



Report – Summary of Key Findings

*Troy Townsend, Melissa Ludwick, Jay McAllister,
Andrew O. Mellinger, Kate Ambrose Sereno*

September 2013 (a revision of the January 2013 Report)

Report – Summary of Key Findings

Table of Contents

Executive Summary	2.3
Introduction	2.4
Participants	2.4
Cyber Intelligence Definition and Analytic Framework	2.4
Baseline and Benchmarking Approach	2.5
Key Findings	2.6
State of the Practice in Cyber Intelligence	2.7
Challenge: Applying a strategic lens to cyber intelligence analysis	2.7
Challenge: Information sharing isn't bad; it's broken.	2.7
Best Practice #1: Aligning functional and strategic cyber intelligence resources.	2.8
Environment	2.9
Challenge: Understanding threats to the software supply chain	2.9
Challenge: Determining where cyber intelligence belongs organizationally	2.10
Best Practice #1: Scoping the cyber environment to the organization's mission	2.10
Best Practice #2: Modeling threats to shape resource allocation	2.10
Data Gathering	2.11
Challenge: Data hoarding	2.11
Challenge: Lack of standards for open source intelligence data taxes resources.	2.12
Best Practice #1: Repurposing search engine referral data	2.12
Best Practice #2: Mind the gaps	2.12
Functional Analysis	2.13
Challenge: Adopting a common cyber lexicon and tradecraft	2.13
Challenge: Filtering critical cyber threats out of an abundance of data	2.14
Best Practice #1: Comprehensive workflow to identify cyber threats and inform customers ..	2.14
Best Practice #2: Producing scripts to automate the filtration of known threat data.	2.14
Strategic Analysis	2.15
Challenge: No industry standard for cyber intelligence education and training	2.15
Challenge: Adapting traditional intelligence methodologies to the cyber landscape	2.16
Best Practice #1: Know your enemy	2.16
Best Practice #2: Global situational awareness.....	2.16
Decision Maker Reporting and Feedback	2.17
Challenge: Communicating "cyber" to leadership	2.17
Challenge: Difficulty capturing return on investment	2.17
Best Practice #1: Failure analysis	2.18
Best Practice #2: Carving channels for communication.....	2.18
Conclusion	2.19

SEI Emerging Technology Center Report Cyber Intelligence Tradecraft Project Summary of Key Findings

Executive Summary

The Software Engineering Institute (SEI) Emerging Technology Center at Carnegie Mellon University studied the state of cyber intelligence across government, industry, and academia from June 2012 to September 2013. The study, known as the Cyber Intelligence Tradecraft Project (CITP), sought to advance the capabilities of organizations performing cyber intelligence by elaborating on best practices and prototyping solutions to shared challenges. Six government agencies and 24 organizations from industry and academia provided information on their cyber intelligence methodologies, technologies, processes, and training. This baseline data then was benchmarked against a cyber intelligence analytic framework consisting of five functions: environment, data gathering, functional analysis, strategic analysis, and decision maker reporting and feedback. The aggregated results of the benchmarking led to the key findings presented in this report.

Overall, the key findings indicated that organizations used a diverse array of approaches to perform cyber intelligence. They did not adhere to any universal standard for establishing and running a cyber intelligence program, gathering data, or training analysts to interpret the data and communicate findings and performance measures to leadership. Instead, pockets of excellence existed where organizations excelled at cyber intelligence by effectively balancing the need to protect network perimeters with the need to look beyond them for strategic insights. Organizations also continuously improved data gathering and analysis capabilities with threat prioritization models, information sharing, and conveying return on investment to decision makers. This report captures the best practices from successful cyber intelligence programs and tailors them to address challenges organizations currently face.

