjects will not be compensated.

F. CONFIDENTIALITY

The researcher will capture general work history, such as years of experience and general progression, such as roles held (vulnerability analyst, incident responder, etc.), as well as general sectors in which the each subject has worked (healthcare, retail, government, defense, etc.) to give background context.

STORAGE:

Audio recordings of interviews will be transcribed in a timely manner (within 3 months) to minimize storage time. Only the investigators of the study will have access to these recordings. Recordings will stored on a secure Purdue storage platform, such as FileLocker or PURR until transcribed. After transcription, recordings will be destroyed. Transcriptions will be de-identified and kept indefinitely for future research.

As signed consent forms are required for this research, the storage of these signed forms will be limited to a locked file cabinet on Purdue's campus. No copies will be made.

MAINTAINING AND/OR DESTRUCTION:

[As above] Recordings of semi-structured interviews will not be digitally copied, and will be deleted from secure storage in a timely manner after transcription. Transcriptions will be de-identified, and kept indefinitely for future research.

G. POTENTIAL RISKS TO SUBJECTS

The potential risk to subjects is minimal, as they will be interviewed regarding their professional expertise in a particular area. A breach of confidentiality is a possible risk. As signed consent forms and audio recordings are identifiable records, safeguards will be included to mitigate this risk.

Signed consent forms will not be copied, and will be kept in a locked file cabinet on Purdue University's campus. Voice recordings will be securely stored on Purdue's secure digitial storage space that is approved by the IRB (FileLocker or PURR) until timely transcription. After this point, recordings will be deleted, and transcriptions will be de-identified and retained indefinitely for future research.

There is no expected need for medical or professional interventions based on the interviews being performed.

H. BENEFITS TO BE GAINED BY THE INDIVIDUAL AND/OR SOCIETY

There are no direct benefits to subjects.

The possible benefits to society include a better understanding of the expertise needed to perform specific steps in the CSIR process. By collecting valuable insights regarding knowledge needed to perform these tasks, the investigators will key requirements for potential customizable solutions *(versus the current generalizable solutions) for cyber security incident response. The overall goal is to help make these incident response processes more effective, and improve overall computer security operations.*

I. INVESTIGATOR'S EVALUATION OF THE RISK-BENEFIT RATIO

The probability and magnitude of possible harms are minimal for this interview study. The researchers are interested in individuals' expertise with respect to specific work functions, and no personally identifiable information will be collected.

J. WRITTEN INFORMED CONSENT FORM *(SEE ATTACHED)*

M. SUPPORTING DOCUMENTS *(check all document that you will be submitting to IRB)*

X__Recruitment advertisements, flyers, emails and letters.

X__Survey instruments, questionnaires, tests, debriefing information, etc.

X__Consent Form, Parental Permission, Assent Form

___Translated consent and recruitment documents

___If the research is a collaboration with another institution, that institution's IRB or ethical board approval for the research or request for IRB deferral.

___If the research accesses the PSYC 120 Subject pool include the description to be posted on the web-based
recruitment program and the debriefing form to be used.

___Local review approval or affirmation of appropriateness for international research.

___If the research will be conducted in schools, businesses or organizations, include a letter from an
appropriate administrator or official permitting the conduct of the research.

___If the study involves an investigational drug/device, include product information or investigator brochure

___Other (please list)

Appendix H.1: Study 2 Interview Consent Form

Purdue IRB Protocol #: 1802020208 - Expires on: 08-MAR-2019

RESEARCH PARTICIPANT CONSENT FORM

Observations of Functioning CSIRTs to Determine Workflow and Critical Handoffs
Barrett S. Caldwell, Ph.D
School of Industrial Engineering
Purdue University

What is the purpose of this study?

This study aims to better understand how information sharing activities (such as handoffs) occur during incident handling. More specifically, the researchers are interested in expertise required to do these tasks, and how a lack of expertise might affect the overall process. This study aims to do this through a semi-structured interview process, known as an Applied Cognitive Task Analysis.

You are being asked to participate because of your expertise regarding Computer Security Incident Response, whether as a past member of a team or a manager of one.

The study is aiming to collect observation data from 8-10 different experts with experience in incident handling functions within Computer Security Incident Response (CSIR).

What will I do if I choose to be in this study?

For this study, the researcher will engage you in an in-person interview that will last no more than 2 hours. There are three parts to the interview. The first part involves you drawing a diagram and denoting certain points while the interviewer listens, takes notes, and asks additional questions. Second part of the interview explores deeper knowledge regarding the elements in the diagram. Lastly, the third part includes a simulated task, which you will be asked to walk through, step-by-step, what you would do and why. Again, the interviewer may ask additional questions for more information.

Note that these research procedures are non-experimental, meaning there are no treatments being tested in this setting. The goal of the interview is to understand the deep knowledge you already have in a particular area.

Data collected includes the general order of events around a particular incident handling function, and where expertise is needed in those functions. The student researcher will be asking you questions, taking notes, and (audio) recording the interview for transcription later.

Please note that your knowledge is not being evaluated, but merely collected to help guide current research regarding less-experienced people who perform CSIR functions.

How long will I be in the study?

The student researcher will conduct the in-person interview over a maximum two-hour period. Some follow-up questions may be needed at a later time, but these would be short and limited to phone or email.

What are the possible risks or discomforts?

The risks for this interview protocol are minimal. The risks are no greater than you would encounter during routine meetings. The researchers are interested in what you know regarding specific functions in a professional setting, and how you know it. None of the questions are aimed at probing personal (non-work-related) experiences. Your decision to participate will not affect their job or relationship you're your employer, and your performance is not being evaluated.

There is a risk of breach of confidentiality, however, safeguards are in place to prevent this risk can be found in the Confidentiality section of the form.

Are there any potential benefits?

There are no anticipated direct benefits you personally. The general benefits to society may include a better understanding of CSIR functions and expertise needed to perform them. Specifically, managers and developers may be interested in understanding the knowledge you provide.

Will information about me and my participation be kept confidential?

Your voice will be recorded during the interview to help the researcher capture your responses. The recordings will be transcribed, then promptly destroyed. During this time, no one other than the research team will have access to these recordings, which will be securely stored on Purdue University's servers or repositories (PURR), or approved digital storage platforms, such as FileLocker. Physical signed consent forms will not be copied, and will be securely stored in a locked cabinet on Purdue's campus. These safeguards should mitigate risk of confidentiality breaches.

Transcriptions will be de-identified and retained indefinitely for use in future research.

Research records may be reviewed by departments at Purdue University responsible for regulatory and research oversight.

What are my rights if I take part in this study?

Your participation in this study is voluntary. You may choose not to participate or, if you agree to participate, you can withdraw your participation at any time without penalty or loss of benefits to which you are otherwise entitled.

If at any time you would like to withdraw from the study, please inform the researcher. Your decision to participate or not in the research will have no effect on the your relationship with your employer.

Who can I contact if I have questions about the study?

If you have questions, comments or concerns about this research project, you can talk to one of the researchers. Please contact the PI for the project, Dr. Barrett S. Caldwell (bscaldwell@purdue.edu, +1 765 494 5412) or the student investigator and primary point of contact, Megan Nyre-Yu (mnyre@purdue.edu , +1 224 622 9765)

If you have questions about your rights while taking part in the study or have concerns about the treatment of research participants, please call the Human Research Protection Program at (765) 494-5942, email (irb@purdue.edu)or write to:

Human Research Protection Program - Purdue University
Ernest C. Young Hall, Room 1032
155 S. Grant St.,
West Lafayette, IN 47907-2114

Documentation of Informed Consent

I have had the opportunity to read this consent form and have the research study explained. I have had the opportunity to ask questions about the research study, and my questions have been answered. I am prepared to participate in the research study described above. I will be offered a copy of this consent form after I sign it.

____ ________________________
Participant’s Signature Date

____
Participant’s Name

____ ________________________
Researcher’s Signature Date

Appendix H.2: Study 2 IRB Approval

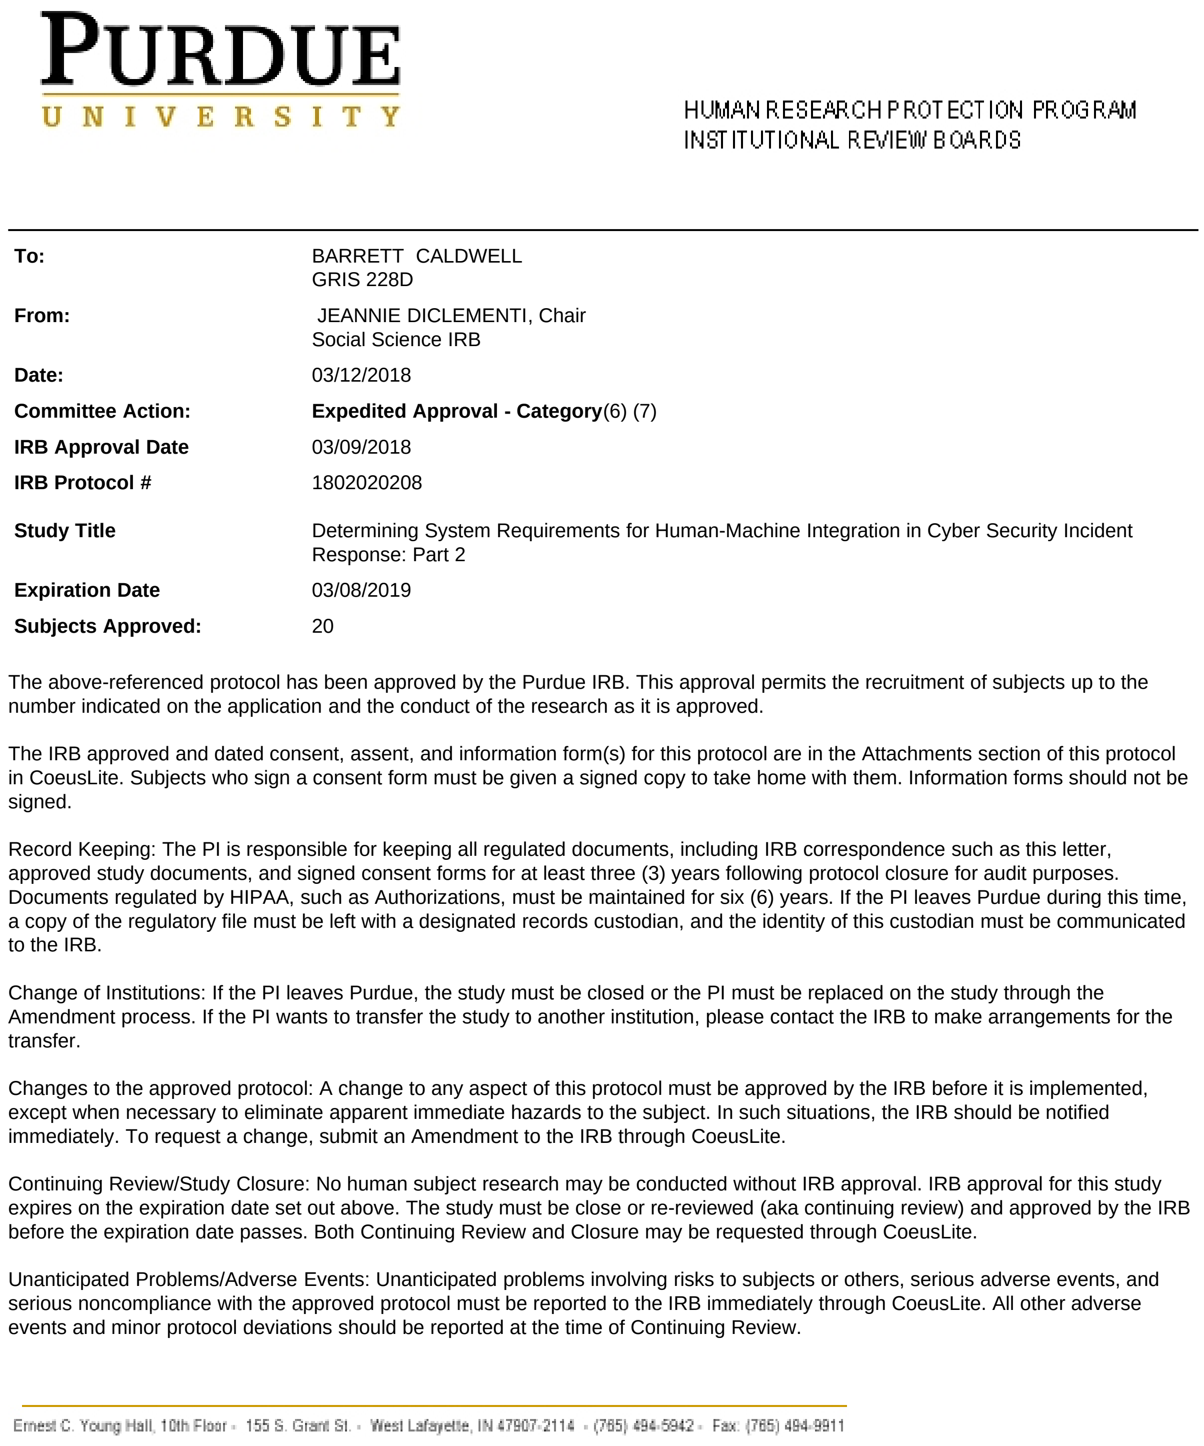

PURDUE UNIVERSITY

HUMAN RESEARCH PROTECTION PROGRAM
INSTITUTIONAL REVIEW BOARDS

To:	BARRETT CALDWELL GRIS 228D
From:	JEANNIE DICLEMENTI, Chair Social Science IRB
Date:	03/12/2018
Committee Action:	**Expedited Approval - Category**(6) (7)
IRB Approval Date	03/09/2018
IRB Protocol #	1802020208
Study Title	Determining System Requirements for Human-Machine Integration in Cyber Security Incident Response: Part 2
Expiration Date	03/08/2019
Subjects Approved:	20

The above-referenced protocol has been approved by the Purdue IRB. This approval permits the recruitment of subjects up to the number indicated on the application and the conduct of the research as it is approved.

The IRB approved and dated consent, assent, and information form(s) for this protocol are in the Attachments section of this protocol in CoeusLite. Subjects who sign a consent form must be given a signed copy to take home with them. Information forms should not be signed.

Record Keeping: The PI is responsible for keeping all regulated documents, including IRB correspondence such as this letter, approved study documents, and signed consent forms for at least three (3) years following protocol closure for audit purposes. Documents regulated by HIPAA, such as Authorizations, must be maintained for six (6) years. If the PI leaves Purdue during this time, a copy of the regulatory file must be left with a designated records custodian, and the identity of this custodian must be communicated to the IRB.

Change of Institutions: If the PI leaves Purdue, the study must be closed or the PI must be replaced on the study through the Amendment process. If the PI wants to transfer the study to another institution, please contact the IRB to make arrangements for the transfer.

Changes to the approved protocol: A change to any aspect of this protocol must be approved by the IRB before it is implemented, except when necessary to eliminate apparent immediate hazards to the subject. In such situations, the IRB should be notified immediately. To request a change, submit an Amendment to the IRB through CoeusLite.

Continuing Review/Study Closure: No human subject research may be conducted without IRB approval. IRB approval for this study expires on the expiration date set out above. The study must be close or re-reviewed (aka continuing review) and approved by the IRB before the expiration date passes. Both Continuing Review and Closure may be requested through CoeusLite.

Unanticipated Problems/Adverse Events: Unanticipated problems involving risks to subjects or others, serious adverse events, and serious noncompliance with the approved protocol must be reported to the IRB immediately through CoeusLite. All other adverse events and minor protocol deviations should be reported at the time of Continuing Review.

Ernest C. Young Hall, 10th Floor - 155 S. Grant St. - West Lafayette, IN 47907-2114 - (765) 494-5942 - Fax: (765) 494-9911

Appendix I: ACTA Interview Aids

The following interview aids are taken from (Militello et al., 1997).

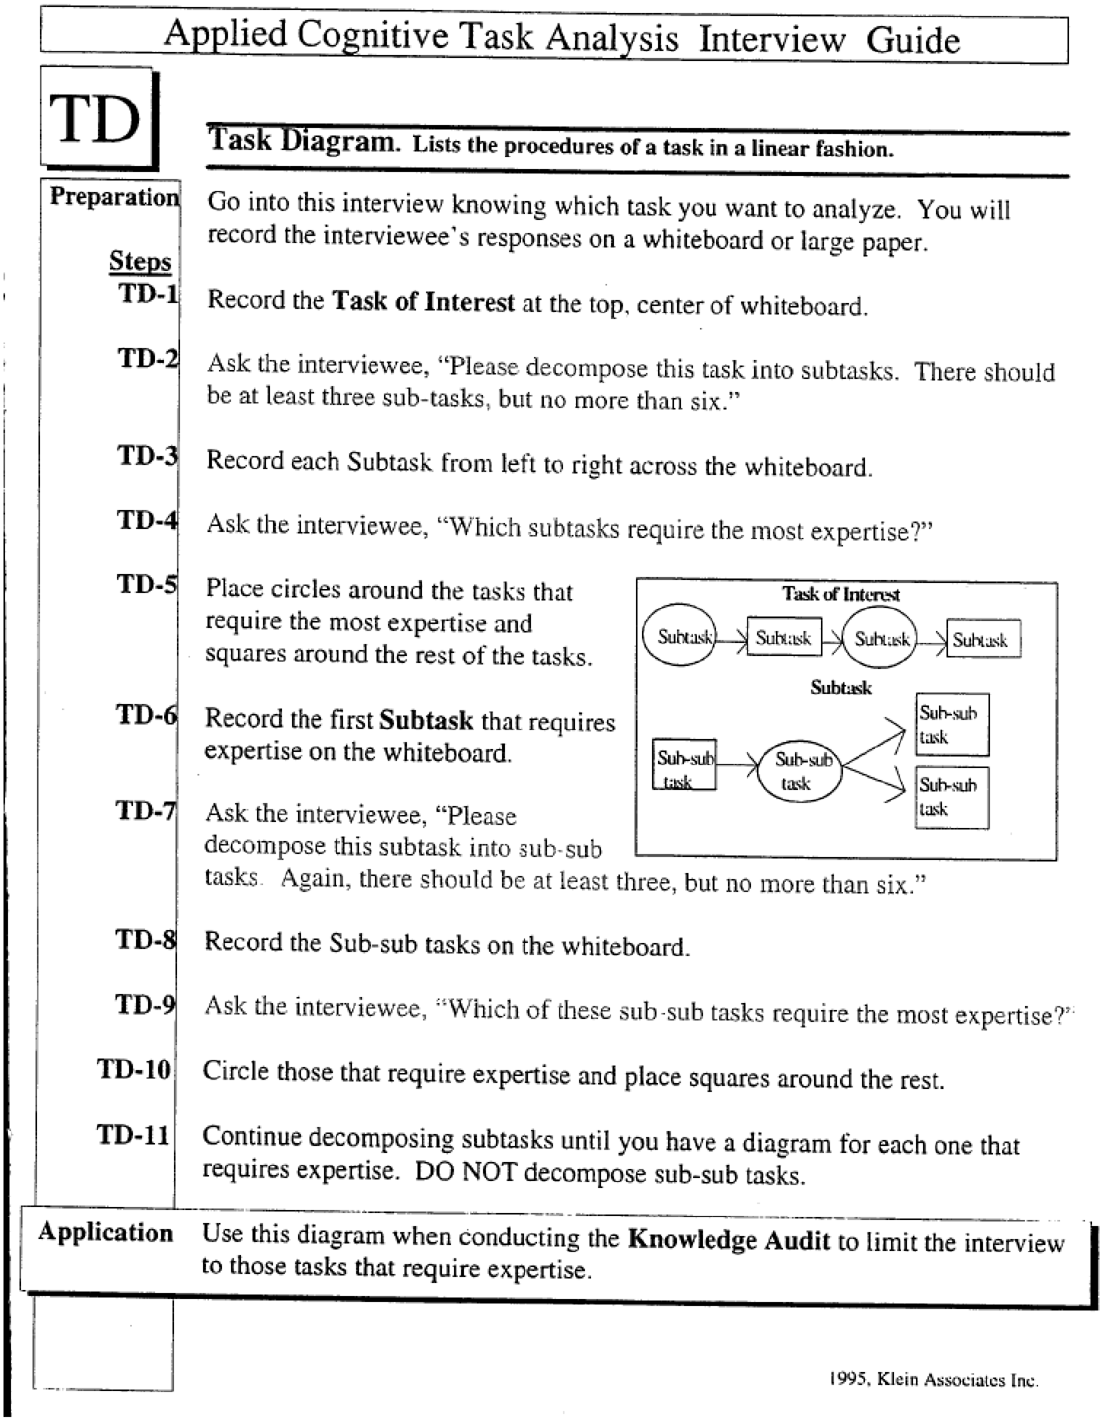

Applied Cognitive Task Analysis Interview Guide

TD

Task Diagram. **Lists the procedures of a task in a linear fashion.**

Preparation Go into this interview knowing which task you want to analyze. You will record the interviewee's responses on a whiteboard or large paper.

Steps

TD-1 Record the **Task of Interest** at the top, center of whiteboard.

TD-2 Ask the interviewee, "Please decompose this task into subtasks. There should be at least three sub-tasks, but no more than six."

TD-3 Record each Subtask from left to right across the whiteboard.

TD-4 Ask the interviewee, "Which subtasks require the most expertise?"

TD-5 Place circles around the tasks that require the most expertise and squares around the rest of the tasks.

TD-6 Record the first **Subtask** that requires expertise on the whiteboard.

TD-7 Ask the interviewee, "Please decompose this subtask into sub-sub tasks. Again, there should be at least three, but no more than six."

TD-8 Record the Sub-sub tasks on the whiteboard.

TD-9 Ask the interviewee, "Which of these sub-sub tasks require the most expertise?"

TD-10 Circle those that require expertise and place squares around the rest.

TD-11 Continue decomposing subtasks until you have a diagram for each one that requires expertise. DO NOT decompose sub-sub tasks.

