ACHIEVING CYBER DEFENSE SITUATIONAL AWARENESS: A COGNITIVE TASK ANALYSIS OF INFORMATION ASSURANCE ANALYSTS

Anita D'Amico [a], Kirsten Whitley [b], Daniel Tesone [a],
Brianne O'Brien [a], Emilie Roth [c]
[a] Secure Decisions, a division of Applied Visions, Inc.;
[b] U.S. Department of Defense; [c] Roth Cognitive Engineering
Contact: anitad@securedecisions.com

A Cognitive Task Analysis (CTA) was performed to investigate the workflow, decision processes, and cognitive demands of information assurance (IA) analysts responsible for defending against attacks on critical computer networks. We interviewed and observed 41 IA analysts responsible for various aspects of cyber defense in seven organizations within the US Department of Defense (DOD) and industry. Results are presented as workflows of the analytical process and as attribute tables including analyst goals, decisions, required knowledge, and obstacles to successful performance. We discuss how IA analysts progress through three stages of situational awareness and how visual representations are likely to facilitate cyber defense situational awareness.

INTRODUCTION

In response to computer viruses, network intrusions, and denial of service attacks, government agencies and large corporations have established network monitoring centers to detect and defend against security breaches in their critical information infrastructure. The US Department of Defense (DOD), which invests more resources in cyber defense than any vertical industry, refers to its cyber defenders using the terms information assurance (IA), computer network defense (CND), and defensive information operations (DIO) analysts.

These analysts comb through thousands of intrusion detection system (IDS) alarms during a single work shift to find real attacks amidst a preponderance of false alarms. They review massive amounts of packet data to identify suspicious and anomalous activities that were undetected by the IDS. They may also engage in forecasting new threats and predicting the activities of known attackers. They strive to attain and maintain situational awareness (SA) of the networks they defend and the attackers that attempt to gain access.

Our research used a cognitive task analysis (CTA) approach to gain insight into the work processes, cognitive skills, and tools that IA analysts rely on to achieve SA and the cognitive challenges and obstacles that impede SA. This paper highlights some findings of a full report; interested readers may contact the authors for further information.

RELATED WORK

Killcrece, Kossakowski, Ruefle, and Zajicek (2003) categorized cyber defense activities to assist organizations in setting up and staffing Computer Security Incident Response Teams (CSIRTs). As the term "response" implies, Killcrece et al. focused on how incident responders coordinate and analyze incidents that have already been detected. Although neither intrusion detection nor cognitive challenges was their main focus, their report offered a categorization of CSIRT functions that guided our early research.

Alberts, Dorofee, Killcrece, Ruefle, and Zajicek et al. (2004) extended the work of Killcrece et al. by documenting workflows considered to be best processes for effective incident management. Their models represent what incident response "should be" and do not necessarily represent the actual experiences of most IA analysts. By comparison, our CTA documented "what is," illuminated the cognitive processes that underlie the work processes, and identified opportunities to improve SA.

Biros and Eppich (2001) studied rapid intrusion detection analysts in the US Air Force, using CTA, and identified five requisite cognitive abilities: recognizing non-local Internet Protocol (IP) addresses, identifying source IP addresses, developing a mental model of normal, and sharing knowledge. We used their work as a starting point for an analysis of additional roles beyond rapid intrusion detection and beyond the Air Force.

METHOD

Participants

We interviewed and observed 41 analysts from one commercial and six DOD organizations responsible for network security. Participants varied in level of expertise from novice to expert and represented a variety of IA roles including all of the analysis roles described below.

Data Collection Methods

The CTA used a bootstrap approach to examine the different roles of IA analysts (Potter, Roth, Woods, and Elm, 2000). We employed a combination of four knowledge capture techniques: structured interviews, observations, review of critical incidents, and hypothetical scenario construction. The hypothetical scenario construction technique was added to the data collection because issues of confidentiality and classification sometimes prevented subjects from sharing details of tough cases. To mitigate this problem, we devised a hypothetical scenario exercise that involved working with analysts to flesh out an imaginary analysis case including typical offensive actions taken by a sophisticated attacker and defensive actions by the IA analyst. The exercise allowed analysts to reveal the kinds of information they seek from available data sources, knowledge of adversary operations and techniques, and types of connections that analysts make between seemingly disparate pieces of information.