Introduction

Cyber intelligence grew from the halls of government into a burgeoning business providing tools and services to industry and academia. As more organizations focused on this topic, varying methodologies, technologies, processes, and training complicated the operating environment. Recognizing a need to understand and improve this situation, the SEI Emerging Technology Center began to study the state of the practice in cyber intelligence in June 2012. This report discusses the Cyber Intelligence Tradecraft Project's (CITP) process and key findings.

Participants

The CITP involved 30 organizations. They included six government agencies with dedicated cyber intelligence missions and 24 entities representing multiple economic sectors, such as academia, defense contracting, energy, financial services, healthcare, information technology, intelligence service providers, legal, and retail. These organizations ranged in size from one employee to global organizations with hundreds of thousands of network users. Their cyber intelligence workforces had diverse backgrounds in intelligence, information security, and the military, and held a multitude of titles, such as chief technology officer, chief information security officer, vice president of threat management, information architect, intelligence analyst, and network analyst.

Cyber Intelligence Definition and Analytic Framework

The SEI Emerging Technology Center developed a definition of cyber intelligence to standardize the scope of the CITP with participants: *The acquisition and analysis of information to identify, track, and predict cyber capabilities, intentions, and activities that offer courses of action to enhance decision making.*

An analytic framework also was created to guide the CITP's baseline and benchmark processes, the foundation of which was based on the U.S. government's traditional intelligence cycle.




Figure 1 – Traditional Intelligence Cycle ²

The CITP's analytic framework promptly deviated from the preceding because the utility of the traditional intelligence cycle is limited when applied to cyber. The traditional intelligence cycle is depicted as a linear process and does not emphasize the inter-related nature of its five functions or their relevance to related functions, namely cyber security. The SEI Emerging Technology Center captured these unique cyber intelligence analysis characteristics by creating an approach that more accurately showed the inter-dependencies and outside influences in the cyber intelligence process. This approach incorporated how technology influences the way analysis is done, and uniquely identified the functions that integrate technology. In particular, the CITP's analytic framework separated analysis into two distinct functions: specialized technical analysis (i.e. functional) and strategic analysis.

² The Traditional Intelligence Cycle was reproduced from a paper authored by Judith Meister Johnston and Rob Johnston, hosted on the Central Intelligence Agency's public website: https://www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/books-and-monographs/analytic-culture-in-the-u-s-intelligence-community/page_46.pdf. Last accessed January, 2013.

This analytic framework utilized five functions to capture inter-dependencies of and external influences on cyber intelligence:

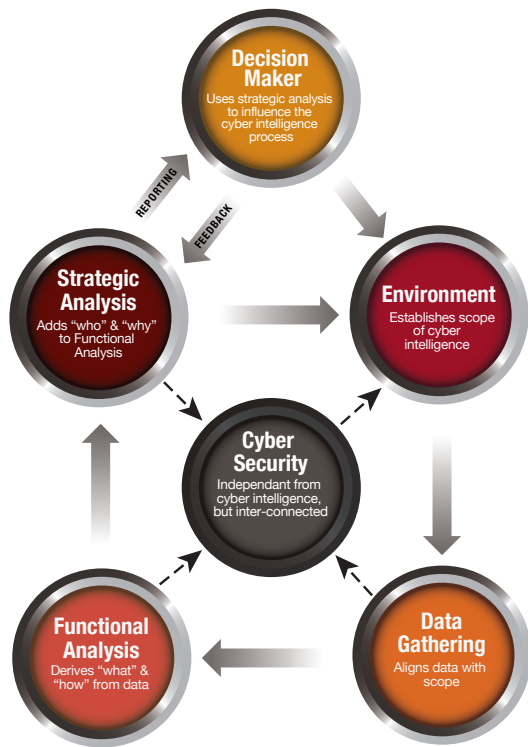


Figure 2 – CITP Analytic Framework

- **Environment:** Establishes the scope of the cyber intelligence effort and influences what data is needed to accomplish it.
- **Data Gathering:** Through automated and labor-intensive means, analysts explore data sources, collect information, and aggregate it to perform analysis.
- **Functional Analysis:** Analysts use gathered data to perform technical and tailored analysis, typically in support of a cyber security mission.
- **Strategic Analysis:** Analysts apply a strategic lens to functional data and report this intelligence to a decision maker or use it to influence the environment. If functional analysis attempts to answer the “what” and “how” of cyber threats, strategic analysis aims to answer “who” and “why.”
- **Decision Maker Reporting and Feedback:** After intelligence is disseminated to decision makers, they provide feedback and/or use the intelligence to influence the environment.