Application Use this diagram when conducting the **Knowledge Audit** to limit the interview to those tasks that require expertise.

1995, Klein Associates Inc.

Applied Cognitive Task Analysis Interview Guide

KA

Knowledge Audit. Contrasts what experts know and novices don't.

Preparation

In the *Task Diagram* you identified the sub- and sub-sub tasks that require the most expertise. Go into this interview knowing the sub-tasks you want to analyze.

Steps

KA-1 Write the **Task of Interest** at top, center of whiteboard. Divide the remaining space into three columns with headings that match the illustration on the right.

Task of interest		
Example	Why Difficult	Cues & Strategies
1. Perceptual Skills *Example of perceptual skills*		
2. Anomaly *Example of Anomaly*		
3. Past & Future *Example.......*		

KA-2 Use the probes listed below to elicit examples of the various aspects of expertise. Record the first **example** in column one. Ask questions KA-3 and KA-4 before moving on to the next probe.

KA-3 For each example, ask, "Why is this task hard for novices or why don't novices know to do that?" Record answers in middle column under the heading **Why Difficult**.

KA-4 For each example, ask, "What cues or strategies do you use in this situation?" Record answers in third column under **Cues and Strategies**.

Expertise

Knowledge Audit Probes

- **Perceptual Skills** — Experts detect cues and patterns and make discriminations that novices can't see. Can you think of any examples here?
- **Anomaly** — Experts can notice when something unusual happens. They can quickly detect deviations. They also notice when something that should happen doesn't. Is this true here? Can you give me an example?
- **Past & Future** — Experts can guess how the current situation arose and they can anticipate how the current situation will evolve. Can you think of any instance in which this happened, either where experts were successful or novices fell short?
- **Big Picture** — If you were watching novices, how would you know that they don't have the big picture?
- **Tricks of the Trade** — Are there tricks of the trade that you use?
- **Improvising or Noticing Opportunities** — Can you recall a situation when you noticed that following the standard procedure wouldn't work? What did you do? Can you think of an example where the procedure would have worked but you saw that you could get more from the situation by taking a different action?
- **Self-monitoring & Adjustment** — Experts notice when their performance is sub-par, and can often figure out WHY that is happening (e.g., high workload, fatigue, boredom, distraction) in order to make adjustments. Can you think of any examples where you did this?

Optional Probes

- **Equipment** — Unless you're careful, the equipment can mislead you. Novices usually believe whatever the equipment says. Can you think of examples where you had to rely on experience to avoid being fooled by the equipment?
- **Scenario from Hell** — If you were going to give someone a scenario to teach someone humility--that this is a tough job--what would you put into that scenario? Did you ever have an experience that taught you humility in performing this job?

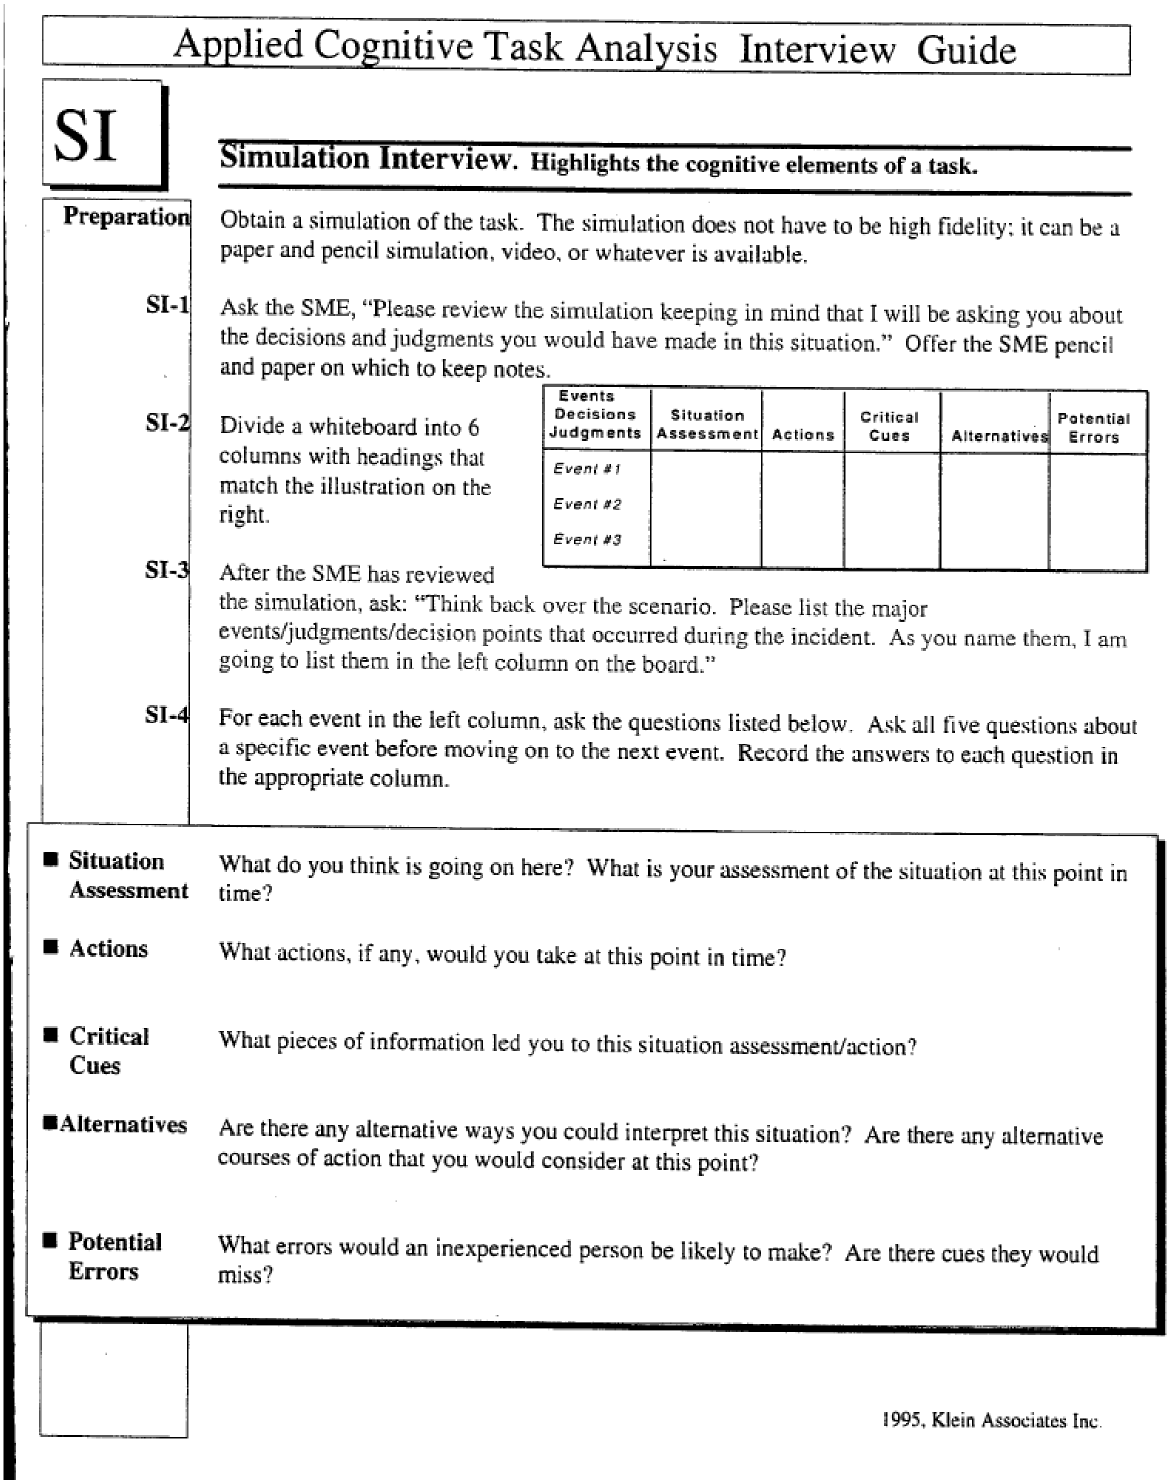

Applied Cognitive Task Analysis Interview Guide

SI

Simulation Interview. Highlights the cognitive elements of a task.

Preparation Obtain a simulation of the task. The simulation does not have to be high fidelity; it can be a paper and pencil simulation, video, or whatever is available.

SI-1 Ask the SME, "Please review the simulation keeping in mind that I will be asking you about the decisions and judgments you would have made in this situation." Offer the SME pencil and paper on which to keep notes.

SI-2 Divide a whiteboard into 6 columns with headings that match the illustration on the right.

Events Decisions Judgments	Situation Assessment	Actions	Critical Cues	Alternatives	Potential Errors
Event #1					
Event #2					
Event #3					

SI-3 After the SME has reviewed the simulation, ask: "Think back over the scenario. Please list the major events/judgments/decision points that occurred during the incident. As you name them, I am going to list them in the left column on the board."

SI-4 For each event in the left column, ask the questions listed below. Ask all five questions about a specific event before moving on to the next event. Record the answers to each question in the appropriate column.

- **Situation Assessment** What do you think is going on here? What is your assessment of the situation at this point in time?
- **Actions** What actions, if any, would you take at this point in time?
- **Critical Cues** What pieces of information led you to this situation assessment/action?
- **Alternatives** Are there any alternative ways you could interpret this situation? Are there any alternative courses of action that you would consider at this point?
- **Potential Errors** What errors would an inexperienced person be likely to make? Are there cues they would miss?

1995, Klein Associates Inc.

Appendix J: Study 2 CDTs

No.	Cognitive Demand	Why Difficult	Cues	Strategies	Potential Errors
1	Knowing whether or not I can handle a particular incident; Knowing if it is normal or not normal.	At the T1 level, there's not a lot of context provided when it comes across my desk. And as a T1, I may not know how to interpret that context anyways; "What is normal?" is a relative question to the environment and operation; "Nerd culture" drives analysts to try things that they shouldn't; I don't always know what is relevant or not relevant to judge	If it's something I've never seen before, then I likely can't handle it. If it's out of my domain, or it looks suspicious, then I probably should escalate it; I rely on my technology to tell me how bad it might be. But sometimes incidents are reported from users, and they're not detected by technology	Policies and procedures help me figure out what I'm supposed to do. I might rely on the previous analyst's assessment if it was handed to me at shift change; If I've seen this before, I might try to use past experience to figure out what I should do next, or if it is normal; Common sense; Collaborate with others in my org to figure out what I should do; Utilize my working relationships to supplement my lack of knowledge in a particular area; When escalating, give the receiving analyst all possible context that I can get, including weird things that are happening, what you've tried/looked at, etc.	I might escalate to the wrong person; I might escalate something that really isn't bad, but I thought it was; I might try to handle it and make a really big mistake, or make it worse; Potentially lose credibility in this kind of decision making.
2	Knowing how to balance the mission and the actions I'm supposed to take.	Sometimes the mission is to keep things running, not respond to a particular incident, even though we know it's happening.	If technology indicates that it's really bad, I would immediately escalate to someone who can better make this decision: to react or not.	I would escalate if I don't know how to do this; I would rely on policy to tell me what I should do; I might rely on a script to help me determine what I should do.	I could react to an incident that ends up taking down some portion of operations (or goes against the mission); Potentially lose credibility in this kind of decision making; Maybe the actions I took weren't appropriate, and what I did to "fix" the problem didn't actually fix it at all; Reinfection could happen.

3	Knowing what kinds of questions to ask, and what pieces of information are valuable	Procedures/policy can standardize the process to the point that I don’t get exposure to asking questions other than “What does the procedure tell me to do?”; I don’t have access to other kinds of context/background that I would need to determine this; I may not have the experience to know what kinds of questions are relevant; I may not have experience in other kinds of positions, like a network administrator, that would give me insights into what those relevant questions are; Knowledge of the architecture is helpful; Practical experience is key; My situation awareness (available context) is really limited with respect to a given incident, and with respect to how it affects other people; Limited ownership at T1 level.	Using the given context to drive what kinds of information I might need to get (if its [x] kind of incident, I should probably look into when it started, on what machine, did the user find it? Did the antivirus pick up on it? Etc); Using other descriptors of the incident to figure out who I might ask for help (It looks like it's malware, and I know [Joe] has experience in this based on my past interactions with him. I should ask [Joe]);	Escalating to someone who does have the access to relevant information; Collaborating with other analysts who may know better or have experience in this particular incident type; Rely on classroom knowledge, and take a learning approach to figuring out how to do it; Consult internal or external knowledge resources, like a wiki, or the internet, to help figure it out; Look at past incident reports (aka incident reviews) to see what was done; Continually educate myself when incidents happen in other companies, the reports are published, and I can learn from that; Ask for feedback; Reflect on my own performance, especially if I struggled to do something, so that I can improve the next time.	I could ask the wrong questions, getting irrelevant pieces of information to influence my decision; I could unknowingly omit information that is relevant to the action decisions; Potentially lose credibility in this kind of decision making.
4	Knowing who to go to for help (collaborate); or who to pass an incident to (escalate)	Organizations might be really big, or really small, where it’s not really clear where the knowledge exists that I need to do this job; Tiered organizations create more separation, and less familiarity at the bottom with who knows what; Working with different parts of the organization over time helps you know who is out there, and who does what	Who has worked on this before?; Who do I know who *might* have experience with this?; Does the org chart or procedure have points of contact?; Who is the most reliable, or has the best track record with me?; Who do I trust?	Standardized inboxes for different tiers funnel all escalated incidents into a single place for T2 analysts to claim; Collaborate with other analysts in my tier; Rely on past experience with particular individuals; Build camaraderie with other analysts, and utilize your network; I tend to escalate it directly to someone if I know them	Escalate to the wrong person; Delay in response; Potentially lose credibility in this kind of decision making.

5	Knowing how to interact with other analysts in other parts of the org or other tier levels	Collaborating through a lot of different channels can create some confusion or uncertainty, and also delay or frustration; Need to pick the most appropriate form of communication for the incident and the person on the receiving end; Collaboration may not be part of the enterprise model/culture	The urgency of the incident; Who do I need to interact with – what might be the most appropriate way of communicating what I need to tell them?; Do I need them to include their inputs in the process, or just get a quick answer?; How far away do they sit?; How much information do I need from the person, and can the channel support that?	If they're close, just talk to them in person; IM or call if they're in a different building; If you've never interacted with them before, learn from the first interaction; Interact with other analysts outside of the normal work context to get to know them and build trust; Communication is key, talking is necessary in some form; Develop good personal skills and communicate regularly with teammates about technical and non-technical things	Can cause peer annoyance or even distrust if the wrong method is used; Issues may not get handled in a timely fashion;

No.	**Cognitive Demand**	**Why Difficult**	**Cues**	**Strategies**	**Potential Errors**
1	Determining whether or not I can handle a particular incident (or if I should escalate it)	Sometimes an incident may seem straightforward, but it's really not; Sometimes its something that's not on the checklist; As a T1, I might not have the exposure needed to make this decision well/consistently, and to help build understanding of false positives, true positives, and different kinds of incidents	Have I handled something similar in the past?; Is there new data that comes in that changes the original decision for me to handle?; Am I getting stuck?	Use ML to filter stuff out that is a known issue; Use checklists to help me figure out suspicious activities versus non; Consult other analysts (including higher tier) to verify that it's something that should be escalated – this happens in less mature organizations; Exposure to more incidents is the best way to build this skill;	Not escalating soon enough; Doing the wrong response activities;
2	Determining what to do about a particular incident	Depends on the kind of incident, the available data, and knowledge of the systems and network; Also must know status of various systems, and how they work; Must know how things look without an anomaly to understand what an anomaly looks like; Need strong 'reference' point to be able to quickly identify anomalies;If the culture doesn't support collaboration, I may have less resources to help me figure this out	Baseline comparison of systems / status to the available data; Specific pieces of information will point to other skills (analysts) that might need to be involved	First, rely on protocols / checklists; Get as much information as possible: Identify the incident type, how bad it is, what systems are affected; As I get more experienced, assess network traffic and other heuristics instead of relying on protocol; Use knowledge of systems to figure out what skills are needed to address the issue; Use knowledge of organization to pull people in;	Not getting the right people involved; Delay in response; Doing the wrong response activities; Prioritizing the wrong incidents;

3	Determining if I’ve provided a quality response; or working efficiently	You don’t know what you don’t know; Feedback isn’t necessarily common; Traditionally, it takes time to learn shortcuts, maneuvers, scripting – time with the tools will help novices get through the learning curve	Get as much information as possible before deciding actions; No news is good news?; Get as much time in the seat as possible	Rely on feedback from peers, other analysts; Rely on knowledge from past performance to navigate the actions I take today; Ask for feedback; Rely on policy and procedures; Reflect often on what I could have done better	Wrong decisions will be repeated; Quality will not improve; Poor efficiency
4	Understanding larger context of an incident in order to guide decisions	May have limited visibility to other things going on, and thus missing context; Technology doesn’t always know what pieces of context are missing; Typically need hands-on experience; Tunnel vision is common amongst T1 level; A lot going on in short amount of time – new data can be overwhelming to a new analyst; May not have a strong understanding of correlations that exist between systems and activities; May not know some of the other controls in place	Get as much information as possible before deciding actions; Scope indicates severity; Are the affected systems critical business systems or not?; Is there a potential for this to affect critical business systems?; How more experienced people talk about a particular incident can hint at the severity; Compare knowledge of network to current state;	Try to get more pieces of the bigger picture by communicating with other analysts/personnel; Use other information, such as peoples’ reaction to it (and my knowledge of how they’ve reacted in the past, to help guide the personal response (if interfacing with an external entity to the SOC – like a client or other business org), such as assuring them or convincing them that its not critical, as well as the technological response	Doing the wrong response activities; Prioritizing the wrong incidents

5	Knowing how to interact with someone important during incident response (including clients)	Personal skills may not be inherent or developed; Knowing how a particular entity normally reacts to incidents/events; Need to maintain trust; Need the communication to be as salient as the severity; Need the communication to reflect the relationship with the person; Confidence in what I’m saying greatly affects this, and less experienced people tend to have less confidence; Environment may not facilitate or allow optimal (and bilateral) communication between analysts	Body language, voice intonation, general emotional response indicates how “amped up” a person is; Their relationship to my or my company	Compare past responses or reactions (and corresponding actual severities) to current situation; Do not include emotional response when giving guidance or feedback because people will react to it; Change the language being used to be more appropriate for that entity (use less technical terms when speaking to non-technical folks; summarize when talking to a c-suite person); Update the tone/language/mode if the person is not responding as needed; Use mentoring to help T1 analysts learn tricks of the trade to overcome environmental boundaries	If they are not reassured, lost trust; Might overreact and cause distress or distrust at the higher levels of management; Unable to build rapport, trust, and credibility with other team members
6	Using proper communication methods to coordinate in a team or with a client	Personal skills may not be inherent or developed; People react before they think, especially if they’re not in a coordinating role like a manager; Depends on the incident and the context; Need the communication to be as salient as the severity; Need the communication to reflect the relationship with the person; Confidence in what I’m saying greatly affects this, and less experienced people tend to have less confidence	Relationship with the person/entity; Severity of the incident; Resources needed (and in what time frame) for the incident; Type of response needed (and in what time frame) from person/entity (inform, acknowledge, or give more substantial feedback like a decision); Person/entity reaction	Change the mode depending on those cues (example, if high severity, use verbal communication like over the phone vs. email/IM/text); Develop, train to, and rely on policy for various situations;	Too many people try to get involved; Some people may get information they were not supposed to have; manager loses control over communication messages and channels; Resource management can be adversely affected;

7	Knowing who to involve in a particular incident	Must know who else in the organization knows how to deal with particular types of situations in the event that they are included in the response;
		Knowledge of organization may not be developed in T1;
			Baseline comparison of systems / status to the available data;
			Specific pieces of information will point to other skills (analysts) that might need to be involved;
				Use knowledge of systems to figure out what skills are needed to address the issue;
				Use knowledge of organization to pull people in;
				Use knowledge of past performance with particular analysts – or my relationship with them – to decide who should be pulled in;
				Knowing availability of analysts that have the skills needed
					Delay in response;
					Involving the wrong people/skills;
					Interruptions;

No.	Cognitive Demand	Why Difficult	Cues	Strategies	Potential Errors
1	Determining the priority of an incident	Automated tools don’t always help the analyst determine this; Novices do not have a deep understanding of the infrastructure and what needs to be protected; Novices may not understand the risk profile of the company	Compare the incident to the risk / risk profile of the company; What do my tools tell me?; Based on what I know, and all the information in front of me, does this seem like a high priority ticket?	Create policies, procedures, and standards based on the risk profile and business plan for novices to follow to help them with this; Develop or install automated tools to help them determine the priority; Escalate everything and skip this step altogether	Incorrectly determining priority (and following activities)
2	Determining if I have the skills or ability to handle an incident	Need a baseline understanding of the alert; Tools often help with this, but do not always; Need familiarity with the tools and environment; Takes time to develop this experience/expertise; Sometimes the potential actions needed can affect system operations, and politically can cause some issues/tensions, which are above the T1 paygrade	What do my tools tell me?; What do my procedures tell me to do?;	Rely on tools and procedures to tell me how I should answer this question (which are derived from business policies); Train procedures regularly so that people are aware and up to date on what the latest policy is; Escalate to someone who is allowed to make business decisions of the needed actions will cause waves	Trying to handle something that I shouldn’t (making a business decision that is not mine to make);