Data Analysis Methods

Workflow diagrams and attribute tables were used to summarize the data and abstract a picture across organizations. Workflow diagrams were used to describe the work processes of each organization. After comparing and contrasting across organizations, we designed a diagram to capture the generalized workflow (Figure 1). Attribute tables were used to capture the analysis details. The tables enumerate important categories about each major IA analysis task. Attributes included: analyst goals, decisions, heuristics, cognitive requirements, general knowledge, site-specific knowledge, rapidly changing knowledge, sources of direction, data sources, tools and visual aids, communication requirements, work environment, and obstacles to performance.

RESULTS

Analysis Roles

Job titles varied considerably from organization to organization, yet there was strong similarity in underlying analysis tasks and separation of duties. We classified analysis functions into six analysis roles that describe significant, distinct portions of the entire analysis process. *Triage analysis* is the first look at the data. Triage encompasses weeding out false positives and escalating suspicious activity for further analysis, all within a few minutes of viewing the data. *Escalation analysis* investigates potential incidents based on escalations from triage and also tip-offs from colleagues and cooperating organizations. Escalation analysis takes hours to weeks, during which the analyst marshals more data, usually from multiple data sources and from inside and outside the organization, resulting in greater comprehension of the attack modus operandi, extent, severity, and goal. *Correlation analysis* searches for patterns and trends in current and historical data. At the community level, correlation analysis includes grouping and investigating related incidents, a task that can take weeks to months. *Threat analysis* is intelligence analysis in support of CND, using additional data sources (e.g., information from hacker websites) to profile attackers; these additional data sources are required to discover an attacker's true identity and motivation. *Incident response analysis* recommends and/or implements a course of action in reaction to a confirmed incident. *Forensic analysis* gathers and preserves evidence in support of a law enforcement investigation.

Killcrece et al. (2003) distinguish between proactive and reactive analysis. Tasks performed in anticipation of a future attack are proactive; tasks performed in response to an attack that has occurred are reactive. Of the six analysis roles identified in this CTA, threat analysis most strongly aligns with the proactive category. The other five are largely or entirely reactive, as currently practiced by IA analysts.

Stages of IA Cognitive Data Fusion

We observed IA analysis progressing through three stages of cognitive data fusion that parallel the Joint Directors of Laboratories (JDL) levels of data fusion (Llinas and Hall, 1998; Waltz, 1998). Figure 1 illustrates these stages.

JDL Level 1: Detection. The initial stage of IA analysis, primarily concerned with initial data inspection and detection, aligns with JDL Level 1. An analyst inspects and associates elements of sensor data at the network packet-level to detect suspicious activities.

JDL Level 2: Situation Assessment. As analysis is escalated, an analyst incorporates other sources of data (e.g. host-based logs, results of web searches, reports from other analysts) into the fusion process to refine his assessment. The analyst extracts features from the data in the form of patterns of suspicious activity related to the attacker (source IP), the target (destination IP), and the attacker's method. The analyst interprets the data within the context of the operational environment.

JDL Level 3: Threat Assessment. Once incidents are confirmed, IA analysis moves to the correlation and threat analysis stages during which an analyst draws inferences about the actual enemy identities, motives, and sponsorship.

Stages of IA Situational Awareness

The IA analytical process also moves through three stages of SA, namely, perception, comprehension, and projection (Endsley, 1995; Endsley, Bolte, and Jones, 2003). All three elements of SA occur for individual analysts within the IA analytical pipeline. However, at each of the three cognitive data fusion stages, an organization's overall SA transitions from one key element to the next.