It is important to note that the analytic framework does not solely exist to address cyber security. Cyber intelligence is a critical component of cyber security, and the two functions are inter-related; however, the CITP focused on cyber intelligence. Cyber intelligence supports a variety of missions in government, industry, and academia; to include national policy, military applications, strategic communications, international negotiations, acquisitions, risk management, and physical security. Throughout the analytic framework, cyber security professionals receive data and intelligence, but the cyber intelligence process operates independently and does not necessarily need to support a cyber security mission.

Baseline and Benchmarking Approach

The SEI Emerging Technology Center employed an iterative process to create a discussion guide that served as a starting point to baseline organizations for the CITP. It reduced biases and was specifically designed to capture entities’ core cyber intelligence functions, regardless of if they represented government, industry, or academia. Using the discussion guide, the SEI Emerging Technology Center typically sent a cross-functional team of intelligence and software engineering professionals to engage with organizations during face-to-face interview sessions. The team interacted with representatives from cyber intelligence and cyber security leadership as well as functional and strategic analysts. During the interview sessions, these entities provided information on the methodologies, technologies, processes, and training enabling them to perform cyber intelligence.

The data gathered during these interviews established the baseline used to benchmark organizations against the CITP’s cyber intelligence analytic framework. For benchmarking, the SEI Emerging Technology Center compiled and reviewed the baseline to ensure it captured the pertinent data. The information then was ranked against 35 assessment factors distributed amongst the analytic framework’s five functions using an ordinal scale of ++, +, 0, -, --, with 0 representing average performance. Due to the variety in the organizations’ backgrounds and sizes, the ordinal scale offered the necessary flexibility for benchmarking, despite its limitations with numerical and interval analysis. Peer and group reviews also ensured consistency throughout the rankings.

The 35 assessment factors were derived from the interview sessions and SEI Emerging Technology Center cyber intelligence and software engineering expertise:

- **Environment:** Top-sight on cyber footprint; cyber intelligence distinction with cyber security; role alignment; personnel to support cyber intelligence; organizational structure; workflow utilization; prioritization of threats; organizational situational awareness; cyber intelligence functional and strategic analysis; scope of past, present, and future analysis; insider threat and cyber intelligence relationship.
- **Data Gathering:** Requirements and sources relationship; information sharing; meeting analytical needs; technology facilitating data gathering; indexing and archiving of data; validation of sources.
- **Functional Analysis:** Workflow exists; timeliness in producing analysis; diversity with incorporating multiple technical disciplines; skills, knowledge, and abilities; tools utilized.
- **Strategic Analysis:** Distinguished from functional analysis; workflow exists; diversity with incorporating multiple technical disciplines; skills, knowledge, and abilities; tools utilized.
- **Decision Maker Reporting and Feedback:** Report types generated; reporting mechanism for actionable and predictive analysis; leadership influences format and production timelines; cyber intelligence influences decision making; feedback mechanisms exist; feedback influences data gathering and analysis; satisfying intelligence consumers; capturing return on investment.

Key Findings

The following highlights the common challenges and best practices identified during the CITP by describing them within the context of the analytic framework's five functions. A stacked bar chart accompanies each function to summarize the baseline of organizations ratings in these areas. Each bar within the charts represents one of the benchmark's 35 factors (X-axis). The height of each color within the bars shows the percentage of organizations (Y-axis) receiving that particular rating and the red-colored diamond symbol displays the median. The ratings range between --, -, 0, +, and ++, with 0 being average performance for that assessment factor.