3	Determine who the incident should be escalated to	Need some knowledge of the individuals and skillsets above me; T1 tends to be segmented or separated; Remote work/physical separation affects team presence while working an incident	What kind of incident? Who has experience with this? Who has interest or responsibility over this? What’s the severity of the incident (in relation to other active incidents)? Who has availability?	Collaborate often with analysts at another level to get a better understanding of the skills and structure at that level; Gain awareness of assets outside the org that might be needed in IR; Gain a shared operating picture with those analysts who also are going to work on this; Escalate through the system > automatic bump up to a general level, where it is claimed by a relevant person	Send to the wrong person; Don't get to know the organization or why things get escalated to particular people (lack of gained experience)
4	What is the best way to escalate/share this? (Method/Mode/Style)	Need understanding of the alert in context to determine urgency/impact > drives method of escalation; How you communicate matters, and T1 analysts don’t always have the background information to properly communicate; Don’t always get feedback that the ticket was received	What is the alert? What is the context of the alert? Based on that, what is the impact to the business? Does that impact warrant high urgency alerting?; What mode does the receiver prefer?; How has this person reacted in the past?; Where is the receiver? What is his/her availability?; Who else needs to be covered on this message?	Use multiple methods of escalation (email, ticketing, text) if you think its important; Set policies for methods/modes; Give only the information that is relevant; Be concise for important people; Respect the other person’s perspective/background when using tone/choice of words; Only cover those who need to be covered; Gain shared awareness to share additional context with the receiver; Share insights with other analysts; Notify several people so there is not a single point of failure; Allow incidents to be “pulled up” by analysts who can better make business decisions	Ticket could be delayed, or missed; Ticket could go to the wrong person; Overreacting to the wrong people can damage your reputation, or their sensitivity to the issue

5	Determining what is normal or not normal	Need a baseline understanding of the alert; Novices do not have a deep understanding of the infrastructure/architecture – including fail safes; Novices may not understand the risk profile of the company; Baseline book knowledge only goes so far, the environment changes very fast; Novices think very linearly, and rely heavily on procedures or structured training/knowledge	What is the context of the incident?; What is the attack vector?; What are the vulnerabilities of the entity being attacked?; What kinds of attacks are predominant right now?; What is the system telling me versus what I'm seeing elsewhere? Does the logic line up?; Where is my information coming from? Is that source consistent/trustworthy?	Use tools to help identify patterns, correlations, etc; Use past experience to guide decision making; Stay up to date on latest news/threats; Update and retrain policy/procedures to ensure analysts have the most current information; Use checklists; Create a baseline; Collaborate with other more experienced analysts; Rehearsing procedures (not just training); Using common sense: If the system says one thing, but you see something else actually happening, that's suspicious; Don't overly trust your systems	Over reacting or under reacting; Responding to the wrong incident;
6	Collecting information from many sources and determining what is relevant or not relevant	Need a baseline understanding of the tools and how they work; Novices do not have a deep understanding of the infrastructure/architecture; Novices may not understand the risk profile of the company; Knowing where information lies (in terms of what needs to be protected); Can be very time-consuming	What is the context of the incident?; What is the attack vector?; What are the vulnerabilities of the entity being attacked?; What kinds of attacks are predominant right now?; Which data sources can I trust?	Utilize shared intelligence from external groups or entities; Utilize threat intelligence tools; Collaborate with other more experienced analysts; Have multiple analysts monitor the same information for shared awareness	Missing needed information to understand big picture or make correct decision regarding escalation or response

7	Understanding the “big picture”	Aptitude is required to do research type skills; Formal education only provides baseline knowledge and practice with analytical thinking; Applied practice is limited outside of an actual job; Applied practice within the job isn’t guaranteed on every kind of problem; “experts” in incident response may only ever work in the area in which they are most comfortable, creating less cross-coverage; Even experts brought in don’t have the context of the particular environment; Access to relevant information may be limited for T1	What is the context of the incident?; What is the attack vector?; What are the vulnerabilities of the entity being attacked?; What else is going on in the network that could possibly be connected to this? What kinds of attacks are predominant right now?; What do all my data sources tell me is happening?	Get as much information as possible; Work to understand the overall structure of the organization and network; Try to work on a lot of different kinds of problems to gain more diverse experience; Aptitude tests to get the right employees; Collaborate with higher level analysts, share information and insights; Have After Action Reviews as a team to review the incident	Make a wrong decision; Make a decision without the needed context; Counteracts other items in play accidentally; Treat the symptoms, not the root cause; Not understanding how your actions affect other people or components
8	Knowing your limits and if you’ve performed a task well	Takes self awareness and feedback; Sometimes we get caught up in what we’re doing that we lose sight of certain things; Novices may not know the effects of fatigue; Stressful environment that is based on excelling – people don't want to admit their weaknesses	No news is good news? Have the expectations been clearly communicated? Do my results indicate I’m doing okay? (response rate, relative impact)	After action reviews, lessons learned; Have a well-constructed team; Ask for feedback from other analysts, leads, or management; Make it a regular practice to give feedback	Mistakes can happen; Mistakes can be repeated

No.	Cognitive Demand	Why Difficult	Cues	Strategies	Potential Errors
1	Determine whether or not I can solve the problem	Assess not only ability of self, but also other context, like time of day, time of year (for retail operations) > compare expected time to resolve to a standard goal time; May not have subject matter expertise in the area where the problem is; May not know when to ask for help;	Have I solved this issue before? Have I solved in a short period of time?	Get rotational experience (no temporal pressure) working on these issues when not time sensitive > get exposure and experience that way; Send it to someone who has the right expertise; Use problem solving skills; Work to understand policy to answer “should I be solving this problem?”; In terms of asking for help, this requires “interacting with people” skills;	Hot issue doesn’t get taken care of in proper amount of time (has business implications);

2	Connecting capability and authority	Need knowledge of priority, capability/systems knowledge, and policy to determine course of action	Have I solved this issue before?
			Who has authority over this type of action?
			What do my documents tell me?
			Does the knowledge database have anything about this?
			How important is this issue?
				Ask a supervisor;
				Train to procedure;
				Train to new developments;
				Use automated systems to help guide self-service items
					May try to do something I don't have the authority to do;
					Waste time

3	Determine who to escalate to	Not all people have the same subject matter expertise; Need to send to the right person to make sure it gets handled;	What kind of problem is this? Who do I know has experience with this? Who has authority over this type of problem? What do the policies say? What does our documentation day?	Have new employees spend some amount of time with each SME so they understand who knows what; Have documentation to support directed communications to SMEs; Train policy so that people know the authority structure; Cover multiple people	Waste time; Incident might slip through the cracks if only one person is covered and they’re not available;

4	Determining all correct actions and executing them	I may not have access, system knowledge, past experience to help direct what I'm doing; Subject matter expertise may be weak in a particular area; Lower tier analysts don't always have time to sit back and reflect, also don't know bigger picture to put their actions into context	Have I seen this before? What did I do in the past? Am I familiar with this tool/system?	Follow provided diagnostic steps provided in documentation; Escalate; Cross-training/rotational experience to help build subject matter expertise; Reflect on past performance to figure out if I did okay, or if I should change something about my approach; Get a lot of hands on experience with different tools	May try to do something I don't have the authority to do; May do the wrong actions; May delay response

5	Determining noise or signal;
	Determining root cause;
	Knowing which data inputs are meaningful
		Need to know how the tool works, plus all the relevant information provided from the tool or user;
		Need to know structure/organization of systems to correlate data to potential causes;
		Novices are linear in their thinking;
		Novices may not understand larger system connections, or attributes of particular systems;
		Novices may not know how the tool works, in terms of what it shows or doesn't show, and assumptions behind it.
			Have I seen something similar?
			How do these data points connect?
			What systems affect each other?
			Could this be related to something else?
			Are there other points of data that would have been looked over by a novice that are not immediately visible/available?
			Are there other things going on in the network/system that could cause this issue?
			How does the technology behind this work?
				System architecture;
				Theory behind how certain technologies work (like internet protocols);
				Use problem-solving approach;
				Use similar situations from the past to guide investigation and decision
					Waste time checking everything when only one thing actually needed to be checked;
					Looking in the wrong place for information needed to solve the problem/ investigate

6	Communicating appropriately with other people/ stakeholders	Novices don’t always know the relationships between organizations to properly manage this communication; Sometimes policy drives this communication, and it might be above T1 level; System has a lot of noise (email), so it’s a channel that is convenient but not efficient; Novices don’t always understand the implications of an incident escalating in terms of legal standards, and may not appropriately adjust style and verbiage accordingly; May be lacking a general ability to communicate with people; May only follow procedure through email	How does the problem affect the stakeholders? Who has jurisdiction over the affected area/action? How do I effectively communicate the problem without causing an emotional response? How do I appropriately talk to someone who isn’t in my same organization (like a contractor)? Who needs to know about this incident?	Communicating up, leave out certain details (but not vital ones) to not overwhelm the executive; When correcting behavior, praise publicly and correct privately; Communicating out to legal, remain calm and use the proper communication channel; Use procedures to manage communication path, not necessarily style; Have a more experienced person facilitate communication with management or external bodies; Use standardized channels; Use more direct communication to head off an issue or de-escalate an issue; Cover multiple people on an incident;	Cause paranoia, over-escalation rate, angst; If from a novice, receiver may not take the communication seriously; If not directed through the proper channels, might slip through the cracks

7	Experimenting with different defense or mitigation techniques	Novices look at the immediate problem under current conditions, and have a hard time extrapolating that into the future/other conditions;
		Assuming that what you have is all you’ve got
			How could this get worse?
			How might it look/change next time?
			What did/didn’t work with similar incidents?
			How could different tools change the outcome?
				If current tools are too costly or painful, it's a good indication that something needs to change;
				Extrapolate into the future how this might affect us tomorrow with a more advanced attack, and assume less time to mitigate
					Wasting time;
					Wasting money;
					“PAIN”

No.	Cognitive Demand	Why Difficult	Cues	Strategies	Potential Errors
1	Gathering information about an alert by doing research and collecting context	Novices may not know where to look for information; May not have practice in investigation; May not have the aptitude for educating themselves; May not connect the dots; Novices may not know which systems can give them the information (or how that tool works); Need to manage multiple channels of information flow (chats, multiple databases, wikis, etc) to maintain awareness, but also help direct	Have I seen this before? What kind of information do I need to solve the problem or answer the next question? What’s happening in the larger network? How do all the technical pieces fit together? How could an attacker use the network in a malicious way such that it would result in what I’m seeing? What security risks or events have happened recently that I need to consider (what is temporally relevant?) What tools can I use externally (sites/google) that might have more info?	Use base knowledge of networks and security to navigate knowns/unknowns; Use deeper knowledge of particular systems to piece together what might be happening; Use all available resources (like knowledge databases) to help me figure out where I need to look or what info I need to collect; Use knowledge of attack strategies to determine what information I need to collect or check; Potentially collaborate with other analysts	Collecting information that is irrelevant; Not learning from the investigation process; May affect the next set of decisions (thinking you have enough of the right information to determine if malicious or not)
2	Do I have enough information to do my assessment?	Sometimes the information exists elsewhere (and T1 don’t know where that information is or who has it), and there are more steps to getting that information; There are some assumptions made about certain areas, and these assumptions can be wrong; May be stubborn and not recognize the need to ask others for more information; May not know who else on the team has the expertise to follow the strategies (right); Novices may thinking something is relevant when it’s not; May not be able to connect the dots in the larger picture	How certain am I that I have covered all the bases? Could I be wrong? Who would be able to validate this? Who already knows this? What kind of information do I need to solve the problem or answer the next question? (“atmospheric”/contextual or technique?) Do I trust my information sources? If it looks innocent, how could it not be innocent? Am I being tricked?	If someone else knows the answer to this, ask them; If I’m uncertain, validate with someone else, Potentially collaborate with other analysts; Use all available resources (like knowledge databases) to help me figure out where I need to look or what info I need to collect; If contextual information is needed from the client and no one here can answer it, escalate to client; If technique information is needed, ask someone internally who’s dealt with this before; Keep a database or document for reference to help find needed internal SME or external POC – kept up by shift so availability is taken into account; Use shared communications to broadcast an issue or question for fast answer from an individual who knows the answer	Not having the right information to determine the next step (if malicious or not); Making a decision that negatively impacts client operations;

3	Is this alert signal or noise? (False positive or true positive)	Novices may not have all the needed information to answer this; May not be able to connect the dots in the larger picture; Naiveté is an issue amongst novices – the don’t know how to be suspicious of information and trust too much; Novices tend to jump to conclusions based on limited experience; Novices may not have the technical knowledge related to technology, attacker techniques, and relevant procedures	Who already knows this? What does “normal” look like for this particular system/client? Is this a repeat event? How could an attacker use the network in a malicious way such that it would result in what I’m seeing? How could the existing data points be connected directly or indirectly that point to malicious intent? If it looks innocent, how could it not be innocent? Am I being tricked? With respect to the client, alert, and other context, is this normal activity?	Use knowledge databases and SOPs to help determine the answer and the following steps; Ask for help if unsure; Use knowledge of attack strategies to determine malicious or not; Discuss with other analysts how things could actually be malicious or connected even though they look innocent or disconnected – driven by actual conversations with other analysts to learn through argument/debate /opinion sharing; Validate information used to make this decision	Dismissing an alert that is actually malicious;
4	Can I handle this?	Novices may not have all the needed information to act; May not fully understand the threat; May not understand the bad actor/tactics/techniques; May not understand procedures; May be stubborn and not recognize need to ask for help; Need to be conscientious of own abilities	How much time do I have? Have I seen something similar? Have I been able to solve these problems in the past (what was my performance)? How sure am I of my assessment that I can handle or not?; What is within the scope of my job responsibilities;	Only take on tasks I know I can handle within a given time; When in doubt, escalate to someone you know can handle.; Can escalate to management to ensure proper resources are assigned (unwritten/unspoken policy to do this); Use feedback from supervisors to determine if I made the right decision in the past, or how I could improve	Trying to handle something not qualified for; Delay actual incident response; Making a decision that negatively impacts client operations;

5	Is the template I’m using to collect information (and decide upon) correct?	Takes time for novices to piece together what is in the template; Might over-trust template; There are a lot of specialties within IR, and the novice may not have knowledge in all of those subject matter areas; Novices may not know all client information/infrastructure to know when the template is correct or incorrect for that particular client	What does the procedure tell me to do? Who might know the client’s information or infrastructure to help me determine? What kind of issue is it with respect to the client, and does that affect what I need to do?	Create procedures/templates that are reviewed regularly to help reduce expertise needed at the T1 level regarding what should be done, in what tool, and who it needs to be sent to; Update knowledge base regularly > Continuous improvement processes	Waste time trying to piece together the process; Potentially miss some steps;

6	What is the best way (mode/method) to contact someone?
	What is the most appropriate style?
		Time zones may be different;
		language might be different;
		relationship with the person matters;
		Clients have different preferences or protocols for communication – need to have awareness of that;
			Who is the person? Do I know them? What is our relationship?;
			Is this an internal or external person?;
			Is there a way I can adjust my language to be more understandable?;
			Is the person technical or non-technical?
				Use phone for talking to a client and crossing technical-to-non-technical boundary;
				Use shared systems for communication; determine whether or not the conversation should be recorded/documented somehow;
				Don’t assume that the person has the same expertise as you;
				Don't assume shared knowledge/ awareness
				Don’t use acronyms with external clients;
					Confusion;
					Delayed response;

Appendix K: Study 2 Codebook

Codebook for Study 2

Note: Codes are not mutually exclusive, hence the limited exclusionary criteria for each code. Raters are encouraged to consider this fact when coding data, and use all relevant codes for a given data point.

Code	C1: SUBJECT MATTER EXPERTISE
Definition	Traditionally defined expertise in a given subject matter area; Usually related to a specific area, but can also be general; Pertaining to domain knowledge.
Examples	• Knowing what a concept means and how it is derived • Knowing the definitions of terms and which terms apply in the current state • Declarative knowledge of the topic being addressed • Knowing a set of facts or theory within a certain domain (e.g. networks, malware, hardware)
Exclusionary Criteria	Does not (by itself) include knowing how these things operate in a particular context

Code	C2: COMMUNICATION EXPERTISE
Definition	The style used to communication with someone; tactics for how they are approached; vocabulary used to communication something; using different styles for different people; being receptive of communication; Not limited to knowledge of channel or mode
Examples	• Knowing to use certain words (or avoid others) with particular people to express facts, thoughts, and opinions • Knowing which tone of voice to use • Knowing which style is the most appropriate for the person you’re talking to • Knowing how to give and receive feedback
Exclusionary Criteria	Does not focus on technical aspects of communication technology

Code	C3: INFORMATION FLOW PATH EXPERTISE
Definition	Concerning the method used to contact someone; Knowing which path is the most appropriate for a given person; Flexibility in exercising that evaluation / knowledge.
Examples	• Knowing which channel is the most appropriate to contact a specific person • Knowing which channels are available for time of day or place • Knowing which channel a specific person prefers • Knowing which channel is the most effective for a person/situation • Knowing when to change the channel based on situation needs
Exclusionary Criteria	Does not include communication style *(see Communication)*

Code	C4: EXPERT IDENTIFICATION EXPERTISE
Definition	Knowing who to go to when you need additional knowledge or expertise in a given area; Knowing who to send something, or who should address a given issue
Examples	• Knowing who can give you the information you need • Knowing where knowledge exists/can be found (which database, SOP, etc) • Social awareness, and ability to make social connections • Knowing who should receive an escalated incident • Determining who is best to ask about one of the other areas, but not others
Exclusionary Criteria	

Code	C5: INTERFACE/TOOL EXPERTISE
Definition	User skill in manipulating technological systems; Familiarity with tools and navigating interfaces
Examples	• Knowing the tools or programs that are relevant to the job • Knowing where to go (in the system) to find the needed information • Knowing the strengths and weaknesses of different tools • Knowing how the tools work and when to trust them • Knowing which tool is the most appropriate for a task
Exclusionary Criteria	Does not include general expertise relating to scripting or coding;

Code	C6: SITUATIONAL CONTEXT EXPERTISE
Definition	Knowing the environmental and situational context and how that affects the outcome
Examples	• Knowing which data sources should be combined to evaluate a decision point • Knowing what normal vs. not normal looks like given all data inputs • Knowing how the combined situational data can affect the system's performance • Knowing when data points should or should not be integrated in a given situation or to resolve a specific ticket
Exclusionary Criteria	(see Policy exclusions); Does not include general rules or laws that always or generically apply and/or are not time- or task-focused

Secondary Codes

Code	C7: POLICY
Definition	Institutionalized knowledge; Driven by rules or procedure that are developed by higher levels of management or company officials
Examples	• Understanding can be gained through company training • Rule-based determinations for how to perform tasks • Published and formally disseminated "standard operating procedures" • Access or clearance needed to execute a particular task or access a tool
Exclusionary Criteria	Does not include *personally-developed* rules or procedures; Does not include general elements of subject matter that do not change from organization to organization (ex: laws of physics, how malware works, etc)

Code	C8: SELF-AWARENESS; CONSCIENTIOUSNESS
Definition	Driven by understanding of self, including limitations and self-evaluation
Examples	• Knowing what you know and don't know • Knowing when to stop yourself • Evaluating your own performance • Executive function (knowing where you are in the task you are performing)
Exclusionary Criteria	Does not include direct feedback from others, or external performance indicators (though this direct feedback can instill self-awareness, if indicated that this is a input for the reflection process)

Appendix L: Study 2 Training Procedure

Training Protocol – MNY Dissertation: Study 2 Analysis

Qualitative Research Ethics:
The data to be analyzed was collected from human subjects, and is therefore subject to ethical considerations regarding protection of participant data. The researcher has taken precautions to remove all identifiable data through multiple stages of transcription and translation of raw interview data into tabular format. Though there are no identifiable markers in the data as it currently exists, it is important for you to recognized these facts and your responsibilities in terms of ethical research. In order to ensure your understanding, please complete CITI training for social research with Human Subjects (https://about.citiprogram.org/en/homepage/).

Qualitative Research Objectives:
The goal of qualitative research is to collect rich data from the field (real-world settings) around a specific phenomenon. It is "naturalistic, participatory, and interpretative" (Kerlinger & Lee, 2000, p. 589). The goal is to conduct true-to-life observation and use description to capture the phenomenon instead of hypothesis testing approaches for evaluating specific conditions or effects (Auerbach & Silverstein, 2003). Using a positivist approach, qualitative research involves recording details of these qualitative interactions generates knowledge (Pelto, 2016). Methods often involve various types of observations and interviews directly pertaining to a particular population or phenomenon, and rely on interpretation of the researchers to synthesize the data into useful findings.

The study which you are helping to analyze is centered around understanding what cyber security experts believe is necessary expertise for novice analysts to have in order to do Tier 1 – Tier 2 (novice to generalist) tasks in computer security incident response. The areas of expertise were established *a priori* to data collection (from literature (Garrett et al., 2009) and previous studies by the researcher), and used to help generate probing questions in addition to interview schedule used. The methodology for data collection closely followed an established interview protocol called Applied Cognitive Task Analysis (Militello et al., 1997), which was designed for less experienced interviewers to conduct a knowledge elicitation exercise. The methodology includes three (3) activities (Task Diagram, Knowledge Audit, and Simulation Interview) to build a Cognitive Demands Table (CDT) that summarizes the results of all three activities. The CDTs from each participant act as the data to be analyzed in this exercise.

The researcher conducted ACTA interviews with five (5) cyber security experts, and directly synthesized the interview activities into CDTs for analysis.

Trustworthiness:
In qualitative research, the researcher is considered a human instrument, and care must be taken when preparing and documenting how the research was conducted. One aspect of qualitative research is establishing trustworthiness in the data. Trustworthiness can be established by having multiple raters code data independently, and comparing the agreement between the coders (i.e. inter-rater reliability) (Goodell et al., 2016).

In order to improve consistency between raters, it is suggested (Goodell et al., 2016) that the lead researcher conduct training to calibrate raters. Therefore, included in this training are the following:

- Codebook of inclusionary and exclusionary criteria for the 8 codes being used
- Discussion around the codebook to ensure understanding between the two raters
- Example exercises, completed by both raters, to help establish consistency

After the above exercises, the researcher will review memo procedures for how she would like data recorded from your coding analysis.