SA Stage 1: Perception. During the first cognitive fusion stage (detection), an analyst acquires data about his environment, which is typical of the perceptual element of

SA. An analyst inspects data to detect interesting activity on the network. As the analysis unfolds, comprehension begins. An analyst combines and integrates data to form a mental model of how the interesting network activity might represent an attacker's action. By testing various hypotheses through additional data and input from other analysts, an analyst modifies and clarifies his mental model. By the end of the first stage, when the analyst decides whether to escalate, the focus shifts from perception to comprehension.

SA Stage 2: Comprehension. During the second cognitive fusion stage (situation assessment), an analyst performing escalation analysis first reviews the suspicious activity (perception). The analyst then combines and integrates his knowledge and experience with additional data sources to determine whether the suspicious activity represents an actual incident. He refines the mental model of the attacker's identity and threat level as he traces the attacker's path through the network back in time. Also in this stage, an analyst performing correlation analysis identifies and reports on patterns of suspicious and anomalous behavior; the reports serve to cue other analysts. This escalation and correlation analysis represents the comprehension aspect of SA. SA Stage 2 also involves limited projection. Escalation analysis includes some postulating about an attacker's actions if left unblocked. Incident responders, in choosing a course of action, project what future actions an unblocked attacker could take and what actions an attacker might take if he realizes he has been discovered.

SA Stage 3: Projection. During the third stage (threat assessment), analysts performing correlation, incident response, and threat analysis review and categorize confirmed incidents across organizations for common features at the community level. By comparing incidents and adding data from intelligence sources, they discern attack patterns. By this point, the analysts at the community level have a *shared* mental model. As comprehension improves, the analysts refine their shared mental model and project into the future to forecast the types of incidents to expect within the community. Furthermore, proactive threat analysis identifies potentially new exploits and attackers that, while currently unobserved, could become active in the future. Based on these projections, tips are fed back into the start of the analysis pipeline as perceptual cues for detection. This forms a perception-action cycle across multiple, distributed actors.

Cognitive Challenges

IA analysts face extensive cognitive challenges, four of which are mentioned here. These challenges reflect the inherent difficulty of the analysis domain (e.g., the popularity and relative insecurity of the Internet) as well as inadequacies in current analysis tools.

Massive Data. IA analysts are overwhelmed with data, one reason being that current-day sensor technology has a high rate of false positives. To process the data, organizations and analysts must use data minimization strategies such as omitting all traffic to and from certain IP addresses. As sensors become overwhelmed with a flood of data, analysts minimize data flow by modifying sensor rules to accept certain types of network activity as legitimate. Data minimization strategies artificially restrict potentially valuable data and consequentially could result in incomplete or inaccurate analysis and failure to detect malicious activity.

IP Addresses	• Is IP external or internal to network? • Is IP on the "Hot IP List"? • Is IP assigned to a suspicious entity or known competitor? • Is the IP address spoofed?
Ports / Protocols	• Is there increased activity on a specific port? • Are there violations of expected port/protocol associations? • Is there a deviation from expected protocol behavior? • Is there activity on a previously blocked port?
Packet Content	• Does the size of data payload exceed a threshold of interest? • Does the packet contain specific strings that match an attack signature? • Are byte counts unusually consistent?
IP Behavior	• Are there other alerts associated with this destination IP, source IP or subnet? • Did an internal IP send outbound traffic in response to a scan or attack attempt? • Is an internal IP sending a lot of outbound traffic and/or content? • Who has the source (destination) IP talked to? Who has talked to the source (destination) IP? • What services is an IP running? • Are there changes to the type/number of peers/protocols in a time period? • Is IP exhibiting multiple roles (e.g., acting as both a client and a server)?
Temporal Issues	• Time: When did the activity occur? What else was occurring at that time? • Duration: How fast did the activity occur? Was there indication of human involvement or was the attack completely automated? • Repetition: Is there evidence of regularly repeating activity?
Across Incidents	• Do disparate incidents share common features (e.g., IP and subnet addresses, port numbers, times of day, names and details of the attack tools, attacker sequences of steps, attacker login names and passwords, attacker hiding places in the file system)?