Figure 3 divides a stacked bar chart by the five functions of the analytic framework to visually show the CITP's baseline. Figure 4 removes the median (the red-colored diamond symbol) and the yellow-colored bar sections depicting the percentage of organizations receiving an average rating in Figure 3 to highlight the variances among entities with ratings of --, -, +, and ++. Figures 6, 9, and 11-13 display a stacked bar chart for the factors within each of the five functions.




Figure 3 – CITP Baseline




Figure 4 – CITP Baseline Variances

State of the Practice in Cyber Intelligence

Most organizations identified cyber intelligence and cyber security as two distinct and capable work functions that interact when necessary to best support their needs. They performed cyber intelligence by trying to understand the internal and external environment, gathering data, and analyzing technical threats, ranging from malware to email phishing. However, their intelligence reporting generally did not include strategic analysis or adequately inform decision makers, limiting its impact beyond the realm of cyber security. This exhibits an endemic problem of functional analysts not effectively communicating with non-technical audiences. It also demonstrates organizations reluctance to share information within their own organizations, industries, and across economic sectors.

Challenge: Applying a strategic lens to cyber intelligence analysis

Despite having a wealth of data available, many organizations struggled with moving beyond the functional analysis of low-level network data to incorporate strategic analysis of threats and threat indicators.

Current state:

- Most organizations had difficulty incorporating strategic intelligence analysis into existing security-focused processes. Correspondingly, entities with poor or no strategic analysis functions struggled with communicating security requirements to leadership, had a more reactionary network security posture, and were less likely to anticipate or be prepared for emerging cyber threats. This can be attributed to an organization-wide lack of support for strategic analysis, demonstrated by organizations not having the resources to index and store data for strategic analysis, perform trend analysis, or look at individual network events in a more strategic context. Some organizations could not obtain resources to address these issues because of an inability to effectively communicate the complexities of cyber intelligence to non-technical decision makers and relate its importance to the organization's overarching goals and objectives. Thus, decision makers did not grasp the benefits of investing in tools and personnel, and cyber intelligence efforts suffered.
- Organizations generally had a mature cyber intelligence workflow that incorporated functional analysis, but only as a means to support cyber security (see Figure 5). The challenge within this version of the analytic framework is communicating the importance and relevance of technical issues to decision makers in compelling terms they understand. Although cyber security benefits from functional analysis, the CITP's findings indicated that the addition of strategic analysis to the analytic framework was the most effective means of bridging the communication gap between cyber security and non-technical decision makers.

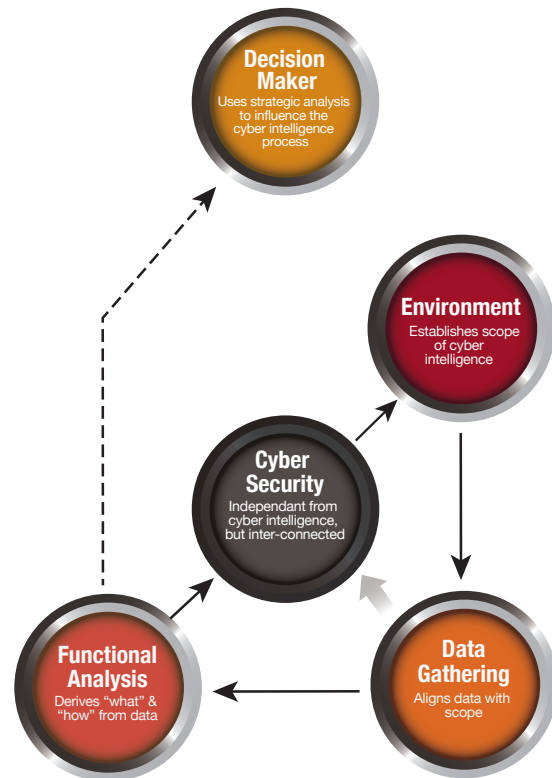


Figure 5 – Cyber Security-Centric Analytic Framework

Challenge: Information sharing isn't bad; it's broken

The highest performing organizations actively shared—not just consumed—data in formal and informal information sharing arrangements.