Top-Down Coding Procedure

1. Sharing Data:
    a. Author will provide you access to a private cloud-based folder where the content for coding is located.
    b. Please download all five (5) documents.
2. Coding & Memo-ing: [INDIVIDUALLY]
    a. You may code electronically, or print out physically and manually code.
        i. ELECTRONIC: MS WORD
            1. Please use highlighting and commenting feature, and use coding number as reference (C1-C8)
                a. NOTE: Items can belong to more than one code. *Please denote all potential codes.*
            2. For your personal notes/reflections, please use commenting feature.
            3. Save document with your initials.
        ii. MANUAL:
            1. Print out all documents.
            2. Use highlighter or pen to denote code-able items, and use coding number as reference (C1 – C8).
            3. *If an item does not fit into a pre-defined code, please use *[#] to denote. These will be addressed later as findings.
3. Post-Processing:
    a. When complete, please scan/upload to shared folder.
    b. Researcher will fill out the electronic worksheet on Google Drive with your coding analysis.
    c. Researcher will do some additional analysis regarding inter-rater reliability.
4. Discussion between Raters:
    a. You and the researcher will meet to discuss your analysis.
    b. Any *[#] findings will be discussed individually.

Thematic Analysis: [TOGETHER]

5. For each code (C1-C8), you and the researcher will further code sub-categories, completing no more than 2 iterations of coding.
    a. This will be done for all items in the codebook, and then for the *[#] items. (C7-C8 included here)
    b. After the above has been completed, you and the researcher will develop themes within each code category.
6. Post-Analysis:
    a. Please shred/electronically destroy any data pertaining to this study.

Appendix M: CDT Coding Example

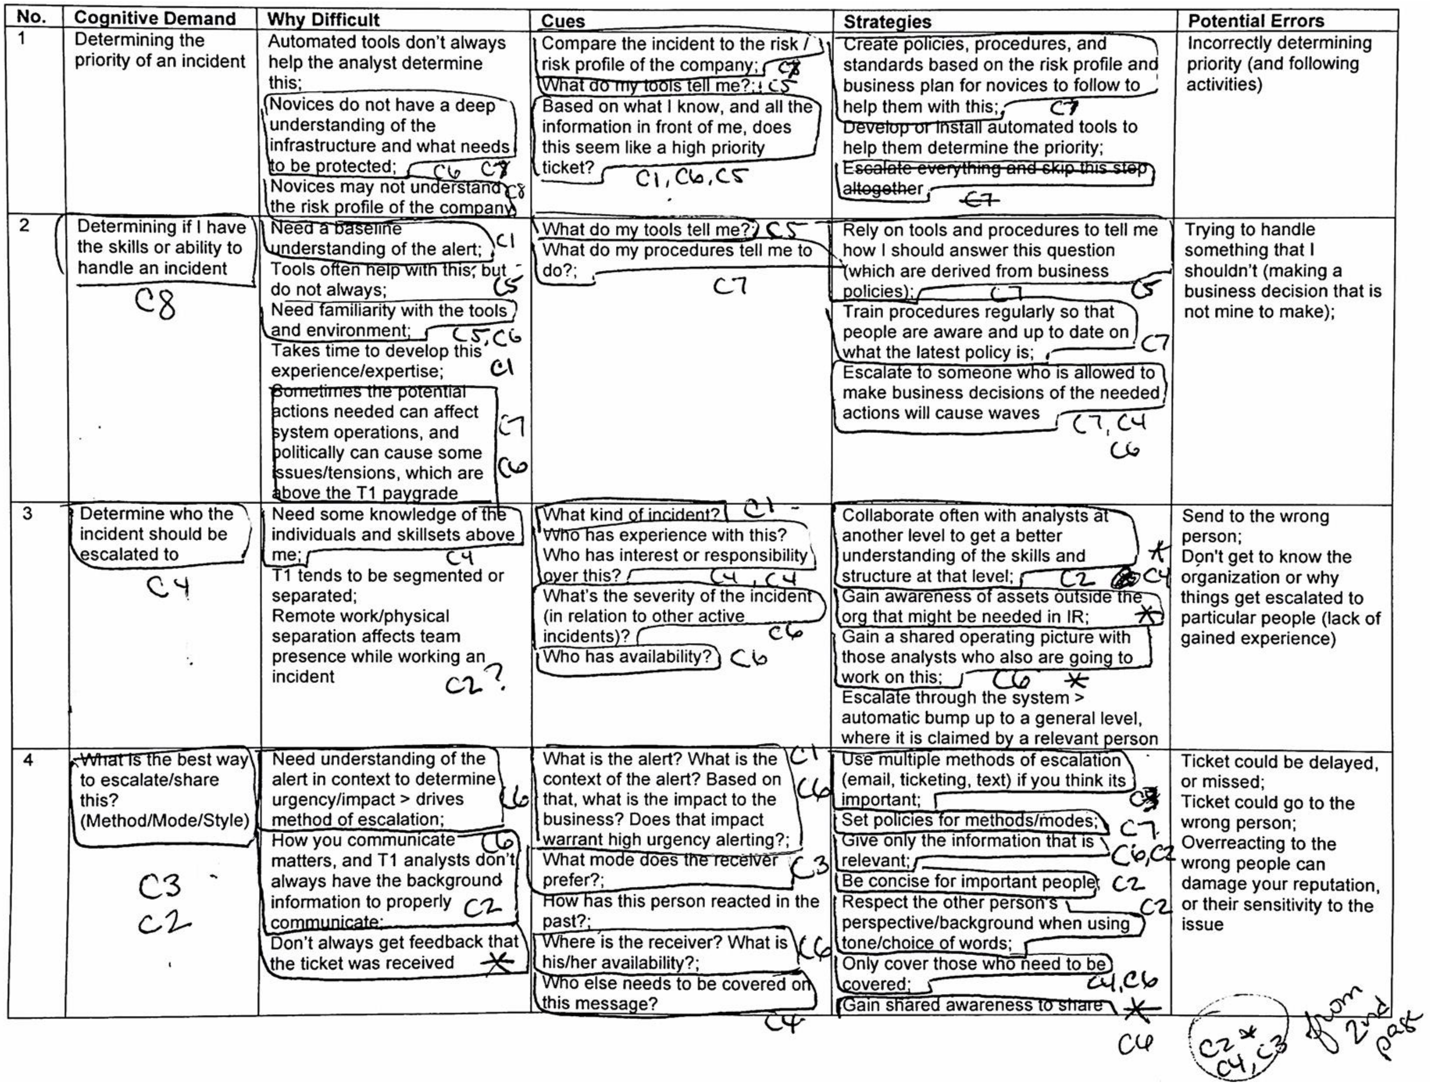

No.	Cognitive Demand	Why Difficult	Cues	Strategies	Potential Errors
1	Determining the priority of an incident	Automated tools don't always help the analyst determine this; Novices do not have a deep understanding of the infrastructure and what needs to be protected; C6 C7 Novices may not understand the risk profile of the company C8	Compare the incident to the risk / risk profile of the company; C8 What do my tools tell me?; C5 Based on what I know, and all the information in front of me, does this seem like a high priority ticket? C1, C6, C5	Create policies, procedures, and standards based on the risk profile and business plan for novices to follow to help them with this; C7 Develop or install automated tools to help them determine the priority; ~~Escalate everything and skip this step altogether~~ ~~C7~~	Incorrectly determining priority (and following activities)
2	Determining if I have the skills or ability to handle an incident C8	Need a baseline understanding of the alert; C1 Tools often help with this; but do not always; C5 Need familiarity with the tools and environment; C5, C6 Takes time to develop this experience/expertise; C1 Sometimes the potential actions needed can affect system operations, and politically can cause some issues/tensions, which are above the T1 paygrade C7 C6	What do my tools tell me? C5 What do my procedures tell me to do?; C7	Rely on tools and procedures to tell me how I should answer this question (which are derived from business policies); C7 C5 Train procedures regularly so that people are aware and up to date on what the latest policy is; C7 Escalate to someone who is allowed to make business decisions of the needed actions will cause waves C7, C4 C6	Trying to handle something that I shouldn't (making a business decision that is not mine to make);
3	Determine who the incident should be escalated to C4	Need some knowledge of the individuals and skillsets above me; C4 T1 tends to be segmented or separated; Remote work/physical separation affects team presence while working an incident C2?	What kind of incident? C1 Who has experience with this? Who has interest or responsibility over this? C4, C4 What's the severity of the incident (in relation to other active incidents)? C6 Who has availability? C6	Collaborate often with analysts at another level to get a better understanding of the skills and structure at that level; C2 C4 * Gain awareness of assets outside the org that might be needed in IR; * Gain a shared operating picture with those analysts who also are going to work on this; C6 * Escalate through the system > automatic bump up to a general level, where it is claimed by a relevant person	Send to the wrong person; Don't get to know the organization or why things get escalated to particular people (lack of gained experience)
4	What is the best way to escalate/share this? (Method/Mode/Style) C3 C2	Need understanding of the alert in context to determine urgency/impact > drives method of escalation; C6 How you communicate matters, and T1 analysts don't always have the background information to properly communicate; C6 C2 Don't always get feedback that the ticket was received *	What is the alert? What is the context of the alert? Based on that, what is the impact to the business? Does that impact warrant high urgency alerting?; C1 C6 What mode does the receiver prefer?; C3 How has this person reacted in the past?; Where is the receiver? What is his/her availability?; C6 Who else needs to be covered on this message? C4	Use multiple methods of escalation (email, ticketing, text) if you think its important; C3 Set policies for methods/modes; C7 Give only the information that is relevant; C6, C2 Be concise for important people; C2 Respect the other person's perspective/background when using tone/choice of words; C2 Only cover those who need to be covered; C4, C6 Gain shared awareness to share * C4	Ticket could be delayed, or missed; Ticket could go to the wrong person; Overreacting to the wrong people can damage your reputation, or their sensitivity to the issue

C2 * C4, C3 from 2nd pass

Appendix N: Tally Example

Phrase	C1	C2	C3	C4	C5	C6	C7	C8	*	AGREED TERM
Knowing whether or not I can handle a particular incident								2	1	Readiness
Knowing if it is normal or not normal.	2					1				
At the T1 level, there's not a lot of context provided when it comes across my desk. And as a T1, I may not know how to interpret that context anyways						2				
"What is normal?" is a relative question to the environment and operation;						1				
"Nerd culture" drives analysts to try things that they shouldn't;								1		
I don't always know what is relevant or not relevant to judge						2				
If it's something I've never seen before, then I likely can't handle it.								2		
If it's out of my domain, or it looks suspicious, then I probably should escalate it;	1							1		
I rely on my technology to tell me how bad it might be.					2					
But sometimes incidents are reported from users, and they're not detected by technology					1					
Policies and procedures help me figure out what I'm supposed to do.							2			
I might rely on the previous analyst's assessment if it was handed to me at shift change;									2	Person Trust
If I've seen this before, I might try to use past experience to figure out what I should do next, or if it is normal;	2									
Collaborate with others in my org to figure out what I should do;		2		1					1	Collaborative Problem Solving
Utilize my working relationships to supplement my lack of knowledge in a particular area;		1		2				1		
When escalating, give the receiving analyst all possible context that I can get, including weird things that are happening, what you've tried/looked at, etc.		1				2				
Knowing how to balance the mission and the actions I'm supposed to take						1	2	1		
Sometimes the mission is to keep things running, not respond to a particular incident, even though we know it's happening.							1			
If technology indicates that it's really bad, I would immediately escalate to someone who can better make this decision: to react or not.				2	2					
I would escalate if I don't know how to do this;								1		
I would rely on policy to tell me what I should do;							2			
I might rely on a script to help me determine what I should do.	1						1			
Procedures/policy can standardize the process to the point that I don't get exposure to asking questions other than "What does the procedure tell me to do?";							2			

Appendix O: SOAR Features Matrix

Technology	Capability	Features	Sit. Context	Subject Matter	Policy	Expert ID	Interface/Tool	Communication	Self-Awareness	Info Flow Path
CyberSponse	Role Based Dashboards	Someone logged into a particular role sees only what is useful to them			x					
		Can be designed by someone with more experience/expertise	x	x	x					
		Integrate information streams	x			x				
	Playbooks	Determine paths or flows			x	x				x
		Integrate information streams	x			x				
		Global monitoring of playbooks	x							
	Multi-Tenancy	Allow filtering of information by, to, and from customer	x		x			x		
		Integrate information streams	x			x	x			
	Incident Management	Create team roles and hierarchies			x	x				
		Control information presentation by role			x					
		Cross-linking modules for analysts reviews	x			x	x			
	Metrics & Reporting	Role-based reporting (assigning)	x		x					
		Incident or user-based reporting	x						x	
	Queue Management	Auto-assignment of incident to analyst			x	x				
		Create custom queues and assign members to monitor	x		x	x				
Demisto	Playbooks	Determine paths or flows				x				x
		Real-time workplan review	x							x
		Codeless playbook creation		x			x			
	Incident Management	Incident repository / knowledge database	x		x	x				
		Evidence board for information presentation during investigation	x							
		Multi-tenancy (data segregation by role)	x		x			x		

		Unified platform - integrated technologies	x				x			
		Incident or user-based reporting / Analyst tracking							x	
		Auto-documentation of incident activities			x					
	Interactive Investigation	Virtual "war room" / ChatOps	x	x		x	x			x
		Correlations & Related Incidents	x							
		Machine Learning Chatbot				x		x	x	
Siemplify	Interactive Investigation	Context /Enrichment / Automated data gathering	x			x				
		Cyber ontologies for making new connections	x			x				
	Case Management	Group potentially related incidents into a case	x							
		Automated Prioritization	x		x					
	Playbooks	Drag-and-drop playbook creation		x			x			
	O/A Platform	Shared workbench between analysts	x	x		x				
		Cross-functional 'war room'	x	x		x	x			x
		Dynamic levels of automation								
		Automated case assignments / escalations	x		x	x				
Swimlane	Security Orchestration	Orchestrate threat management across disparate platforms	x			x	x			
		Collect and consolidate all relevant alarm and event data	x			x				
		Automatically initiate actions on any third-party system			x	x				x
	Playbook/ Workflow Automation	Standardize IR process within single platform			x		x			
		Build with expert logic	x	x		x				x
	Adaptable Case Management	Access highly contextualized incident data in a single interface	x				x			
		Enforce process standardization and compliance			x					
		Dynamic levels of automation								
	Metrics & Reporting	Granular reporting of performance	x						x	
Phantom	Information Aggregation	Unified platform - integrated technologies	x				x			
		Flexible data sources and flows (push/pull) to aggregator	x				x			
		Enforce policy decisions			x					
	Playbooks	Visual playbook editor		x			x			

		Auditable response actions (documentation)			x					
		Canned playbook actions in API		x						
	Mission Control	Combined screens into single dashboard	x				x			
		Access to event history, contextual info, and interactive data views	x				x			
	Mission Guidance	Intelligent assistant to offer suggestions for learning / validation	x	x		x			x	
	Threat Intelligence	Query contextual information	x				x			
	Activity Feed	All current activities displayed	x				x			
		Allow multiple analysts to act on the same incident	x							x
	Case Management	Map tasks to SOPs	x		x					
		Pull in incidents into a single case	x							
		Case templates with industry standards			x					
D3 Soar	Security Orchestration	Determine paths or flows for tasks				x				x
		Allow users to approve / apply unique expertise		x		x				
		Integrate technologies	x				x			
	Playbooks	Drag-and-drop playbook creation		x			x			
	Automation	Flexible scripting / Easy-to-use APIs					x			
		Full lifecycle automation for certain threats	x	x			x			x
	Incident Response	Standard-based playbooks w/ expert input; Customizable		x		x				
		Dynamic workflows; integrate stakeholders	x			x				x
	Case Management	Timeline and link analysis	x							
		Audit logs, chain of custody, sign in/out logs			x					
		Role-based access controls			x					
	Reporting & Dashboards	User-based dashboard control	x							
		Executive reporting	x		x			x		
LogRhythm	SmartResponse Automated IR	Semi-automated, approval-based operation	x	x						
		Automatically initiate actions by incident type	x	x		x				x
		Canned playbook actions	x	x			x			
		Flexible scripting / Easy-to-use APIs					x			
		Integrate technologies	x				x			
		Sophisticated approval scenarios			x	x				x

		Full-chain execution	x				x			
		Remote execution (3rd party system)	x				x			
		Audit logs			x					
	Case Management	Centralize case evidence			x					
		Leverage different skillsets by task		x		x				x
Syncurity	IR-Flow (product)	Unified platform - integrated technologies	x				x			
		Data enrichment (context from multiple tools)	x			x	x			
		Auto rank / identify critical alerts		x	x					
		Filter false positives	x	x						
		Users validate what IR-Flow is unsure of	x		x	x				
		Consistent, repeatable workflows			x					x
		Codify best practices / policy / procedure into workflows	x	x	x	x				x
		Custom playbooks by incident type	x	x		x				
		Auditable system of record			x					
IBM Resilient	Automated response	Automated triage	x	x						
		Automated data enrichment	x	x			x			
		Integrate technologies	x				x			
		Help prioritize incidents for analysts		x	x					
	Playbooks	Guide analysts through response; guide with procedure and timelines		x	x		x			
		Determine analyst role/responsibility; Right analyst for the job	x		x	x				
		Documentation throughout the response process			x					
		Codeless playbook creation		x			x			
		Prompt post-incident review							x	
	Collaboration	Enable centralized communication	x	x		x	x			x
		Identify when people outside the SOC need to be involved	x		x	x				
	Privacy module	Navigate regulation and policy (internal/external) around incident		x	x					
		Provides breach response plans based on legal expert advice		x	x					

Appendix P: Concept of Operations (CONOPS) Document for SOAR 2.0

1. Scope

The purpose of this document is to provide a high-level view of user needs and expectations of SOAR platforms, derived from Studies 1 and 2 of this dissertation. The CONOPS is meant to later guide requirements development, but does not explicitly define them here. The IEEE guide for CONOPS was used to drive components of the document (*IEEE Guide for Information Technology - System Definition - Concept of Operations (ConOps) Document (1362-1998)*, 1998). Other articles and standards were used to supplement certain sections, which are cited accordingly.

2. Summary of Current System

The current system is defined as the aggregate of SOAR technologies currently available on the market. A table of features included in current SOAR platforms can be found in Appendix O. The major functional components of SOAR include (Neiva et al., 2017):

1. Orchestration, which includes integration of different security technologies such that operations can flow across them
2. Automation, which executes tasks independent of a human operator
3. Incident management & collaboration, or tracking an incident from detection to resolution by multiple human parties
4. Dashboards & reporting, which includes user interfaces and data collection and aggregation for reporting and audit

These major system components are not independent of each other. Orchestration allows automation to execute tasks without interruption; automation allows for faster incident management; and dashboards encourage collaboration by providing shared, analyzed data. From what the researcher can ascertain from SOAR documentation, external policies and procedure are embedded into the coded procedures or visualized ontologies within the system. The platform maintains a link to the manufacturer, such that the manufacturer can push updates to each customer. The researcher also points out that SOAR platforms are rapidly growing in capability, and that interfaces to external systems or procedures may be currently in development, or part of customization of the platform to a client.

A list of capabilities, functions, and features of the current system can be found in Appendix O, and further reading on general SOAR functions and features (from which the list was derived) include (Bedell, 2019; D3 Security, 2019; Demisto, 2018; Foroushani, n.d.; IBM, 2017; LogRhythm, 2019; Neiva et al., 2017; Siemplify, 2018, 2019). Current platforms aim to increase the overall efficiency of incident response, as well as reliability (in creating more standardized response protocols), and auditability. The main two metrics that many organizations prioritize in terms of security response are meant time to detect and mean time to respond. Literature about SOAR points out needs in scalability (C. Brooks, 2018; Oltsik, 2018a), but it is unclear how well current platforms meet this need. Many SOAR platforms require some sort of maintainability from people who can create, modify, or overhaul protocols in code, but advertise that drag-and-drop capabilities have been developed to overcome some level of coding expertise.

The researcher has not validated aspects of usability of the platforms with actual users, in part because many of these platforms are new, and have relatively low adoption compared to the size of the market. To that point, the researcher has not interacted with current users of SOAR, but did interview managerial personnel who were prospective buyers of a platform. The one manager who explicitly mentioned SOAR platforms by name did so with the emphasis on SOAR as a solution to increase the maturity of the team, allowing lower-tier incidents and filtering to be completed by the technology, and allowing analysts to work on higher-level tasks.

3. Justification for and nature of changes

3.1. Justification of New System

The below subsections provide justification for modifications or additions to an existing system. A complete overhaul of current SOAR technologies is not necessarily needed, as these platforms are still in early adoption phases in the market, and business literature heralds the potential of SOAR in meeting business (customer) needs regarding coverage and efficiency (C. Brooks, 2018; Engelbrecht, 2018; Oltsik, 2018b). The findings of the analysis conclude that, while no *new* system is needed, additional features and capabilities are needed in order to also meet the needs of end-users, especially in relation to the next phase of pain points that will become apparent in the market once the immediate concerns are addressed. Essentially, the labor shortage currently affects all levels of security operations, but the lowest level of it has a very high cognitive load with filtering and triage. Once that is addressed with current SOAR, the next layer of the shortage will become more pressing: having a shortage of T2 and T3 experts. This document asserts that SOAR may also be able to facilitate learning and development of future T2 and T3 experts alongside current goals of efficiency and reliability, and that these become the focus of user needs and justification for new features.

3.2. Needs Summary

The summary of needs includes components identified in (The MITRE Corporation, n.d.) that provide context around what users need from a SOAR system. Much of the current drivers of SOAR development are from a business perspective (Neiva et al., 2017). The goal of this section to balance current features with analyst-driven needs with qualitative research conducted in Studies 1 and 2 with different levels of security technologies.

Defining the **enterprise and operational context** of SOAR technology helps set the stage for how and when these platforms are deployed in security operations. SOAR platforms can be deployed in any security operations setting, regardless of sector or size of the company. Companies are feeling much pressure from labor shortages, which means they have less analysts who can work through the deluge of alerts detected by their respective portfolios of security appliances (i.e. network monitoring, email monitoring, etc). While some companies have some form of SIEM to help aggregate log data from said appliances, there are still too many alerts for the human operators to handle. Thus, many companies view SOAR as a solution to this problem, as SOAR has the ability to automate low level filtering, triage, and mitigation for routine tasks. One barrier to implementing SOAR is that not all companies have standardized processes or procedures regarding incident response, which is a requirement for customizing SOAR in the first place. In summary, the operational context is inundated with noise, does not have enough humans to do the work, and does not necessarily have strong procedures.