Table 1: Typical IA Analysis Questions

Fusion of Complex Data. Indicators of an attack may occur in a variety of data sources, which an analyst must mentally fuse. Armed with tools to sort and filter data, he seeks to answer certain analytic questions (See Table 1) and to formulate associations between multiple data sources. However, today's tools are limited, especially in support for data fusion and correlation across incidents; therefore, the analysts are the correlators of the heterogeneous data.

Building Site-Specific Knowledge. Site-specific knowledge is at the core of many tasks during the detection and situation assessment stages. Analysts must first formulate over time a mental model of what is normal for each site they are monitoring and then look for deviations from normal. There are several challenges with this process: it often takes months to determine what is normal for a particular monitored network; what is normal changes over time, so an analyst must constantly adapt his model of normal; and it is difficult for an analyst to articulate his model of normal to others, thereby slowing the learning rate of new analysts and impeding the communication of results to others.

Maintaining Multiple Mental Models. Analysts simultaneously track many potential attackers and incidents. For each track, an analyst creates a mental model of an attacker, and for each model the analyst hypothesizes the attacker's identity, motive, and modus operandi. At various points, particularly at the threat assessment stage, separate mental models will converge, and the analyst then combines the similar tracks into a unified attacker profile. The challenge for analysts is to formulate and maintain dozens of these mental models with few external aids.

DISCUSSION

The results of this CTA identified a broader set of roles and more comprehensive collaborative process than described by previous research. Consistent with the findings of Biros and Eppich (2001) (who primarily studied what we term triage analysis), we identified cognitive tasks that are essential to the detection and situation assessment stages of IA cognitive data fusion: identification and recognition of IP addresses; developing and updating a mental model of normal network traffic; and acquiring and sharing a variety of knowledge including knowledge of emerging vulnerabilities and new exploits. Our results expand upon the types of knowledge requirements and point to a wider set of cognitive activities in which IA analysts engage.

The detection stage of IA cognitive data fusion maintains its focus on the role of rapidly detecting suspicious activity within a specific set of data. As analysis progresses to situation assessment, other analysis roles are activated and more data is analyzed. Escalation analysis considers additional hypotheses and fuses more data. Correlation analysis explores for previously unrecognized patterns and trends. Threat analysis introduces intelligence about the attacker, and incident response analysis adds evidence collection issues to the mix. The cognitive and collaborative requirements are more complex at this stage and are broader than those described by Biros and Eppich (2001) or Alberts et al. (2004). For example, analysts reach outside their own enclave to other organizations for evidence to evaluate their hypotheses, and they collaborate with system administrators at remote sites to assess the potential impact of an attack.

Unlike previous literature, our CTA documented activities and cognitive requirements at the threat assessment stage of cognitive fusion, when incidents are correlated across organizations in a community. During this stage, goals shift to revealing the true identity and motivation of the attackers and to predicting their next actions by correlating activity over long periods of time (weeks to months) and space (across multiple organizations) and by projecting into the future. Given the scarcity of research about this third stage, it is fertile ground for further inquiry and analysis of our data.

One goal of this CTA was to examine the current and potential use of visualization aids in IA analysis. Overall, we observed relatively little visualization use in IA analysis. During our data collection, we observed that visualization tools were used only for correlation analysis and threat analysis and for creating post-analysis pictures to be included in informational briefings. We also observed that analysts used visualization tools borrowed from non-IA applications, adapted those tools to fit their needs and often modified their data to fit the tool. We concluded that, to be effective for IA analysis, visualizations should be chosen or designed to align with an analyst's role and specific data formats.

The distinct tasks and cognitive requirements of each analysis role and analytic stage indicate a need for role-based visualization aids. Because triage analysis is focused on specific data and must be performed within minutes, visual aids should cue the analyst's attention for where to search in the voluminous data. Techniques such as highlighting data from "Hot IPs" or summary graphs showing most active ports and IP addresses could focus the triage analyst's attention on potentially suspicious activity. By comparison, correlation and escalation analysis is less time-constrained and requires exploration of both recent and historical data as well as data drawn from several sources in order to discover patterns. Visualizations such as time walls or draftsman's plots, which depict a matrix of scatter plots that permute the variables mapped to the axes, may enhance the analyst's ability to see patterns and time trends across multiple data sets. Threat analysis may be facilitated by animations and replay of events across the network, from which the analysts can deduce the progression of an attack and infer the speed and direction of the attacker's next action.