Current state:

- Government organizations in the CIP demonstrated excellent internal information sharing practices. Many codified processes that require internally distributing artifacts to other departments, such as draft analytical products, network security data, and indications and warnings information. However, they consistently cited access to data from external organizations as a challenge. Organizational culture is the largest road block to success in this space, as mature technology solutions are available to overcome classification and need-to-know restrictions on information sharing.
- Information sharing for the organizations in industry and academia varied significantly. They generally failed to share data in a meaningful way, resulting in a reactive, patch-and-remediate cyber security posture. Similar to those in government, the most significant barrier to external information sharing in industry and academia was cultural; organizations remain reluctant to share "sensitive" network data and intelligence indicators with competitors. Conversely, entities that overcame this reluctance and routinely provide and consume data identified these practices as a major reason for their ability to stay ahead of cyber threats. Examples of data being shared include indicators of malicious activity, draft analytical reports, and contextual data surrounding malware and bad IP addresses.

- Initiatives sponsored by government and industry attempt to facilitate information sharing, but with limited success. Industry-sponsored information sharing generally is open only for select audiences and requires a financial commitment to join. Many of the government-sponsored arrangements tend to be redundant; they report the same data, but in different formats (one agency reports in .PDF, another in XML, another through RSS feeds), and with a range in timeliness. Information sharing relationships with the government also have the perception of being a “reporting” mechanism, which has dissuaded organizations from being more engaged.

Best Practice #1: Aligning functional and strategic cyber intelligence resources

High performing cyber intelligence programs employed a mix of functional and strategic analysts. For three organizations in the CITP, one government and two commercial, functional analysts were physically co-located with strategic analysts. Cyber intelligence is too big a topic for any one person to cover adequately. The nuances of technology, the intricacies of network defense, and the complexity of adversary intentions and capabilities makes it difficult for any one person to fully understand the cyber landscape. For this reason, successful cyber intelligence programs adopted a collaborative culture, so that experts can interact and share ideas.

Organizations that adopted this best practice were able to generate timely intelligence products, better communicate technical issues to senior leadership, and adjust data gathering tools to meet analysts’ needs more efficiently. The close interaction between functional and strategic analysts allowed them to more effectively understand complex technical details. This, in turn, provided analysts a better understanding of the threats and risks, benefitting their ability to communicate these concepts to leadership. The SEI Emerging Technology Center observed that organizations not employing this best practice incurred delays in reporting due to lags in collaboration either by email or phone calls. Other alternatives included paying to collaborate with third-party intelligence providers that offered technical expertise, or engaging in an online collaboration portal where participant expertise was difficult to verify.

Analysts also benefited from being co-located with their counterparts because it enabled them to seamlessly communicate data gathering requirements to the people that have access to the collection tools. Functional analysts typically had the ability to adjust data gathering tools or resources so that others could receive the data they needed. One organization in the CITP had strategic analysts sitting next to their functional counterparts responsible for a unique data-gathering tool. As the strategic analysts received new requirements, or wanted to pursue interesting data, they asked the functional analysts to collect this data, and received it almost instantly.

Best Practice #2: Information sharing in the financial sector

Financial sector organizations exhibited the strongest information sharing culture, processes, and mechanisms. Internally, they have formal communication channels between cyber security experts, analysts, and the various business divisions within their organizations. Analysts produce a range of intelligence products, each one designed to meet the needs of internal decision makers; from strategic summaries for executive leadership to organization-wide products educating the workforce on pertinent cyber threats. Strategic cyber intelligence analysts also work closely with functional analysts to understand the scope and nature of cyber threats, which better allows them to communicate risks and impacts to internal business operations.