While SOAR in its ideal deployed state has the capability to meet immediate **needs** of companies performing incident response, Studies 1 and 2 indicate that there are secondary concerns by users that by replacing Tier 1 analysts with automated activities that the pipeline for development of future Tier 2 and Tier 3 analysts may be truncated. All of the experts interviewed had some aspect of Tier 1 ‘desk work’, which they reflected was somewhat boring and monotonous, but valuable in gaining a wide range of experience quickly and getting practice in problem solving needed at the higher tiers. *Analyst users need to be able to follow along with what automated tools are executing and understand incoming information, decisions made, policy enacted, and consequences of each decision made.* As tasks become more automated, the ‘explainability’ of the platform is critical not only for analyst understanding, but also for their development.

Furthermore, *analysts need directed feedback regarding problem-solving and performance.* Studies 1 and 2 indicate that analysts may not receive this feedback directly from peers or higher tier analysts, or even leads and managers. Though literature supports that incident response is collaborative, little evidence of in-person collaboration was observed in current SOCs on a daily basis. Many collaborative interactions with other analysts were through chat features and email, which could limit organic conversation and collaborative problem solving due to physical separation and lack of synchronicity of the communication mode. Analysts often stated that not receiving feedback during incident response regarding receipt and resolution was annoying, and caused many to track the ticket manually through resolution, and in the background of their ongoing other activities. Feedback on how they handled a ticket was not common unless something was glaringly wrong or missing, and after action reviews did not necessarily involve Tier 1 analysts who worked on a ticket. Essentially, feedback is key to learning and development, as well as ensuring awareness of ongoing incidents. Current systems do not support this type of feedback, nor do organization or culture necessarily.

3.3. Conditions/Scenario:

The **conditions** under which this need exists are currently defined by computer incident response settings in which novice, mid-level, and expert analysts are interacting with a SOAR platform. The SOAR technology will employ machine learning to help adapt and respond to incidents, with some tasks fully automated, and others partially automated. The human may or may not see the steps being enacted in the platform, but may need to confirm certain decisions made by technology. The human operators sit at desks with 2 or more screens, though much of the SOAR activities are visible on one screen. The other operators may not sit in the same room, or even in the same building, and the environment is relatively quiet. The analyst may be conducting some sort of incident response with the SOAR platform, but as a novice, may not understand the rules behind the platform, only whether or not it was correct based on his or her own past experience and limited knowledge of the network. The analyst makes a dozen or more decisions for SOAR every hour, and many activities are predetermined in playbooks.

3.4. Growth/Extensibility:

As mentioned above, the skills shortage may take years to overcome, especially if the minimum needed expertise to fill a position increases. Combined with the increasing threats, the need for **growth and scalability** of current SOAR platforms is critical (Oltsik, 2018a). As the internal on-the-job pipeline is replaced with automation, the need for additional SOAR features to support training and development will increase to ensure that analysts at the lowest level can continue to develop and progress into higher tiers. Different types of learning and interaction should be supported through the platform to adapt to different types of learners and ensure effective growth. Measures of comprehension might be helpful (complete with sensors in the platform) to be able to gauge when analysts are ready for advancement.

3.5. Independent of Solution Approach:

In terms of a successful operation, current needs indicate that analysts need more interaction to organically learn, whether this comes from people or automation. While literature describes a collaborative environment to conduct investigations, collaboration done in certain fashions may not encourage or ensure learning, especially if collaboration in incident response refers to handing off an investigation from one person to another (not really collaboration). Interaction designed with humans in mind can ensure effective response as well as enrich investigation and review steps with valuable feedback for awareness and learning.

4. Scenarios of User Needs

The T1 analyst job is changing to become automation supervisor with scripted responses when automation does not step in to execute tasks. The T1 analyst may have some background in security, but this job is essentially a shoe in for other security jobs if they can excel in this environment and learn quickly from their experiences. The environment is somewhat overwhelming, as the company they work for may have many different appliances and tools contributing to their alert systems, and the architecture of the entire system may not be well understood. The analyst largely follows rule-based operations, executing tasks based on predefined playbooks. They may not be part of after action reviews, or get feedback regarding incidents they touch. They may not directly interact with any experts or higher-tier analysts during their normal daily duties. “Learning” may be restricted to learning the playbooks and system names. Progression of the T1 analyst is not clear; they may not have an idea of what positions they are eligible for, or which roles they should pursue. Thus, they do not have direction on what skills they should develop in order to reach them. The goal of this analyst is to stay in the same company and gain additional skills so they can move into a higher role with more complex tasks.

5. Concepts for Proposed System

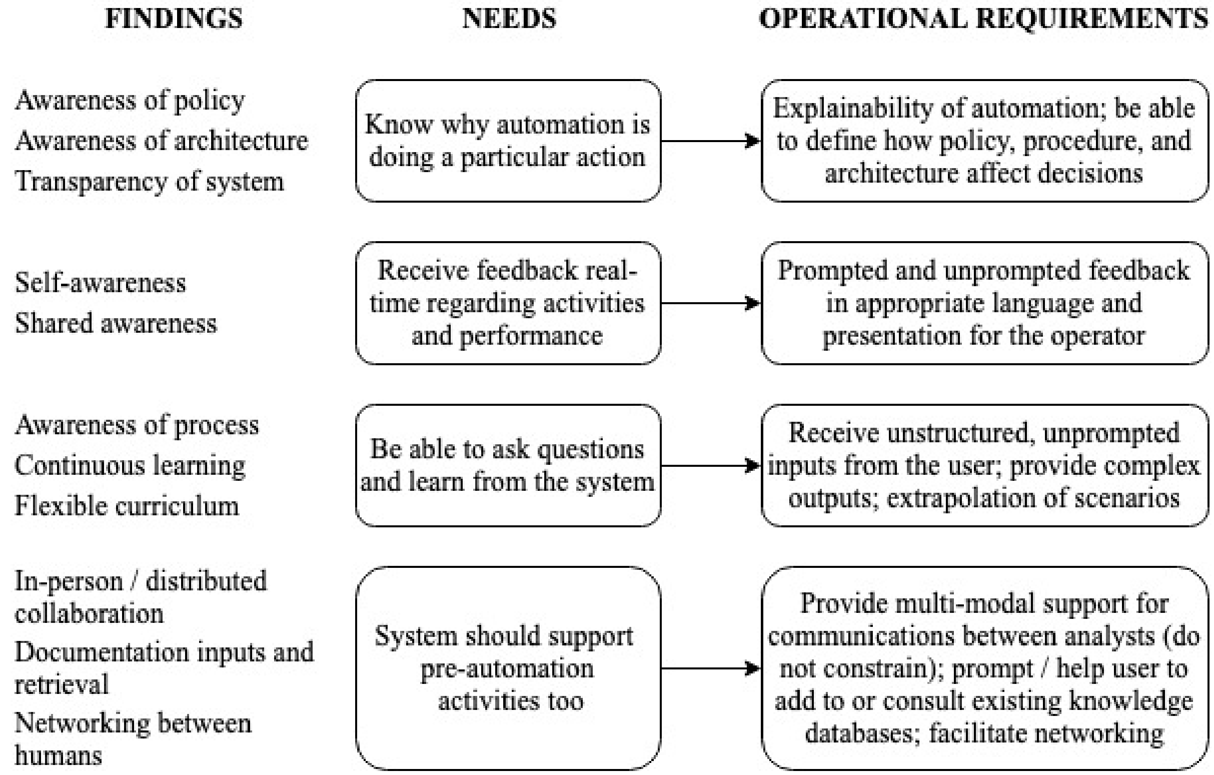
FINDINGS
NEEDS
OPERATIONAL REQUIREMENTS
Awareness of policy
Awareness of architecture
Transparency of system
Know why automation is doing a particular action
Explainability of automation; be able to define how policy, procedure, and architecture affect decisions
Self-awareness
Shared awareness
Receive feedback real-time regarding activities and performance
Prompted and unprompted feedback in appropriate language and presentation for the operator
Awareness of process
Continuous learning
Flexible curriculum
Be able to ask questions and learn from the system
Receive unstructured, unprompted inputs from the user; provide complex outputs; extrapolation of scenarios
In-person / distributed collaboration
Documentation inputs and retrieval
Networking between humans
System should support pre-automation activities too
Provide multi-modal support for communications between analysts (do not constrain); prompt / help user to add to or consult existing knowledge databases; facilitate networking

Figure.N.1. Concepts for SOAR 2.0

5.1. Operational Requirements

Studies 1 and 2 indicate considerations for developing a new system or new capabilities, particularly from the perspective of analysts currently working without SOAR and from experts who have different levels of experience with automated tools. One primary concern was that the learning opportunities in current T1 settings would potentially be lost. In addition, valuable experiences regarding professional development of the lowest tier may be affected by the implementation of broad scale automation like SOAR, which could truncate the talent pipeline further. The above needs stem from these findings, and help formulate operational requirements for the next generation of automation in cyber security.

1. Explainability of automation: be able to define how policy, procedure and architecture affect decisions
2. Provide prompted and unprompted feedback regarding performance and problem-solving activities in appropriate language and presentation for the operator (expanding upon "continuous & shared learning" in (Oltsik, 2018b) and "analytics support" in (Neiva et al., 2017))
3. Receive unstructured, unprompted inputs or queries from the users
4. Provide complex outputs (more than common information retrieval)
5. Provide examples or extrapolation of scenarios
6. Provide multi-modal support for communications between analysts (do not constrain) (expanding upon "journaling and evidentiary support" and "Case management" in (Neiva et al., 2017))
7. Prompt / help user to add to or consult knowledge databases (expanding upon "case management" in (Neiva et al., 2017))
8. Facilitate human-to-human networking

5.2. Modes of Operation

The concept for this proposed system includes three main modes of operation to support current and future needs of cyber security analysts.

1. **Automated:** The automated mode is largely behind the scenes. This mode of operation is not immediately visible to the analyst, but is accessible (transparent) and can be explained by the system itself should the analyst want to know what it is doing and why. The activities within this mode are largely predetermined and consistent with current SOAR development, with the addition of explainability.

2. **Coordinated:** This mode of operation supports interaction with humans directly, and includes orchestrated playbooks (predetermined paths for decision making and actions), as well as communications with other analysts to collaborate during incident response. This mode should accommodate different types of inputs from different users (written, spoken, etc.), and support shared situation awareness between humans and between human and system. The coordinated mode should also support connecting analysts that might be working on something similar, or providing expertise and network context about other analysts during collaboration.

3. **Directed:** This mode of operation is to provide unprompted feedback to the analyst regarding performance or activities. This could include insights regarding investigations or after action reviews, as well as performance metrics and what impacts them. Directed mode could also support informing the analyst of other potentially impactful activities in the larger network, system, or environment, such as a recent incident that was resolved, or a new report that the operator might want to read. Directed mode should be sensitive to how

interruptions impact performance and formulating the appropriate language and presentation for the respective analyst.

5.3. User Classes

The proposed system should support different classes of users, defined by their respective levels of expertise and responsibility. Roles (formal or informal) within the organization may need to be incorporated into the system to help the system itself provide needed support.

1. **Novice:** This class of users needs the most feedback and learning support to help advance to the next level within the organization. Their formal education may be theoretical and systems-based, but not include a broad range of scenarios or understanding of company-specific policy. This class of users may need help in developing flexible problem-solving, as well as feedback to increase self-awareness. This class of users may also need metrics (and corresponding sensors) to define comprehension in addition to performance.

1. **Generalist:** This class of users is considered vetted with respect to policy and scenarios, but may still need system support to help with continuous learning. This might include prompts from the system about other methods (or perhaps new methods) of tackling a particular problem, analytics support regarding trends across incidents, or external information (reports, news articles, after action reviews) that might be pertinent to their continued learning.

2. **Specialist:** Specialists are experts in a particular area within cyber security. They tend to focus on only certain classes of incidents, but at some point were generalists, and could feasibly step in for these types of tasks. The system should support eliciting knowledge from this class of users to help feed the explainability engine for the Novice (and Manager, if needed) class. This class of users can help “teach the system”, so the system can “teach the novice”, all the while expanding on the knowledge base of the firm.

3. **Managers:** The manager class likely needs aggregated information regarding system performance, as well as some prioritization of what areas need attention. This could be analyst performance, tool/appliance performance, and overall system status. In addition to traditional metrics reporting, this class of users may be interested in understanding Novice class growth and development for managing human capital.

6. Summary of Impacts

6.1. Operational impacts

With the added capabilities outlined, analysts and managers at multiple levels can gain from interactions with the system, particularly in a learning/teaching capacity. If lower level analysts are able to learn and develop, even in the absence of current ticket rates and response tasks, the pipeline of talent to higher-level positions can be secured. As the system can still perform automated tasks, the key benefit of increased time and rate of response is preserved, allowing operations to continue as normal. The additional interactions may slow down analyst incident rate (which is already compensated by automation), but have an added educational benefit to increase the value of each interaction.

6.2. Organizational impacts

The capabilities outlined above are meant to help alleviate longer term talent pipeline pressure on low, mid, and high level analysts, as well as decrease the onus on analyst to continue education outside their normal daily duties. The impact would not be immediate, but gradual over time, with the immediate pressure on low tier hiring alleviated by the current version of the system. Additionally, the "coordinated" mode should help facility collaboration between novices, generalists, and specialists, increasing organic networking value-added interaction.

7. Analysis of the proposed system

7.1. Summary of improvements

The current version of the system aims to solve the issue of information overload on a small workforce through automation of tasks and creation of canned playbooks and responses for novices to follow during incident response. Improvements upon previous version of the system will focus on increasing the value of interaction between the analysts and the system, such that both can learn from each other. Additionally, improvements will build upon facilitating and refining activities that the humans already do between each other, which allow the full value of the skilled, diverse workforce to be realized.

7.2. Disadvantages and limitations

Reiterating that these ideas have not yet been validated by users, it is critical that this CONOPS draft be reviewed with a wide set of users, including but not limited to users who currently engage with SOAR platforms. The needs were distilled from two separate studies that included novices, experts, and managers in three firms, and should be validated against a wider set of firms and participants. This could be done through surveys or, if resources allow, focus groups.

Development of the proposed improvements require a much better understanding of human behavior to consider and accommodate with code, and further research in these areas are encouraged if the improvements are to be executed with a high degree of rigor. Incorporating sensing and comprehension assessment, parsing and responding to unstructured question inputs from the user, and supporting shared situation awareness (beyond the perception level (Endsley, 2018)) are each large scientific undertakings. It is recommended that developers and behavioral scientists be included in these studies and discussions to further define costs and effort required for construction of all improvements.

7.3. Alternatives and trade-offs considered

One alternative to this set of improvements is that, instead of pursuing value-added human-machine interaction, to eliminate the human in the system altogether. This alternative is currently not being fully considered because of the changing nature of the environment, and the reliance on human skills at mid and high-level incident response.

Another alternative to embedding educational interactions is that education continue in the direction it is currently heading: formal certifications and degrees, with intermittent courses and self-learning throughout tenure at an organization. While this does not help alleviate turnover (L. Hoffman et al., 2012), it takes the onus off software companies and SOAR customers to manage knowledge growth and preservation in an organization. Moreover, the improvement itself does not prevent turnover, but rather encourages organizational growth and development such that analysts might not want to leave. Should the operational firefighting at the T1 level be resolved with current SOAR technologies, the focus should shift inward to the organization toward more traditional practices of retaining human capital, regardless of the improvements proposed.

Appendix Q: Functional Analysis for SOAR 2.0

The purpose of this document is to provide a functional definition of the proposed system, creating a baseline of requirements for what the system should do. This document provides the first iteration of requirement allocation, with the assumption that, consistent with design theory and principles, the final version of the system will go through many more iterations and evolution as the system is designed, developed, tested, and validated. Using the operational requirements from the CONOPS document, high-level functions are formulated to correspond with each requirement. Each function is supported with examples from literature that describe instances or outlines of similar functionality to support feasibility and definition. This document outlines functional requirements, breaking down each need into multiple high-level functions that could be required in order to meet the need.

This document also provides a first iteration of function allocation from two perspectives. First, the human factors perspective of function allocation is to determine which tasks should be done by which entity: the human or the machine (Lehto & Buck, 2008). Function allocation between human and system provides consideration of strengths and weaknesses of both, such that allocation between human and system teammates can be performed at subsequent stages of development. Second, the systems engineering perspective of function allocation, not entirely unlike the human factors approach, is to determine which subsystem should be responsible for a given function based on factors such as reliability, maintainability, availability, life-cycle cost, performance, producibility, and more (Blanchard & Fabrycky, 2006). For this project and at this level of conceptual development, functions are grouped to show shared capability and potentially same subsystem, such that later steps in system development can draw upon these similarities in requirements. The factors typically included in the systems engineering approach to requirement allocation are not included in this analysis, but act as a guide for actual development of the below functions.

Lastly, this document provides some guidance regarding prioritization based on the functional requirements described below, and the expected effort needed to develop them. Some functions are less complex than others, and could feasibly enter development. Other functions would require additional research and development for both capability and lower-level requirements. The prioritized capabilities are summarized in the last section.

1. Summary of Needs

The CONOPS document describes three main classes of operational requirements from users that highlight how needs of different stakeholders in CSIRTs can be met by new system features and capabilities. These are:

1. Explainability and transparency of automation systems: helping the user understand what the system is doing and why it is doing it.
2. Bi-directional, value-added human-automation interaction: utilizing the human user as not just a recipient of outputs, but as an input feature that has knowledge or needs knowledge / validation.
3. Facilitation of analyst collaboration and networking: providing additional communication support (not just a platform or channel) to help analysts connect and build trust with each other and the system.

The above classes touch on a deeper need from a base of human users with respect to 'smart systems', especially as the field and technologies evolve. Future versions of today's automation systems will allow humans and systems to work *together*, creating a team or hybrid system that can leverage the knowledge base of the system with the learning, flexibility, and social strengths of the human user. Furthermore, one industry report indicated that training and learning are among the top three priorities of cyber security analysts (VIB & Demisto, 2018). The above requirements support this statement by developing human analysts alongside automation. The needs above aim to recognize immediate needs of human-automation interaction in CSIRTs, but these can be extrapolated to other domains to build capacity in human-machine teaming.

2. Mapping Needs to High-Level Functions

The following section connects each need to a subset of functions. Each function is supported with literature and a breakdown of functional requirements in terms of inputs, outputs, controls, and mechanisms (Blanchard & Fabrycky, 2006).

2.1. Explainability and transparency

The need for explainability and transparency of the system refers to the ability of the system to show and explain what is doing and why. Current SOAR technologies are equipped with some level of transparency, depicting flows of tasks that it is executing in the background while the human user engages with the system in other tasks. However, it is not evident that there is flexibility in this depiction regarding different aspects of the tasks being executed or in the system's ability to portray why it is executing particular tasks, or in a particular order (or even the next level of detail down, such as 'with what subsystem it is executing a task?').

Considering that many of the analysts interacting with SOAR are in a supervisory capacity, it becomes important to ensure that the user understands the 'what and why' of system activities. Not only is this the foundation of trust in a system, it can also help the user learn rules (policy and procedures) as executed in the network, and eventually may help them understand potential weaknesses in rules and logic (Abdul, Vermeulen, Wang, Lim, & Kankanhalli, 2018; Core et al., 2006). In a constantly evolving landscape, the system should be adaptable beyond annual human reviews of rules and procedures (which companies currently struggle with regarding review of their own written procedures (VIB & Demisto, 2018). That is, how policy is executed in the system may change over time, and the system should be able to detect when changes are needed, and what those changes might be.

Even at higher levels of incident response, explainability and transparency are important, from both the user's and the system's perspective. If a system is using expert outputs to train its algorithms, then it might also be important for the system to be able to explicitly correlate factors that lead an expert to make a particular decision (inputs), effectively fueling the explainability piece that the system picks up from the human expert. Building on the idea of a true hybrid system (human-machine team), the explainability concept can go both ways between human and automation, with some added capability to reduce load of user inputs.

Two main functions stem from the need for explainability and transparency. First, the system should be able to support definition of logic structure behind actions that it executes and proposes. This definition should extend beyond a flow diagram and have an additional level of detail. Relevant policies and procedures should be referenced. Essentially, the explicit knowledge of the expert who designed the playbook or automated task should be captured in the logic structure definition. Second, the system should be able to explain, in the appropriate language and presentation to the user, any point of that logic structure back to the user. A low fidelity example of this might be to display the logic structure to the user, highlighting the step in question, and providing some additional interaction or drop down menus to describe policy, procedures, or context relating to that activity to answer the 'why'.

There is overwhelming literature support for the need of explainability and transparency in automation applications, also called Explainable Artificial Intelligence (XAI), an overview of which can be found in (Abdul et al., 2018). XAI literature has grown since DARPA's interest in its development (Gunning, 2017), and has even been considered a "grand challenge" in machine learning research (Bonacina, 2017) as computer scientists advance towards algorithms and systems that can depict an describe reasoning beyond the black box approach currently employed. This literature base is extensive, touching on input sending (Chakraborti et al., 2018), interpretability and comprehensibility (Doran, Schulz, & Besold, 2017), levels of explanation and output requirements (Doran et al., 2017; Waltl & Vogl, 2018), output formats (Doran et al., 2017), and learning and training contexts for XAI (Core et al., 2006; Gomboc, Solomon, Core, Lane, & Lent, 2005). Moreover, researchers have gone so far as to propose XAI architectures (Gomboc et al., 2005), development questions (Gunning, 2017), and metrics (R. R. Hoffman, Mueller, Klein, & Litman, 2018). These sources depict a ripe opportunity for this additional capability in CSIRTs, especially as experts are currently driving automation design. This situation offers a unique opportunity for a system to learn from said experts during its own development.