The development of visualization aids based on the findings of the CTA is the focus of the next phase of this research.

ACKNOWLEDGMENTS

The research described in this document was sponsored by the US Department of Defense and the Advanced Research and Development Activity (ARDA) under contract # F30602-03-C-0260, with the Air Force Research Laboratory as the contracting agency. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of ARDA or the US Government.

REFERENCES

Alberts, C., Dorofee, A., Killcrece, G., Ruefle, R., and Zajicek, M. (2004). *Defining Incident Management Processes for CSIRTS: A Work in Progress,* Technical Report CMU/SEI-2004-TR-015, ESC-TR-2004-015.

Biros, D. and Eppich, T (2001). Human Element Key to Intrusion Detection. *Signal*, August, 2001, p. 31.

Endsley, Mica R. (1995). Toward a Theory of Situation Awareness in Dynamic Systems, *Human Factors*, 37(1) 32-64.

Endsley, M.R., Bolte, B., and Jones, D. (2003). *Designing for Situation Awareness: An Approach to User-Centered Design*. New York, NY: Taylor and Francis. pp. 13-18

Killcrece, G., Kossakowski K.P., Ruefle, R., and Zajicek, M. (October 2003). *State of the Practice of Computer Security Incident Response Teams (CSIRTS)*, Technical Report CMU/SEI-2003-TR-001, ESC-TR-2003-001.

Llinas, J. and Hall, D.L. (1998). An Introduction to Multi-Sensor Data Fusion. IEEE Report 0-7803-4455-3/98

Potter, S. S., Roth, E. M., Woods, D. D., and Elm, W. (2000). Bootstrapping multiple converging cognitive task analysis techniques for system design. In J. M. Schraagen, S. F. Chipman, and V. L. Shalin (eds.) *Cognitive Task Analysis* (pp. 317-340). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Waltz, E. (1998). Information Understanding: Integrating Data Fusion and Data Mining Processes. IEEE Report 0-7803-4455-3/98.