Externally, these organizations are very active, benefitting from their involvement with the Financial Sector Information Sharing and Analysis Center (FS-ISAC). The financial services organizations in the CITP unanimously agreed that FS-ISAC indications and warnings directly enhance their network security. Additionally, the FS-ISAC facilitates numerous analytical exchanges, allowing participants to better understand the capabilities and techniques of cyber actors targeting the financial sector. It also fosters informal collaboration among members, despite the sector’s overarching competitive environment.

Environment




Figure 6 – Environment – CITP Baseline

Understanding internal and external environments allowed organizations to establish the scope of their cyber intelligence effort. The internal environment usually consisted of determining where the cyber intelligence program should exist and how to allocate resources. In some instances, aligning functional and strategic analysis efforts according to threat prioritization models aided resource allocation. The internal environment also included studying participant’s global cyber presence, what infrastructure was accessible through the Internet, and how to identify what data needs to be collected to maintain network situational awareness.

Externally, the environment involved knowing the entities capable of affecting organizations’ networks by focusing on system vulnerabilities, intrusion or network attack vectors, and the tactics, techniques, procedures, and tools used by relevant threat actors. It tended not to gauge the threat emanating from software supply chains, but in certain cases did track external factors affecting organizations’ different business units using open source monitoring. By investing the time and energy to define the environment, organizations significantly improved their data gathering efforts, resulting in more efficient and effective cyber intelligence programs.

Challenge: Understanding threats to the software supply chain

The unknown provenance of software complicated organizations’ ability to define the cyber environment.

Current state:

- Software development is a critical component of the networked world. Businesses, government, and individuals completely rely on software to perform daily tasks. Error-free and reliable software is a necessity for software found in commercial enterprises, industrial control systems, and military technology. When buying software, or having it coded for a specific purpose, these customers generally do not know the individuals performing the actual coding (much of software coding is out-sourced), the code’s reliability, or to what extent it has been error tested by developers. This puts customers from government, industry, and academia in the position of having to accept supply chain risks when contracting for software development, exposing them to potential security compromises that could cost proprietary information, R&D resources, business models, or future profits.
- Many organizations in the CITP did little to no vetting of software for security and counterintelligence purposes prior to acquisition. Although some stated they vet software vendors to ensure acquisition of the best available product on the market for their enterprise, they did not focus on understanding the software’s coding or any potential vulnerabilities associated with it.

Challenge: Determining where cyber intelligence belongs organizationally

Where the cyber intelligence function was organizationally situated affected its focus, performance, and effectiveness.

Current state:

- To fully leverage the benefits of a cyber intelligence program, it should be organizationally situated to inform leadership, strategic decision makers, network security personnel, and influence network and organizational policies. In practice, nearly every organization in the CITP housed its cyber intelligence function in a different location. Organizations in industry placed cyber intelligence personnel in a variety of offices including risk management, security operations, threat intelligence, or network management. Entities that aligned the cyber intelligence function more closely to security operations and network management relegated their analysts to more functional, reactive tasks supporting cyber security. Others that housed the intelligence function in areas such as risk management or threat intelligence fostered an environment where cyber intelligence fed strategic decision making, and had equal bearing to other strategic-level business units. The challenge inherent with these models was forming the relationship with network security, so that data was shared and intelligence products benefited from the technical expertise of the security staff.
- For government organizations, locating their cyber intelligence programs was less of a problem because they simply located them where financial resources existed to sustain the programs. Nevertheless, variances still were observed. In one case, financial resources dictated that a cyber intelligence program operate within a geographically focused non-cyber intelligence unit. In other instances, cyber intelligence was interspersed throughout multiple divisions of the same organization, augmenting non-cyber intelligence analysts, such as counterterrorism analysts, with a cyber component for their specific area of expertise.