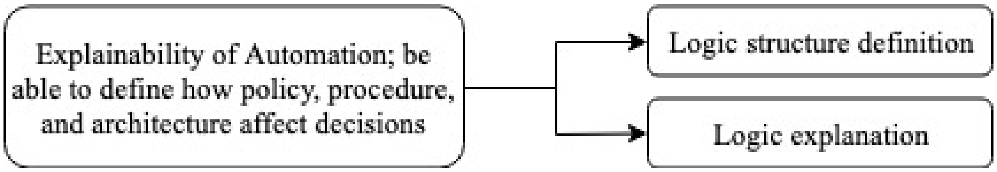
Explainability of Automation; be able to define how policy, procedure, and architecture affect decisions
Logic structure definition
Logic explanation

Figure.O.1. Explainability and Transparency Functions

2.1.1. Logic Structure Definition: Functional Requirements

Functional requirements of the logic structure definition are depicted in Figure O.2 below. In order to deliver explainability and transparency to the user, it is important that the system itself have an "understanding" of the logic employed during incident response. SOAR platforms incorporate knowledge and expertise from cyber experts in developing playbooks and automated tasks.

However, the ‘why’ of those design decisions may not be captured. It is this content that is critical to knowledge transfer to analysts such that they understand and expand upon the rules and policies guiding automated activities. To achieve this, inputs requirements include an additional level of knowledge regarding ‘the why’ during task and playbook design. This includes policy, procedure, and context around each decision point and activity that the system executes.

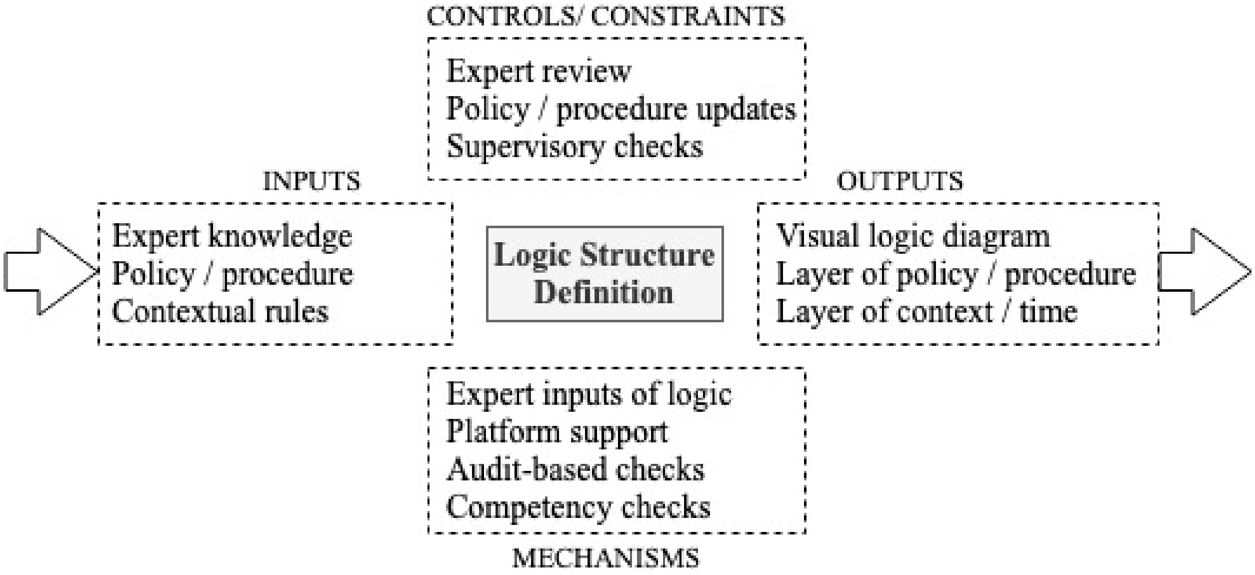

Figure.O.2. Logic Structure Definition Functional Requirements

Several controls and constraints should exist to ensure proper functioning of the logic structure definition. Expert review of this extra layer of knowledge is critical to validate content, and policy and procedure updates should trigger an update to this layer of data. Furthermore, checks by the human supervisor should be routine, such that the system can be ‘challenged’ regarding why it is doing a particular activity. This allows the human user to remain at a necessary level of skepticism, as unquestioning trust in an automated system can be dangerous (Doran et al., 2017), especially if it is dynamically learning and adapting.

Mechanisms in delivering logic structure definition include continuous expert inputs of logic, including passive capture of rules during expert incident response. The SOAR platform should support this added layer of detail, as well as the ability to review and change it. Furthermore, audit-based checks of the system, and competency checks of the user, will help validate and maintain the content and its usefulness.

Finally, the outputs of the system include some visual component of the logic structure, complete with layers that show or reference policy and procedures that affect outcomes of a particular step. Layers should also include context and time stamps, such that the user can validate environmental and other contextual information along with recentness. These outputs complement the inputs of the function for logic structure explanation.

2.1.2. Logic Structure Explanation: Functional Requirements

Figure O.3 illustrates the functional requirements for logic structure explanation. In order to develop the capability of explaining rules and logic to a user, the system has a series of required inputs. First, and most fundamentally, the logic structure definition is a prerequisite, such that there is content to be explained. Next, the system should be able to receive inputs or detect which action, decision, or point therein needs to be explained to the user. Additionally, the system should be able to determine (through explicit or implicit cues) what level of explanation is needed, and the most appropriate style of presentation to a given user. For instance, if the user is an analyst who has been working as a T1 incident responder with supervisory responsibility over automation for 6-8 months, he/she might have an idea of why the system is doing what it is doing, but want a deeper, more specific explanation regarding policy. This user has adequate vocabulary in security and networks, general knowledge of the systems in place, and a tendency to use or create graphical resources more often than written resources. This user would need an in-depth explanation of policy, preferably in a graphical format, with technical vocabulary.

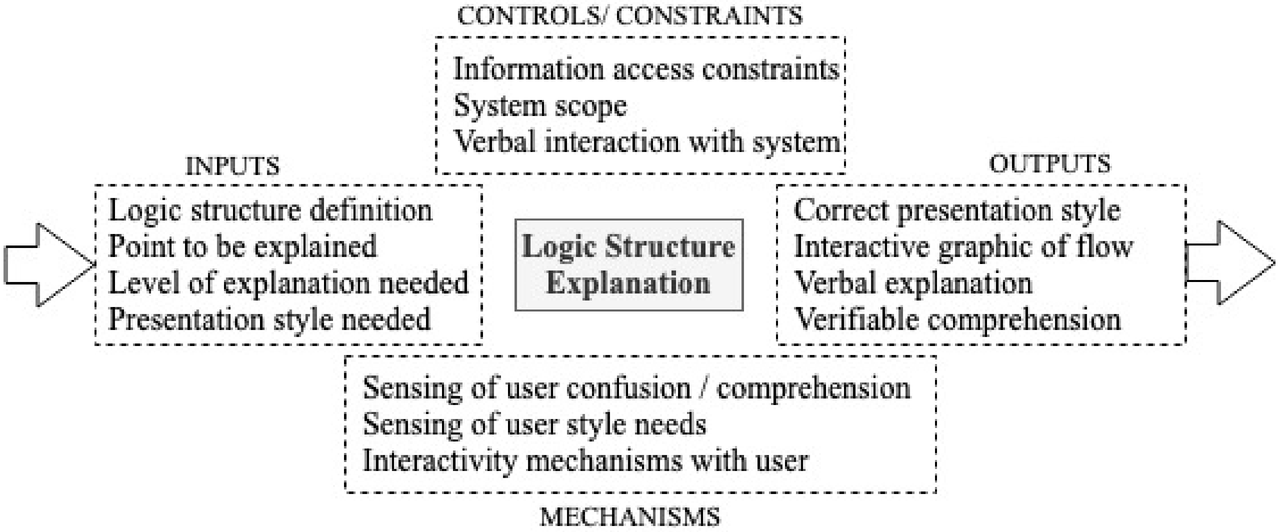
CONTROLS/ CONSTRAINTS
Information access constraints
System scope
Verbal interaction with system
INPUTS
Logic structure definition
Point to be explained
Level of explanation needed
Presentation style needed
Logic Structure Explanation
OUTPUTS
Correct presentation style
Interactive graphic of flow
Verbal explanation
Verifiable comprehension
Sensing of user confusion / comprehension
Sensing of user style needs
Interactivity mechanisms with user
MECHANISMS

Figure.O.3. Logic Structure Explanation Functional Requirements

The controls and constraints around this function include pre-existing policy constraints, scope of the system, and existing capabilities of human-computer interaction. First, if the reasoning behind a decision or activity includes classified or executive level knowledge, the user may not have access to this information, and communicating the ‘why’ might be restricted to lower-level explanations. Furthermore, the system scope cannot include all potential answers to all potential questions. Lastly, current capabilities for human-system interaction act as temporary constraints, as well as development ideas, for how the system can interact with a human user. Interaction constraints may lead to the human and system not understanding each other’s needs and messages.

Mechanisms that drive this function mainly include aspects of interaction and sensing. First, the systems should have some mechanism for sensing user confusion or comprehension to validate user understanding and system determination of appropriate topic, level, style. Building upon this mechanism, the system should have the capability to sense different aspects of the user’s style to drive output delivery. This might include format, graphical vs. verbal, and language and vocabulary. Finally, the system should have mechanisms to support multi-modal, and even non-verbal, interaction with the user, which includes sensing, input format, output format, and feedback design.

The outputs of this function include multi-modal representations of the logic structure. This might include interactive graphic flow to walk through with the user, verbal explanation with different levels of detail, and verifiable comprehension. It is also critical that the system present outputs in the correct style of outputs needed for the user, which could be dynamic.

2.2. Bi-directional, value-added human-automation interaction (HAI)

In order to increase the value of interaction between the human and automation, it is imperative that the information flows are bi-directional and dynamic. Current human-automation interaction is limited in CSIR by programmed, or even learning, automation that largely operates separately from the human, feeding information to the operator when input is needed for a decision or validation. However, this interaction could be much richer and feed in both directions. Human-automation interaction need not be limited to the graphical user interface designed by the developer of the software. Increasing the modes of interaction to accommodate verbal and non-verbal cues beyond the screen and keyboard open up additional opportunities for the system and the human to work with and learn from each other.

The follow section describes two groups of functions (Figure O.4) that build upon the idea of that HAI can be bi-directional and value-added, such that the system and the human can work together as a true hybrid (centaur) team. The first group focuses on communication inputs and outputs, and some of the functions needed to support the capability. These include capabilities such as: sensing to determine when feedback or interaction is needed, speech recognition (for verbal inputs), and feedback formulation and presentation (for verbal or graphical outputs). The second group of functions overlaps with the first, and adds ‘intelligence’ functions for communicating more complex information. These include parsing questions from users (what does the operator mean, want to know or what is their intention?), sensing to determine comprehension (did they understand the system’s answer?), an analogy engine (for providing examples, scenarios, or analogies in teaching), and scenario extrapolation (being able to abstract the scenario and define in different contexts with expected outcomes).

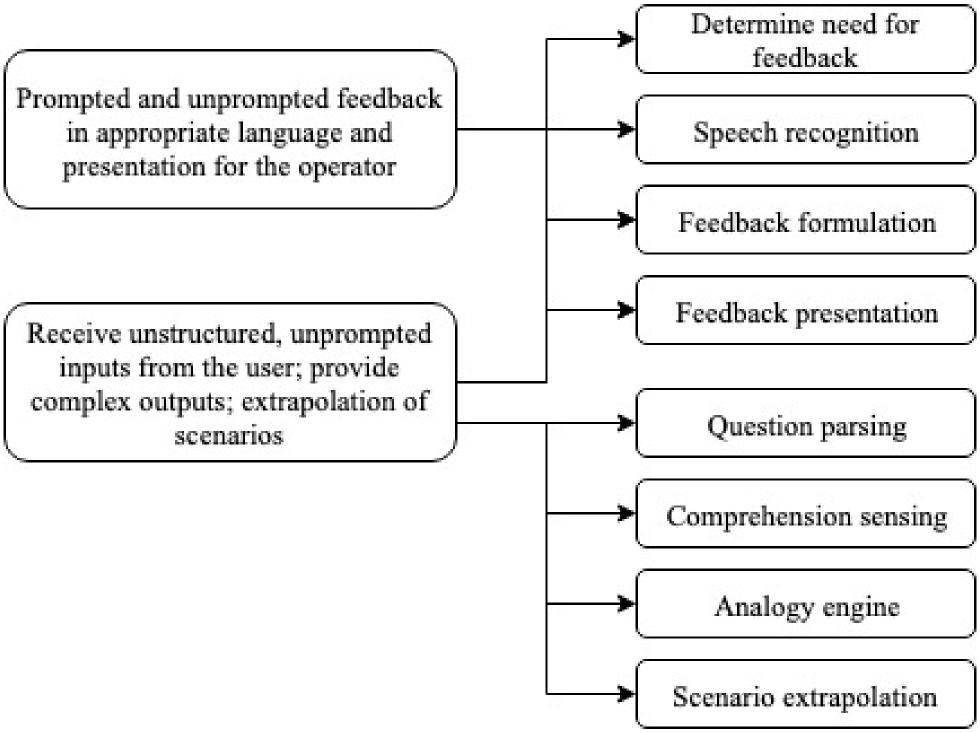
Prompted and unprompted feedback in appropriate language and presentation for the operator
Receive unstructured, unprompted inputs from the user; provide complex outputs; extrapolation of scenarios
Determine need for feedback
Speech recognition
Feedback formulation
Feedback presentation
Question parsing
Comprehension sensing
Analogy engine
Scenario extrapolation

Figure.O.4. Bi-Directional Value-Added Human-Automation Interaction Functions

Some of these functions are already in development for other applications, and research has identified similar traits as necessary for intelligent assistants (Budiu & Whitenton, 2018) like Amazon Echo and Google Home. General capabilities for intelligent hybrid agents have been defined for human-automation symbiosis (Romero, Bernus, Noran, Stahre, & Fast-berglund, 2016). A tangential topic that should be considered in this class of functions is the appropriate level of automation (LOA), the interaction requirements for that level (Sheridan & Verplank, 1978), and a design method that supports human-automation interaction (Johnson, Bradshaw, & Feltovich, 2017). Nevertheless, literature in these functional areas has been growing in both theory and development. Speech recognition and natural language processing are well-researched areas (Lee, Soong, & Paliwal, 1996; Xiong et al., 2018; Yu & Deng, 2016) which are both precursors for question parsing (Pearson, n.d.). Another aspect of speech recognition is building vocabulary sets that the system can understand, which can be a large undertaking (Warden, 2018), but there is evidence of human-robot interaction that has ongoing speech interaction capability (Sheridan, 2016; Vlahos, 2015). This includes some aspects of reaction (feedback formulation and presentation), which require sensing (Piasecki, Fendley, & Warren, 2017; Romero et al., 2016; Schilberg & Schmitz, 2017) and potentially learning from human feedback (Knox, Stone, &

Breazeal, 2013; Sheridan, 2016). Finally, higher functions of artificial intelligence relating to formulation of analogies and scenario extrapolation have been suggested (Hoffman, Klein, et al., 2018; Sheridan, 2016), but are not yet developed or standard.

Clearly the development of this class of functions is both complex and interdisciplinary, drawing on expertise from multiple sub-fields of computer science and human factors to determine balance and interaction dynamics between the human and the system. Due to the size and scope of this set of functions, it is recommended that this set be investigated separately with a systematic literature review of relevant domains in addition to industry review of current development across human-automation interaction (HAI), human-robot interaction (HRI), and human-system interaction (HSI) areas. The priority of the study should be to determine additional structure and steps for developing the capabilities suggested here.

2.3. Facilitation of analyst collaboration and networking

One area of need from users was around collaboration and networking. Understanding where resources are (including human and non-human), and which resources are appropriate are both important pieces of collaboration that require more than just a chat platform. Moreover, as CSIR operations become more distributed, the system in between humans will need to do more to facilitate collaboration that traditionally happens in person. The following section describes a class of functions for a system to help facilitate collaboration and networking in SOCs (Figure O.5). Three functions are broken down into requirements, while the other requires additional research to define and conceptualize.

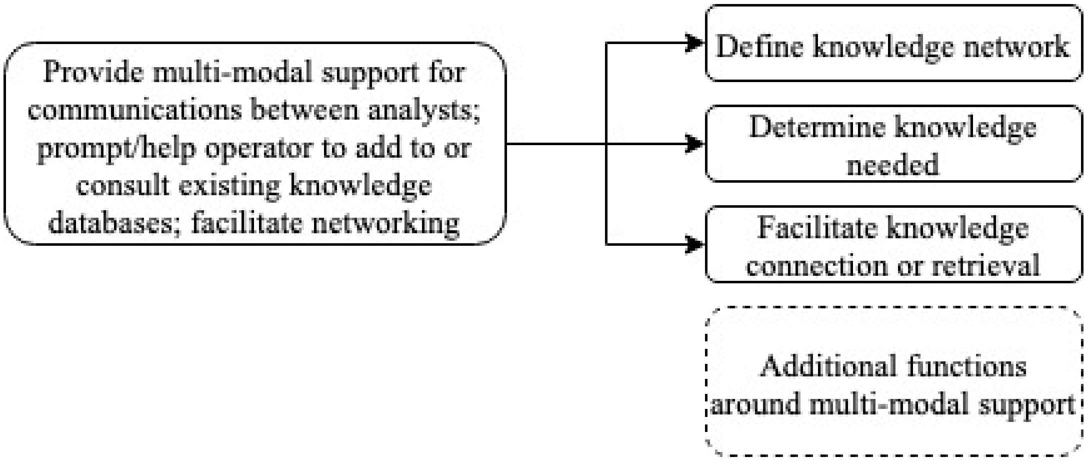
Provide multi-modal support for communications between analysts; prompt/help operator to add to or consult existing knowledge databases; facilitate networking
Define knowledge network
Determine knowledge needed
Facilitate knowledge connection or retrieval
Additional functions around multi-modal support

Figure.O.5. Functions to Facilitate Collaboration

One important component of collaboration is knowing where knowledge or information exist within the network; this is essentially knowing ‘who or what to collaboration with’. SOAR platforms are currently capable of capturing some aspects of knowledge from human interactions, and capable of being programmed to use or update certain resources. The next step in building capacity in this area is to go a step farther in defining an actual knowledge network that is both accessible to the user and dynamic as knowledge evolves in the organization.

Next, the system should be able to detect or determine what knowledge or expertise is needed at any given point in any given task. With the current deficit of expertise across the cyber security domain, it becomes increasingly important to utilize existing expertise to build knowledge foundations in systems and human assets. Being able to explicitly identify these needs helps both the system and the operator better define knowledge deficits (opportunities for development), and opportunities for collaboration and knowledge sharing.

Third, the system should be able to facilitate knowledge retrieval, being transparent about what knowledge is needed, where it is located, and how to access it. After identifying the expertise needed and who has that expertise, the system might prompt the operator to contact the available humans with the needed information, or create the connection directly between the two humans while supporting shared awareness by sharing screens. This facilitation step goes beyond having an available chat platform, and acts as an operator assistant to connect resources for faster, more effective collaboration and knowledge sharing.

Finally, the system should support additional capability in multi-modal communication and information sharing between humans working in the system. Shared awareness can be secured with more than just shared screens, but also with voice and video support, or even tactile feedback to assist coordination between two humans who are not collocated, but working on the same incident. Moreover, this multi-modal communication can also exist between the system and the human, such that the interactions between the human and system are more natural and fluid. This level of capability propels SOAR development firmly down the path of “SOAR as a teammate”, working towards more human-like interactions with automation by supporting richer communication between them.

Knowledge networks are not a new concept, and include two main perspectives and multiple approaches (Armistead & Meakins, 2002). The first perspective is the idea of a knowledge network from the standpoint of the human entity, which is also called a transactive memory system (TMS). Simply put, this can be thought of as a network that portrays ‘who knows what’ in an organization. TMS can be preserved in mental representations or virtual representations, and can also be called ‘knowledge directories’. Creation of these directories can be transferred to information systems that are designed to support knowledge sharing (Jackson & Klobas, 2008). Research to support deeper aspects of knowledge transfer in a real network of people can also be considered useful in supporting TMS development and knowledge sharing practices between humans (Reagans & McEvily, 2003). The second perspective stems from the data standpoint, which focuses on creating knowledge networks from databases (Chen & Lynch, 1992) and passively from humans (Lin et al., 2009). These are also called ‘knowledge graphs’, and have been popularized by large search tools such as Google and Wikipedia. Knowledge graphs help increase the speed of knowledge searches by creating connections between relevant pieces of information (Paulheim, 2016; Pujara, Miao, Getoor, & Cohen, 2013).

By merging the above perspectives, future SOAR platforms can create (through automation) representations of knowledge within an organization and help facilitate knowledge finding and sharing between human and non-human entities. The explosion of search engine capabilities add to potential ideas for how to make this idea come to fruition as algorithms can help navigate what a human is looking for, or what they might need from the search. Anticipating the human’s needs helps build an understanding of knowledge deficits, and complements the construction of the knowledge network.

2.3.1. Knowledge Network Definition: Functional Requirements

Functional requirements for defining a knowledge network are depicted in Figure O.6. In order to define the knowledge network within the organization, information regarding the organizational structure and system architecture may be useful in building a realistic layer of physical or social navigation of the overall organization. Additionally, predefined knowledge classifications and ontologies help construct common language and models for what the knowledge network should include. Finally, incident handling data as a continuous input to the function will help the system update the knowledge network as different humans work on particular types of problems, gaining both experience and knowledge over time.