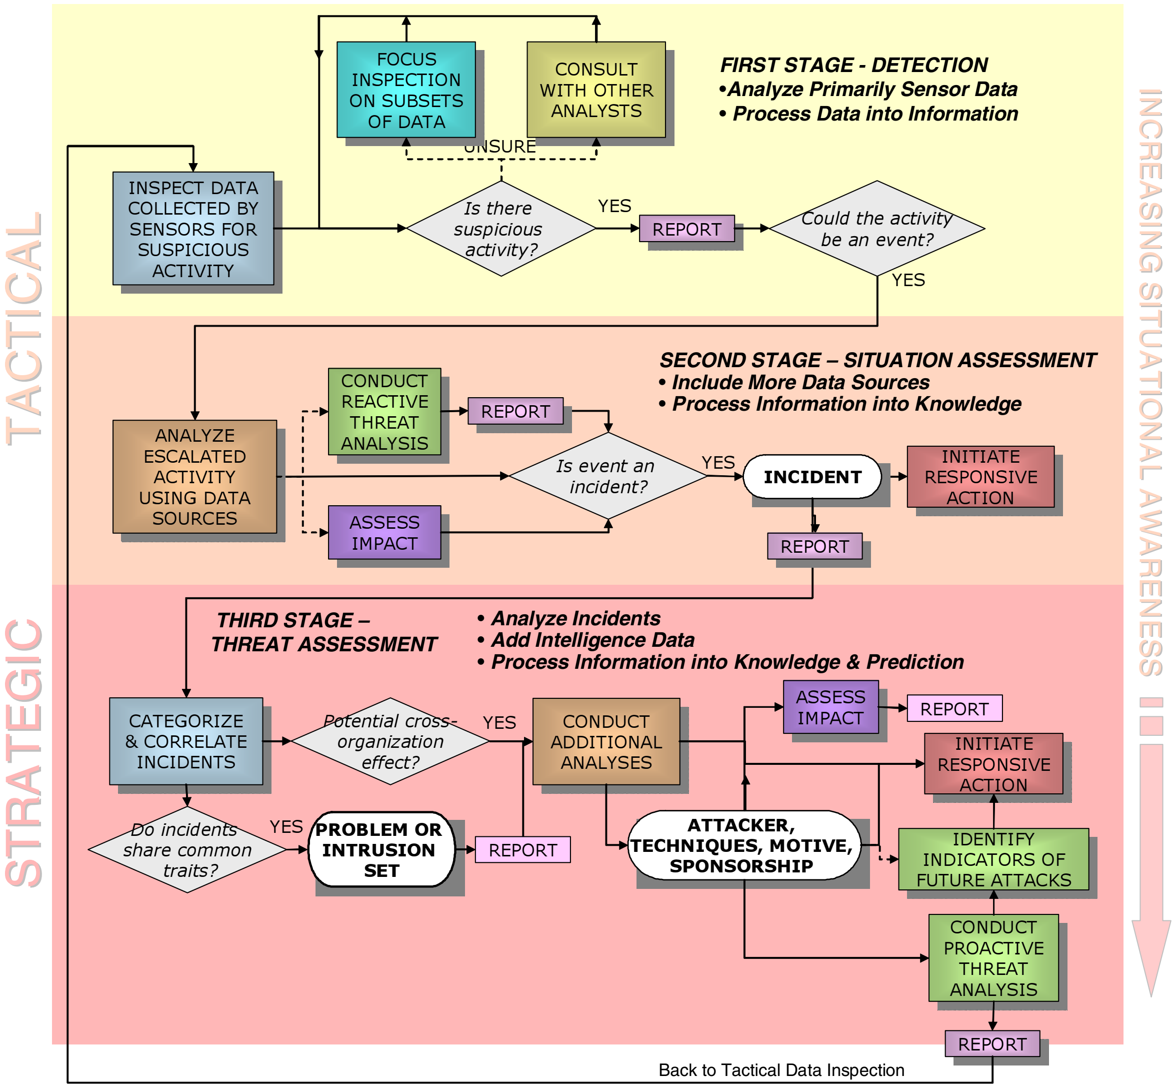
TACTICAL
STRATEGIC
INCREASING SITUATIONAL AWARENESS
FIRST STAGE - DETECTION
•Analyze Primarily Sensor Data
• Process Data into Information
FOCUS INSPECTION ON SUBSETS OF DATA
CONSULT WITH OTHER ANALYSTS
UNSURE
INSPECT DATA COLLECTED BY SENSORS FOR SUSPICIOUS ACTIVITY
Is there suspicious activity?
YES
REPORT
Could the activity be an event?
YES
SECOND STAGE – SITUATION ASSESSMENT
• Include More Data Sources
• Process Information into Knowledge
ANALYZE ESCALATED ACTIVITY USING DATA SOURCES
CONDUCT REACTIVE THREAT ANALYSIS
REPORT
Is event an incident?
YES
INCIDENT
INITIATE RESPONSIVE ACTION
ASSESS IMPACT
REPORT
THIRD STAGE – THREAT ASSESSMENT
• Analyze Incidents
• Add Intelligence Data
• Process Information into Knowledge & Prediction
CATEGORIZE & CORRELATE INCIDENTS
Potential cross-organization effect?
YES
CONDUCT ADDITIONAL ANALYSES
ASSESS IMPACT
REPORT
INITIATE RESPONSIVE ACTION
Do incidents share common traits?
YES
PROBLEM OR INTRUSION SET
REPORT
ATTACKER, TECHNIQUES, MOTIVE, SPONSORSHIP
IDENTIFY INDICATORS OF FUTURE ATTACKS
CONDUCT PROACTIVE THREAT ANALYSIS
REPORT
Back to Tactical Data Inspection

Figure 1: Three Stages of Cognitive Data Fusion and Situational Awareness in IA Analysis