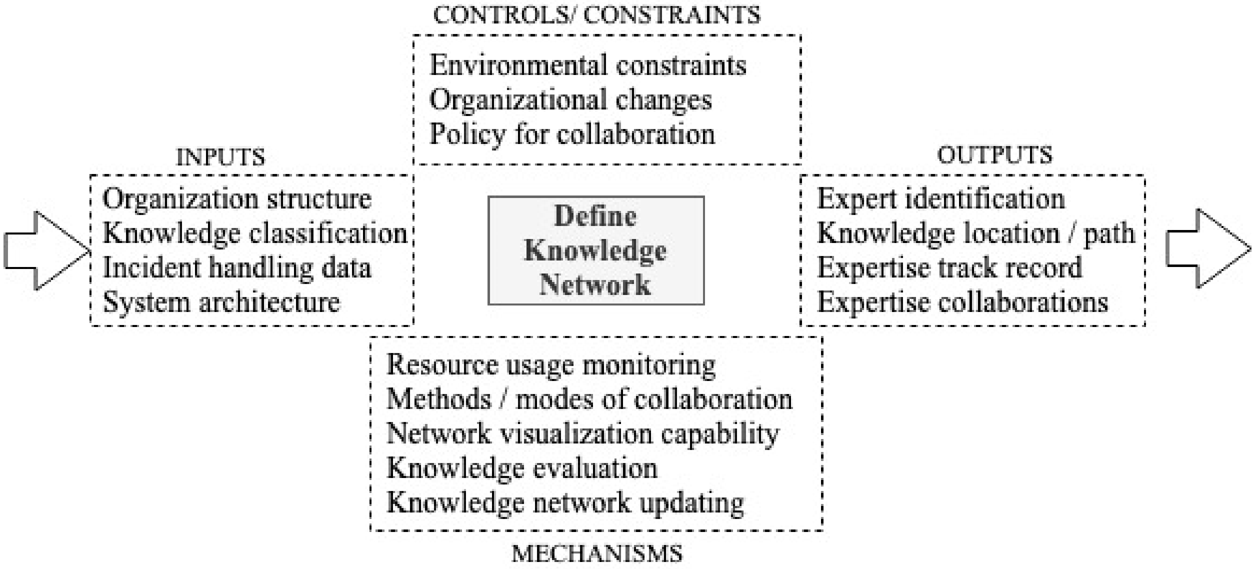

Figure.O.6. Functional Requirements for Defining a Knowledge Network

This function does have certain controls and constraints that help define the scope. First, environmental constraints, such as system structure, may prevent SOAR from connecting or adding on to knowledge networks on its own, essentially acting as barriers to particular areas. Without a connection to a knowledge source, the system is not able to incorporate content into the knowledge network. Next, organizational change in structure may affect the ‘social navigation’ or knowledge assets in the system. If there is high human turnover or significant knowledge sources leave the organization, the system will need to update the knowledge network accordingly, and learn new sources of and paths to that knowledge. Finally, policy around collaboration may prevent the system from connecting resources that might need each other. In the case of security clearances, a clean handoff may be required between human entities because one human does not have the required clearance to work on the next stage of investigation or response. These rules would need to be incorporated into how the system defines the network.

Mechanisms to support this function help with both operation and adaptation. First, it might be relevant to include monitoring of resource usage, which helps the system learn from how humans locate and utilize resources. Methods and modes of collaboration help define links between human nodes, and act as pathways for connection. A capability to visualize at least some dimensions of the knowledge network may help the human operator understand the ‘bigger picture’ of where knowledge exists, but at the very least helps the managers of the organization identify critical knowledge assets. Constant evaluation and consumption of knowledge data will help the system adapt and learn from the human users, including who is progressing (building expertise), or how they are utilizing knowledge they find. Finally, mechanisms to support regular automatic updates help keep the system up to date.

The outputs of this function aim to identify where knowledge exists in the overall organization (include who/what *and* where), the path to that entity through various modes, the history (reliability) of that knowledge asset, and identification of when, how, what, and in what context expertise is shared in the organization. These outputs support both construction of the knowledge network and potential metrics to manage it.

2.3.2. Knowledge Deficit Determination: Functional Requirements

Functional requirements for determining knowledge needed are depicted in Figure O.7 below. In order to detect what knowledge or expertise is needed by the operator, several inputs are needed. First, some detection of the context, including relevant keywords, from the incident help set the stage for what the operator might be looking to do, or what additional information they might need. A solid foundation of semantics and classification around types of knowledge in this context is also needed to help the system navigate shared language and models of connections between different types of knowledge with the analyst. Next, historical information regarding the analyst’s performance on similar incidents might provide clues regarding what that person has done in the past, and what they might be looking to do again for the incident at hand. The analyst’s physiological behaviors inside and outside of the system might also act as indicators of confusion or confidence, which can help the system to better interpret and anticipate needs.

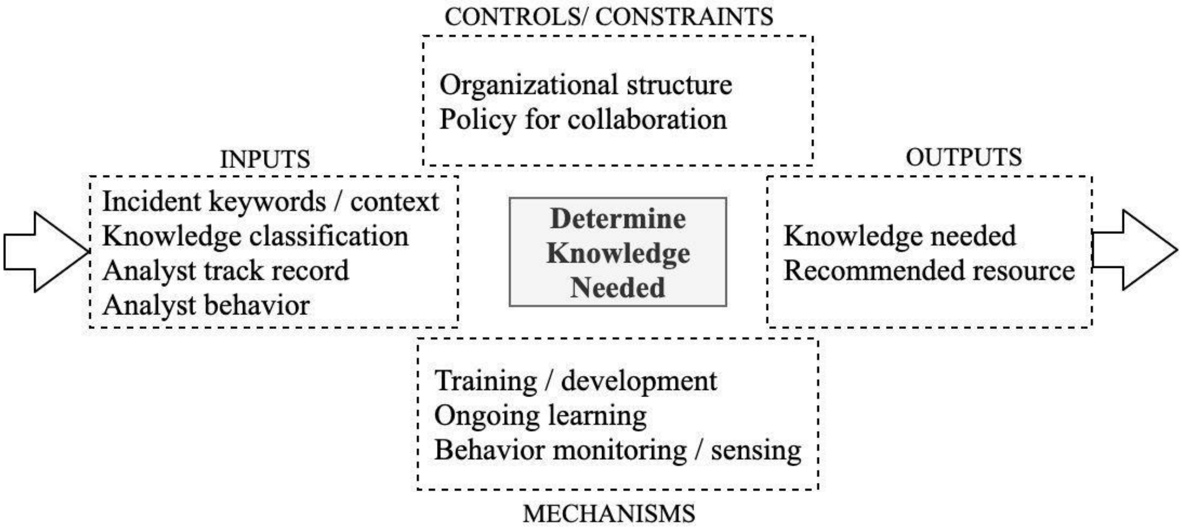

Figure.O.7. Functional Requirements for Knowledge Deficit Determination

Controls and constraints for this function are similar to defining the knowledge network. Organizational structure may act as a constraint or a control, creating boundaries for with whom or what an analyst may share knowledge. Likewise, policy creates similar boundaries with respect to access to certain knowledge assets depending on clearance or rights. These constraints are placed upon the outputs of the function regarding a recommended resource for the analyst to seek out to find the knowledge needed.

The mechanisms to support this function are mainly from the perspective of sensing and updating models and devices used to detect knowledge needed (or might be needed) by the analyst. Training and development of the analyst inside and outside the system can prompt comprehension evaluation or validation such that the system can update how much knowledge an analyst has in terms of training and practice. The system itself can learn from these interactions, and update its own processes and the knowledge network. Finally, mechanisms to support sensing and behavior monitoring of the analyst provide different types of inputs for the system to use to anticipate user needs.

The outputs of this function are identification of knowledge needed, and a recommended resource for where to find it (based on the outputs of the knowledge network function). Successful operation of this function can be validated using outputs of the knowledge connection function (below), and through external means (short surveys for users) during stages of system update or calibration.

2.3.3. Knowledge Connection: Functional Requirements

Figure O.8 depicts the functional requirements for facilitating knowledge connection. The inputs to this function include identifying sources to connect (from the previous two functions), identifying the availability of the resources and appropriate channel for connection, and integrating contextual information (i.e. urgency, geographic differences). These inputs help determine who or what need knowledge from each other, how they should share it, and pertinent details that might affect the ‘who, what, and how’.

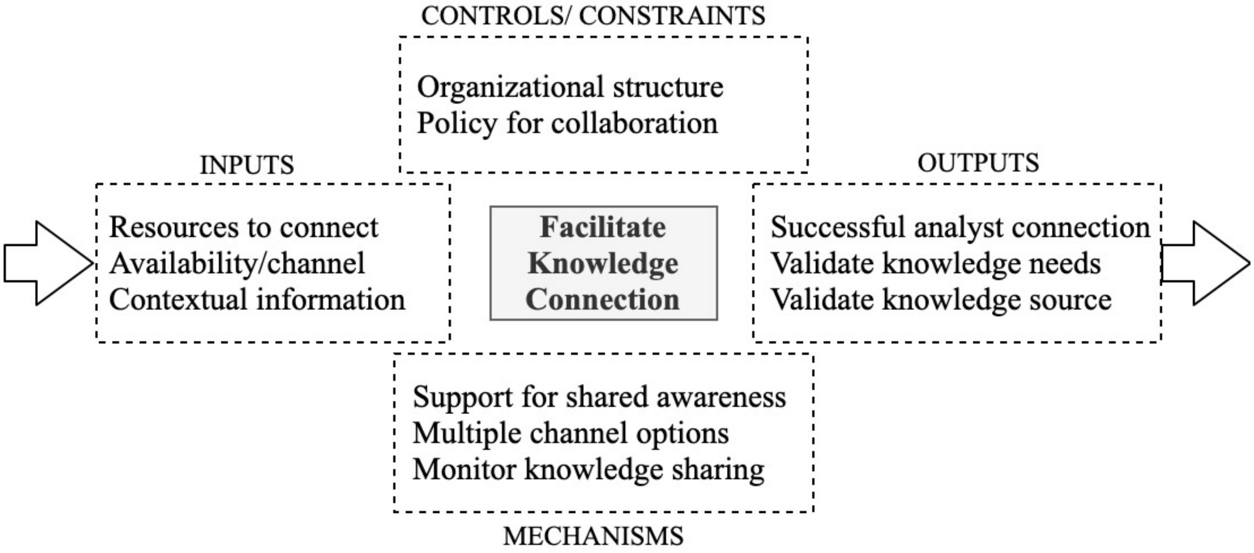

Figure.O.8. Functional Requirements for Facilitating Knowledge Connection

Organizational structure and policy for collaboration remain constraints for this function, as availability and context are affected by these factors. Mechanisms needed to support this function include monitoring of knowledge sharing patterns and connectivity between analysts and resources, maintaining multiple channels for connecting resources, and supporting shared awareness when two human analysts need to share knowledge. The last support mechanism should aim to provide a common operating picture between analysts to minimize the time needed to get up to speed and to facilitate stronger collaboration. Channel options should allow the system adequate flexibility in terms of time, method, and salience, such that the resource connection is efficient and effective. Monitoring of knowledge sharing patterns will help the system learn about how particular analysts tend to seek out or share information with other analysts and databases, overall enriching the data in the knowledge network definition function.

The function should successfully facilitate connection between an analyst and a resource (human or non-human) while validating the knowledge needs and correct knowledge source. Theses validation points act as potential metrics to ensure that analysts are getting the information they need, when they need it, and from a reputable source.

3. Partitioning of High-Level Functions

In systems engineering literature, allocation starts with grouping of functional requirements, also referred to as partitioning. Allocation commonly involves identifying common systems that functions would potentially use, which can somewhat lead system designers to start defining "form" of the final system. The intent of this document is not to determine form, thus designation of potential systems is not included in this document. The function allocation process should also be iterative (Blanchard & Fabrycky, 2006), and the grouping of functions presented above is a first attempt to partition functions. Additional detail will be provided during development of the operational concept.

At this stage of technological development in CSIR, there is intense focus on overall reducing tasks allocated to the human in the system due to overload, burnout, and labor shortage. While this Functional Analysis is too high level to delve into assigning tasks to humans and automated agents, it is worth noting that allocation between human and system will evolve over time, especially as capability increases. SOAR 2.0 might involve dynamic function allocation, allowing the human and system to change task responsibilities with time and context and adapting to expertise and capacity growth.

4. Prioritization of Functions

This Functional Analysis presents ideas for conceptual development regarding additional capabilities for SOAR platforms to better meet the needs of CSIRTs. However, in proposing these functions, it is also clear that at this stage of development for machine learning and artificial intelligence, not all of the above functions can be developed in a short time frame. Thus, there is a need to prioritize the functions by feasibility and scope, such that software developers have some options to work on while the state of the art advances to support the other functions. Feasibility is defined here as the ability to produce the specified function in a relatively short and reasonable period of time (5 years or less), and with a small enough scope that it could be done on a reasonably sized cross-functional team. That is, the researcher estimates that these capabilities that can be developed and realized in traditional industry project timelines with industry-sized teams. The other functions have a much larger estimated scope for research and development, but are ripe opportunities for future activities.

The prioritized function classes from this analysis are *explainability and transparency* and *facilitating collaboration*. Considering the current DARPA focus on XAI, there is a wealth of ongoing research, methods, and tools available to developers to start creating this capability in SOAR. The researcher recommends using available resources to incorporate best practices and cutting edge approaches in order to develop this set of functions. However, facilitating collaboration through knowledge networks will require some additional conceptual development in order to meet the unique needs of this environment. This set of functions will be explored in the next phase of this research: the operational concept.

Appendix R: Operational Concept Definition: Facilitating Collaboration

1. Purpose / Goal of this document

This document serves as the operational concept for one particular set of functions identified in the Functional Analysis (Appendix Q). Though literature diverges on the depth of information needed in an operational concept document, the researcher aims to clearly define the concept of “facilitating collaboration” for SOAR technologies with respect to interfacing with and assisting human analysts. This document is the first draft of this concept, and acts as a ‘living document’, to be revised in future efforts (outside the original dissertation) with additional research and development activities. This document does not include all recommended validation activities or evaluation of alternatives to move forward with full development, and has limited focus on acquisition or form, which are both traditionally part of an operational concept from the Department of Defense DoD perspective. However, it does provide the foundation of the concept by providing additional details around how the functions should operate and interact with users. Furthermore, this document shows that the Systems Engineering conceptual development process is worth pursuing to bridge human-sourced data with robust system (SOAR) development.

2. Concept Scope

Currently, SOAR platforms support human-to-human communication through chat platforms, as well as exchange (and logging) of documentation used during collaboration. Users expressed additional needs regarding connecting them to knowledge needed, as well as richer collaboration between human analysts. Understanding where resources are (including human and non-human), and which resources are appropriate are both important pieces of collaboration that require more than just a chat platform. *Facilitating collaboration* is a set of functions that will help define and identify knowledge sources within a network, matching the human user in need of knowledge to the appropriate source, and facilitating communication and shared awareness between them. Research indicates that this functional set is the baseline for teaming, especially between humans and automation (Lathrop, 2017), which can help propel automation developing firmly into the stage of human-system collaboration.

As indicated in the Functional Analysis document (Appendix Q), literature supports the idea that knowledge or expertise networks have already been developed, as have methods to connecting users to information during active searches (e.g. Google). This evidence indicates that this concept is indeed feasible from a technical perspective, which is further supported by indicators of maturity in the cyber domain (Lathrop, 2017). Literature has also proposed like-minded functions in cyber security (Abbass, Petraki, Merrick, Harvey, & Barlow, 2016; Lathrop, 2017; Sycara & Lewis, 2004), as well as broader frameworks for augmented cognition in team environments (Cuevas et al., 2007). Estimates are not available for economic and time-related factors for deployment in CSIR.

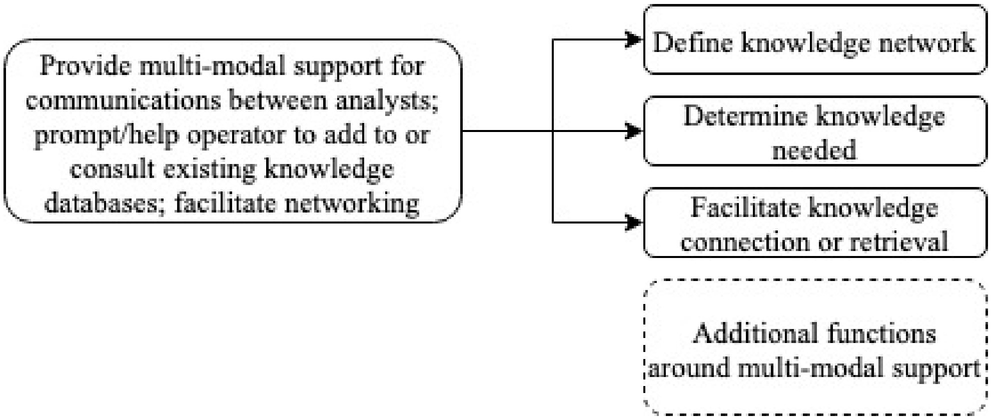
Provide multi-modal support for communications between analysts; prompt/help operator to add to or consult existing knowledge databases; facilitate networking
Define knowledge network
Determine knowledge needed
Facilitate knowledge connection or retrieval
Additional functions around multi-modal support

Figure.P.1. Functions to Facilitate Collaboration

3. Conceptual Context & Boundaries

The concept of facilitating collaboration is set in the context of distributed, tiered CSIRTs using SOAR technologies to conduct incident response activities in 24/7 operations. CSIRTs are assumed to be within an organization, but could potentially include inter-organizational collaboration in future versions of this concept. Based on this assumption, the boundaries of the concept are confined to one organization, and all systems within that organization's network of operation.

Operationally, this concept includes capabilities to identify, store, connect, and distribute expertise when and where needed. Organizations may not currently have a definition of the expertise within their networks, be it human or machine-based expertise. Thus, one supporting concept to facilitating collaboration is capturing the transactive memory of an organization and storing it for use by the SOAR platform. SOAR platforms are currently capable of capturing some aspects of knowledge from human interactions, and capable of being programmed to use or update certain resources. The next step in building capacity in this area is to go a step farther in defining an actual knowledge network that is both accessible to the user and dynamic as knowledge evolves in the organization.

Another operational concept is to be able to detect or determine what knowledge or expertise is needed at any given point in any given task. This has also been identified in literature as *cognitive automation* (Onken, 2003). Being able to explicitly identify these needs helps both the system and the operator better define knowledge deficits (opportunities for development), and opportunities for collaboration and knowledge sharing. This concept includes parsing user inputs, much like a Google search engine, but might also include other types of inputs. Eye-tracking, pupillometry, and other human-computer interaction patterns may help indicate what a user is trying to do, what they are looking for, or if some information they are processing is confusing or difficult. By combining some of these inputs, the SOAR platform might be able to deduce or even predict when a user will seek out additional expertise, and even what that expertise might be.

Within the scope of this operational concept, the system should also be able to facilitate knowledge retrieval, being transparent about what knowledge is needed, where it is located, and how to access it. After identifying the expertise needed and who or what sub-system has that expertise, the SOAR platform might prompt the operator to contact the available humans with the needed information, or create the connection directly between the two humans while supporting shared awareness by sharing screens. This facilitation step goes beyond having an available chat platform, and acts as an operator assistant to connect resources for faster, more effective collaboration and knowledge sharing.

By merging the above perspectives, future SOAR platforms can create (through automation) representations of knowledge within an organization and help facilitate knowledge finding and sharing between human and non-human entities. The explosion of search engine capabilities add to potential ideas for how to make this idea come to fruition as algorithms can help navigate what a human is looking for, or what they might need from the search. Anticipating the human’s needs helps build an understanding of knowledge deficits, and complements the construction of the knowledge network. This operational concept is not unlike models proposed in semi-autonomous vehicles, such as the driver-adaptive decision model in (Onken, 2003).

4. Context of Use / Operational Scenarios

4.1. Overview

The functions supporting collaboration should be able to quickly identify from the human user what expertise or knowledge is or might be needed, identify and prioritize potential sources of that knowledge by availability and history (including performance and/or access frequency), and facilitate a connection between the user in need of knowledge and the knowledge source. If between two humans, this facilitation would include creating a communication connection (potentially a screen-sharing connection) to ensure that analysts are on the same page during incident response.

The sequence of events for this set of functions includes:

1. Create / maintain a knowledge network of humans and non-human resources for different knowledge areas
2. Monitor analyst activity / Update knowledge network based on performance
3. Detect when help is needed from analyst [DECISION]
4. Determine what help is needed by analyst [DECISION]
5. Determine viable sources for knowledge needed
6. Prioritize / recommend sources [DECISION]
7. Facilitate connection with source [DECISION]

Defining data flow of this function set first requires definition and description of two major classes of data. The two classes differentiate sources of data (system or human), and indicate in the descriptions potential interactions between sub-systems and between humans and sub-systems. The data would be ingested by the SOAR platform to carry out the sequence of events indicated above.

The first class of data can be described as *historical and ongoing incident data***,** which provides a wealth of information regarding which analysts know what, and where else knowledge is located in the network. While some aspects of this can be explicitly defined, they can also be detected through performance monitoring, credential updates, and interaction of analysts, all of which help indicate vetting of an individual with respect to a category or class of knowledge. Furthermore, monitoring what sources analysts access that are non-human help indicate where they tend to find knowledge pertaining to types of incidents. This class of data provides points to be correlated regarding knowledge area, location, and quality, all of which can help the system create a model of the knowledge network in an organization.

The second class of data involves *monitoring user behaviors* within the SOAR platform, as well as other interactions outside of it. For instance, eye-tracking and pupillometry would help the platform determine what an analyst is looking for and where they might be confused in the process. With current capabilities around playbooks, the system can effectively track where in the process the analyst is, and use the physiological data to predict what specific aspects of the process need outside knowledge or intervention. Other inputs from the analyst include sources accessed, keywords of searches, historical methods employed by the analyst, and potentially analyst preferences of tools, techniques, and interfaces. Furthermore, the platform might also be able to cross-reference the individual's schedule and location / time zone to help determine availability.

From the sequence of events above, the following steps constitute as decision points based on human interactions. These decision points need not be only within the system, but could also include explicit user inputs to guide decisions made:

1. *Detect when help is needed from analyst:* the system could use interactions with the system, as well as physiological data, to detect when help might be needed from the analyst. This could prompt a question to the user to confirm if help is needed.
2. *Determine what help is needed by analyst:* like the previous decision point, predictive capabilities using inputs could help guide what type of help is needed by the analyst, with confirmatory activities to validate the system's determination.
3. *Prioritize / recommend sources:* the system would need to make some decision regarding sources to present to the analyst, much like a recommender system in other contexts.
4. *Facilitate connection with source:* like the prioritization and recommendation step, the system would need to conclude similarly the mode and presentation requirements to best connect an analyst to a knowledge source. If user inputs are taken into account, the system could learn from these interactions to guide user preferences.

4.2. Performance

Performance of the system should aim to minimize response time to a detected threat in the network. In order to measure effectiveness of the collaboration functions, a baseline of mean time to responds (MTTR) with a SOAR platform should be collected to indicate how the organization reacts with current technologies to threats. With the added capabilities outlined in this operational concept, the system should have a reduced response time as it relates to time to bridge knowledge gaps within the organization.

With respect to reliability, availability, and maintainability, the added capabilities should not impact the baseline metrics for these performance criteria. However, the new functions do require additional points to monitor. For instance, the knowledge network should have loops to continuously query and update the model based on new information. Continuous data consumption regarding ticket information and analyst performance should keep the model up to date, which also supports the maintainability portion. Metrics regarding reliability should also include some user inputs regarding human perspectives of performance (usability, usefulness, correctness, etc.). Finally, in order to ensure minimal disruption, it is also important to monitor cognitive workload of the human during automation assistance for collaboration.

4.3. Measures of Effectiveness

Regarding Measures of Effectiveness in the traditional Systems Engineering sense, this document is limited in defining these due to the static nature of the functional architecture and the lack of executable models (Levis & Wagenhals, 2000). However, this section will address what might be useful to define and consider in the future as measures of effectiveness based on proposed human-system interaction.

With respect to key performance measures of human-automation teaming, it is especially important to include standard indicators of effective human-computer interaction. Examples of system performance to ensure usability can be found in (Nielsen, 1994). The system should work quickly to present the needed information to the user in a style or format that is appropriate, and efficiently and effectively facilitate a connection between a user and a knowledge source. Thus, metrics might include time to detect that the user needs knowledge, accuracy in knowledge prediction, accuracy in knowledge source identification, and efficiency and effectiveness of connection. Other proposed metrics for human-automation interaction include performance with respect to task complexity (Budiu & Whitenton, 2018), helpfulness to the user (Budiu &

Whitenton, 2018), analyst workload (Lathrop, 2017), human trust in the automation (Lathrop, 2017), effectiveness of learning of the system, and learning and comprehension of the user as they access new knowledge. Lastly, literature regarding the role of automation and expertise storage and retrieval can help guide expertise-based metrics and considerations for development (Buchanan et al., 2018).

In summary, effectiveness measures that support *usability* and *process efficiency* are of utmost importance. The human user should quickly gain situation awareness of incoming information, efficiently fill knowledge gaps, make a decision, and mitigate a threat. The system should aim to minimize disruption and be as helpful to the human as possible, while also learning from interactions to minimize explicit human training of learning models.

4.4. User and Organizational Issues

The *user types* of this functional set are mainly defined by levels of expertise. Users with lower levels of expertise (mainly subject matter or situational context expertise) will interact with the SOAR platform to conduct their tasks in incident response. If additional knowledge or expertise is needed, the platform would determine (through these functions) who or what other resource is needed, how to best connect them, and facilitate connection. Thus, the other user type is defined as experts in various areas of knowledge that would be contacted or connected through the platform to a novice that would need their help in understanding some aspect (or context) of an incident.

Training within the system should be minimal to reduce impact on actual work. Early in deployment, the system can have built-in survey mechanisms to "train" the system more efficiently, and rely less on human programming. After initial development and deployment, the system should be able to function autonomously as an assistant, but continue learning on its own with respect to the knowledge network and facilitating connections.

The human user maintains incident response decision authority for non-automated responses. The system can propose actions or facilitate connections, but the human maintains the responsibility of making a decision. The system should monitor and aim to reduce user workload during incident response, quickly help establish situation awareness (shared, if needed), and ensure trust and efficiency in the human-automation teaming activities.

5. Functional Architecture

The functional architecture is a visual representation of how the proposed system (or capability) will perform, and is built upon the user needs and operational concept (Levis & Wagenhals, 2000). Functional architecture is commonly determined before, and then intermittently, with physical architecture, such that the two aspects fit together to form a technical overview of how the system will function and how it might be constructed. Architecting is the foundation of model-based systems engineering (MBSE), and provides the means to create computational models of proposed systems for further analysis and design (Carson & Sheeley, 2013; Levis & Wagenhals, 2000). The researcher notes that some literature has described architecting as more of an art than a science, and can be considered a creative process (Emes et al., 2012). Thus, the functional architecture presented here should be considered a first draft of how the new system capability might function.

Using the Structured Analysis approach, the functional architecture is comprised of four (4) models and an integrated data dictionary to help describe different aspects of the system's operation and data flow between functions (Levis & Wagenhals, 2000).

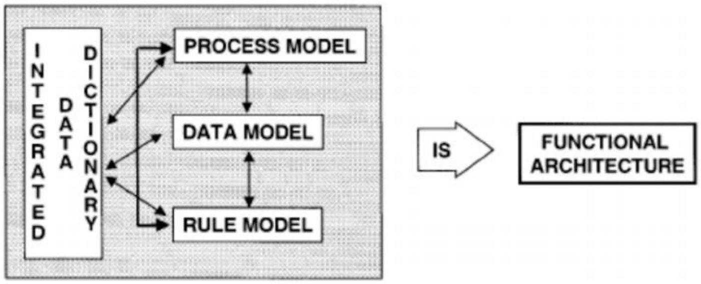

Figure.P.2. Components of Functional Architecture from (Levis & Wagenhals, 2000)

5.1. Activity Model

The activity model below depicts the higher-level function of "facilitate collaboration" as a set of three sub-functions. Inputs to this function include: 1) knowledge of the organizational and larger socio-technical system structure, 2) where knowledge sources exist (including what kind of knowledge and how much), 3) incident data that shows which analysts used which tools (or collaborated with what other analysts) to work on an incident that required a certain type of knowledge, and 4) user behaviors. These inputs support the creation and maintenance of a knowledge network that can employ machine learning to catalogue new data into a schema of knowledge existing within the organization. For instance, these inputs indicate where analysts go to get certain types of knowledge (including accessing other analysts), their performance on tickets (based on time, completeness, and if it needed to reopened), and if a user is confused or struggling to find an answer.

The outputs of the function include the knowledge network, identification of knowledge needs, and data supporting knowledge sharing patterns in the organization. The supporting functions are further described and discussed in the Functional Analysis document, but the interaction between functions is clearly depicted in Figure.P.3. Constraints include organizational or system structure (which could prevent analysts from accessing knowledge), policy regarding knowledge sharing in an organization (i.e. data classification for different clearance levels), and time.

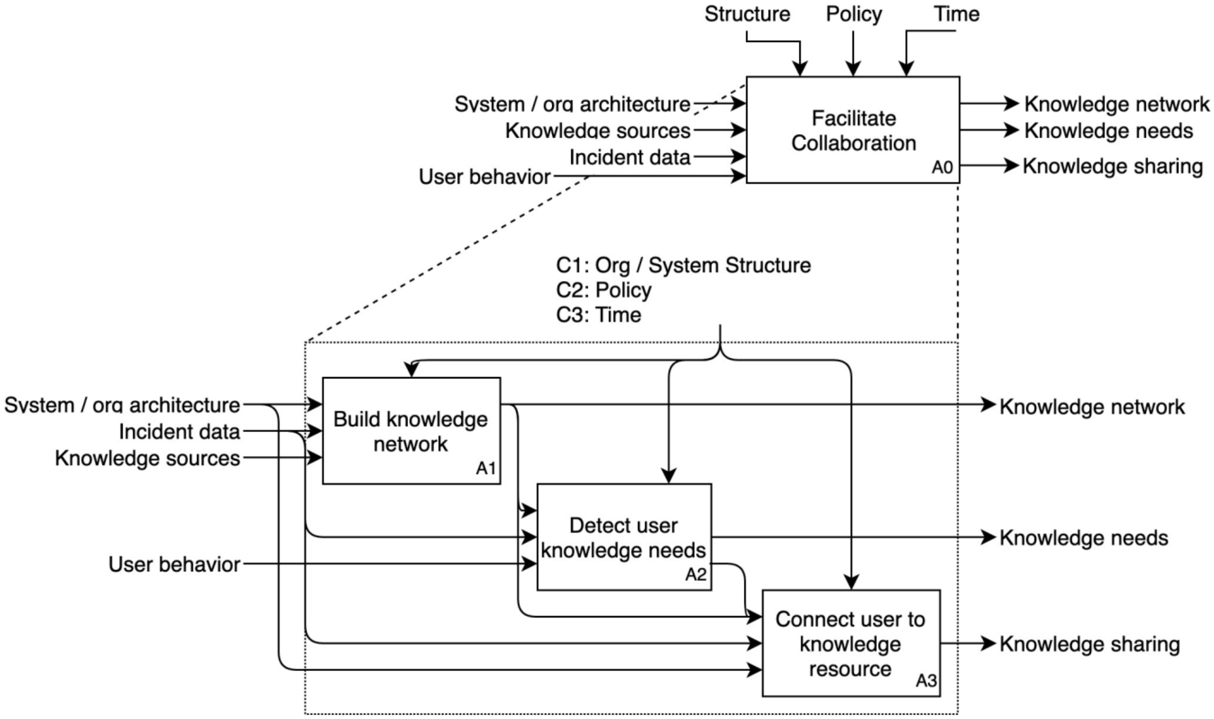

Figure.P.3. Activity model: Facilitating Collaboration

5.2. *Data Model*

The "source" is any entity that has or needs knowledge in a particular area of expertise (see *expertise* below). Entities can be human or non-human, and include higher tier analysts, databases, wikis, or experts outside the immediate SOC environment (e.g. legal, forensics, etc). Source attributes include modes for contact and contextual factors relating to contact, as well as behaviors that might indicate that knowledge supplementation is needed, and knowledge estimates from previous experience.

Behaviors, as mentioned above, may help predict when a user is in need of assistance related to finding and using knowledge. Behaviors can be determined unobtrusively from mouse movements, eye-tracking, and pupillometry, or help can be prompted from the user. Confirmation from the user that assistance is needed can also help validate behavior-deducing functions. Correlating behaviors with steps in the playbook for a particular incident can also help estimate what type of assistance is needed.

Incident data includes attributes such as type, expertise required, and urgency, and can also be correlated with playbook data. The incident data supports decisions regarding if assistance is needed (and on what timescale), what kind of assistance is needed, and how to best deliver it. Incident data is already recorded in most organizations through ticketing systems, which can be integrated with SOAR platforms.

Expertise data has dependent entities of four of the six dimensions of expertise (subject matter, communication, interface/tool, situational context). The other two dimensions, expert identification and information flow path, are integrated into this function to supplement those particular areas for analysts. Expertise data include a type (classifier) and amount. Expertise data can be estimated from past experience, vetted by known experts (i.e. checking tickets), or measured through the interface itself during incident response through factors like time, decisions made (compared to vetted experts), and robustness of solution. Further development of this data class is recommended.

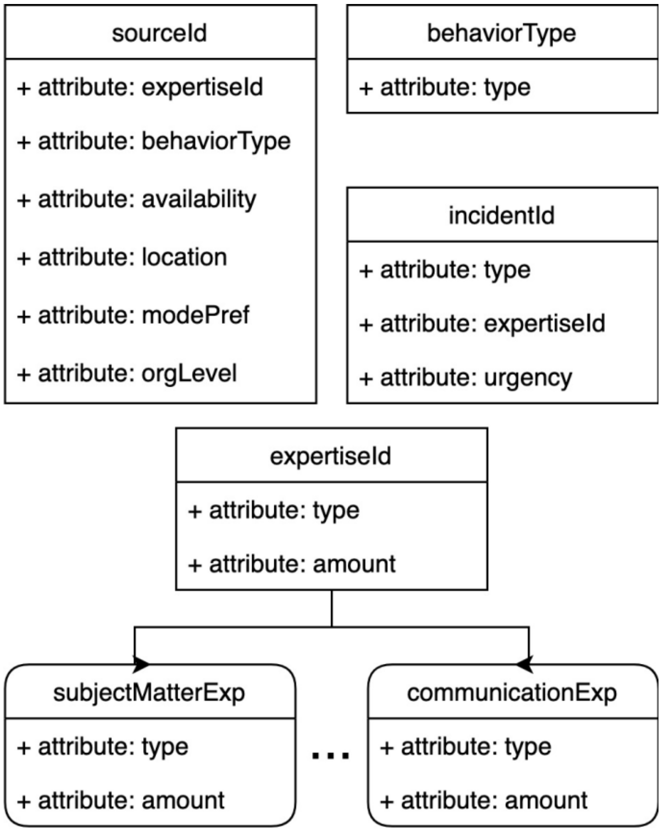

Figure.P.4. Data Model: Facilitating Collaboration

5.3. Rule Model

The function includes four (4) key decisions: 1) if help is needed, 2) what kind, 3) best source for helping, and 4) how to connect need to source. Each decision requires multiple inputs. Figure P.5 depicts decision inputs and outputs for the function.

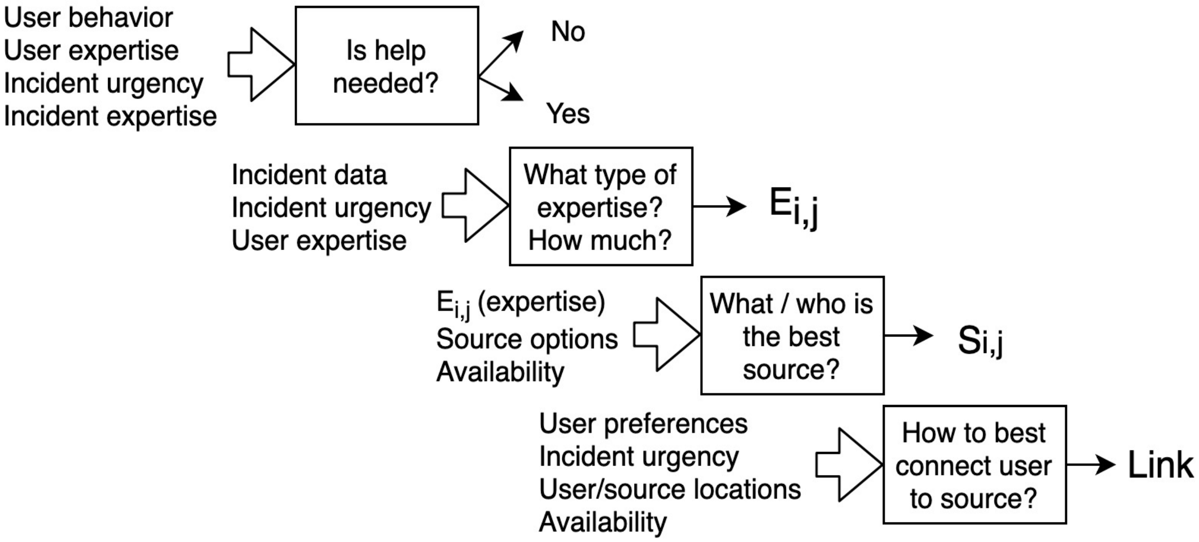

Figure.P.5. Rule Model Decisions

Some potential inputs to deciding if a user is in need of help include: user behaviors, user expertise, incident urgency, and incident expertise needed. Different inputs regarding user behaviors, such as mouse movement patterns, eye-tracking, and pupillometry, could indicate if an analyst is confused or stuck in solving a problem. Combined with known factors of the incident type (and expertise required), urgency, and user's experience with a particular incident type, the function would be able to decide whether or not assistance is needed. An example of how these inputs are related is depicted in Table P.1 below.

Table.P.1. Decision Factors: Determine Assistance

Behavior	User Expertise	Urgency	Incident Expertise	Help Needed?
Normal	Type: X, Level: High	Low	Type: X, Level: Low	No
Stuck	Low	High	High	Yes

Potential inputs to determining what kind of expertise is required include: incident type, incident expertise needed, urgency, and user expertise. Inputs might include multiple attributes, such as a measure or level of expertise that is needed. This function should correlate factors to predict what kind of assistance is needed by the user, whether its true expertise or access to certain information that could help with an investigation. An example of how these inputs are related is depicted in Table P.2 below.

Table.P.2. Decision Factors: Assistance Type

Incident Type	User Expertise	Urgency	Incident Expertise	Expertise Required?
Non-routine	Type: X, Level: Low	High	Type: X, Level: High	Type: X, Level: High
Routine	Type: Y, Level: High	Low	Type: Y, Level: Low	None

Potential inputs to determining the best source of expertise for the user to utilize include: expertise needed, source options, and source availability. In combination with the previous decision, this set of inputs helps the function decide what resource can best support the user's needs. This decision could actually result in multiple prioritized outputs, allowing the user to make the ultimate decision. An example of how these inputs are related is depicted in Table P.3 below.

Table.P.3. Decision Factors: Assistance Source

Expertise Needed	Source Options	Availability	Source Recommended?
X	1	High	1, Yes
Y	2	Low	2, No

Lastly, the function should determine the best way to connect the user to information. Potential inputs to linking user to source include: user preferences, urgency, location, and availability of the source. Contextual information helps determine availability of the analyst in need and the potential resource, which could be critical in deciding the best (most available) resource to connect to. An example of how these inputs are related is depicted in Table P.4 below.

Table.P.4. Decision Factors: Connection

User Preferences	Urgency	Location	Availability	Link Created?
Phone	High	Shared	High	Yes
Instant message	Medium	Distributed	Low	Yes

5.4. Dynamics Model

The states included in this function included *monitor, learn, assist,* and *connect,* as shown in Table P.5 and Figure P.6. The Monitor state is the base (or idle) state in which the system is ingesting information about ongoing incidents and learning from user behaviors. When not in any other state, the system defaults to Monitor. Incident data informs patterns in incident type, involved analysts, solutions, and outcomes. User behavior data indicates when a user might be in need of assistance. The next states from Monitor are Learn and Assist.

TableP.5. System States During Facilitate Collaboration Function

Current State	Event	Action	Next State
Monitor	New incident	Analyze incident data	Update / Learn
	User busy	Analyze user behavior	Assist
	System idle	Idle	Monitor
Update / Learn	New pattern	Update knowledge network	Monitor
Assist	User confusion	Propose help / idea	Connect
Connect	Source found	Facilitate connection	Update / Learn
	System idle	Idle	Monitor

The Learn state is the active state of learning for the system. In this state, the system updates models of knowledge in the network and gleans understanding from patterns in incident data. Learn triggers changes to “who works on what type of incident”, “how quickly did they resolve it”, “which incident types use which resources”, and “what patterns exist between users and resources”. The Learn state may also include some validation of expertise (i.e. resolution checking by a more experienced analyst) before changing user knowledge information. After the Learn state, the system returns to Monitor.

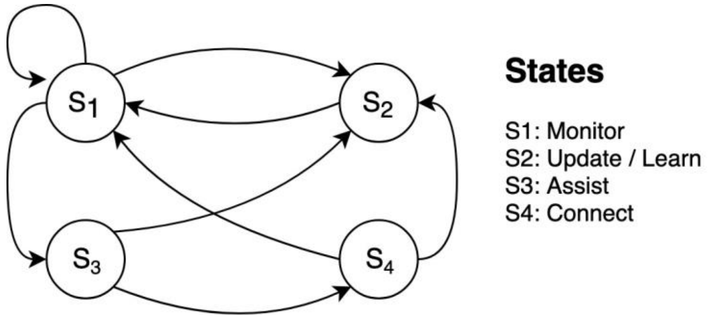

Figure.P.6. Facilitating Collaboration State Dynamics

When analyzing user behaviors, the Monitor state might transition to the Assist state when detecting confusion from the human. For instance, eye tracking might indicate the user is staring at a particular set of information and not acting, or continually glancing back and forth between two pieces of information for some time. This might indicate that the user is thinking, but not concluding the next action. In the same sense, monitoring the user’s pupils might indicate when the user is thinking hard or processing high amounts of information; and monitoring mouse movement might also indicate intent of the user as they user the pointer to direct their own focus. Furthermore, superimposing this data over playbook activities could indicate where the user is in the incident response process, further informing intent. These are just some examples of data that might be ingested and analyzed to deduce the user’s state of thinking and intent. When user confusion or need is detected, the system searches for potential resources that might help them overcome the problem, and proposes a connection. Thus, the following state is the Connect state.

The Connect state is triggered by a resource being identified for a user in need, and confirmation from the user that it is a source they would like to connect with. The Connect state actively links a user in need to the resource through an appropriate mode, which could be supplied by the system’s learning about how the knowledge network is connected (e.g. phone, email, instant message, platform chat). If the resource is not a person, but rather a database, the user will also receive assistance regarding navigating that database, if needed. The Connect state concludes with Learn and Monitor, returning back to the original state of Monitor after the new data is included during Learn.

5.5. Integrated Data Dictionary

Some of the terms, functions, and inputs for the four models are described explicitly within the respective sections. Other terms not explicitly defined are included in this section to help integrate terms across the models and solidify the functional architecture.

Link mode is the mode or channel by which two entities can be connected. This includes phone, email, instant message (within or outside the SOAR platform), and so on. Link mode may not always be facilitated by the system, but the system can recommend the mode, and perhaps prepare the entity being queried. For instance, the system might give a person a notification that someone will be calling about help on a particular incident, and offer to screen share or brief the person in advance).

Availability indicates the activity level of the entities to be connected. This can be a clear determination, referencing meeting schedules or out of office settings; availability can also be more nuanced, using monitoring features to detect if an entity is deep in another problem and does not want to be disturbed. Availability should be considered relative to urgency. If an incident is high urgency, then the entity's priorities might shift.

Urgency indicates how urgently the ticket needs to be addressed, which can be driven by company policy, classifications of severity, etc.