Best Practice #1: Scoping the cyber environment to the organization’s mission

Cyber intelligence programs that incorporated the overarching goals of the organization into their cyber environment saw benefits in structuring data gathering requirements with the scope and focus of their analytical efforts. One organization in industry made cyber security a part of its business culture. This resulted in an extra emphasis being placed on the cyber intelligence component as a mechanism to identify potential threats that may impact this organization. Cyber intelligence analysts were kept apprised of new products being released and of other strategic business decisions so that they could be more productive in their analysis and focus their efforts on only the most relevant threats. This strategic insight was particularly valuable as it helped the analysts manage the collection and monitoring of more than 400 open source resources supporting approximately 1,500 products of interest to the organization. Because this entity’s leadership prioritized security across all products, cyber intelligence was ingrained with product development from the conceptual phase to public release.

Best Practice #2: Modeling threats to shape resource allocation

Cyber security resources remain limited. Organizations that attempted to broadly protect their data from all cyber threats tended to inefficiently invest these resources, making them slower to adapt to the changing trends and techniques of cyber threats. Entities with more effective cyber intelligence programs implemented a tiered threat model that helped determine the severity and priority of threats and potential targets of threat actors. These organizations were found to be more agile, able to appropriately and quickly respond to pertinent threats because they had been ranked and prioritized according to a specific threat model on a regular basis. In the financial sector, organizations used tailored threat matrixes. A simplified version of one of these matrixes is depicted below:



Figure 7 – Variation of Conventional Risk Matrix

Various threats can now be plotted on this matrix to provide an organization’s leadership, security staff, and risk managers a visual aid in understanding the severity of a particular threat.



Figure 8 – Tiered Cyber Threat Model

When deciding to invest in security, understanding the threat and its potential risk to the organization become strong influencers in the decision making process.

Data Gathering




Figure 9 – Data Gathering – CITP Baseline

To excel in performing cyber intelligence, analysts used their understanding of the cyber environment to influence how data was gathered. Data gathering consisted of identifying data sources, collecting the data, and aggregating it to support future analysis and to address basic cyber security issues. Effective data gathering contained both internal (e.g., netflow, logs, user demographics) and external sources (e.g., third-party intelligence providers, open source news, social media), and focused collection on the pertinent threats and strategic needs analysts identified while learning about their organization’s environment. Without clearly defining the environment, data gathering became disorganized. Entities collected too much unnecessary data and not enough substantive information for functional and strategic analysts to conduct any meaningful analysis on critical cyber threats.

Challenge: Data hoarding

Organizations knew they needed data for functional and strategic cyber intelligence analysis; however, the lack of planning and ineffective use of technology resulted in collecting and storing far more data than they could process.

Current state:

- Organizations were inundated with data. Some in the CITP collected so much data that they simply discarded it without looking at it. Other entities saved data, but did not use it effectively, as it continued to idly accumulate in their servers. Multiple organizations also collected relevant data from multiple sources, but failed to correlate it.

- When acquiring open source information, many organizations employed an ad-hoc approach to collecting data from threat intelligence services, colleagues, and a collection of personally selected open source websites and RSS feeds. Many of the entities in the CITP that subscribe to threat intelligence services found the relationship frustrating to manage, as they had to specifically and frequently tell the services what issues/keywords they wanted data on. As the organizations’ intelligence needs evolved, there was latency in communicating the new needs to the services and getting the corresponding intelligence information. The services inconsistent sharing of keyword-specific information and lack of timeliness in doing so were ongoing issues, forcing organizations to tackle open source data collection using any possible means. In many instances, the solution involved establishing traditional RSS feeds using applications such as Google Reader. The problem with this approach was having to manually sift through hundreds of articles for relevant data, which inhibited consistent and accurate information collection. Furthermore, this data was notoriously difficult to correlate with other data sources (network data, social media, chat rooms, geopolitical news sites) and complicated trend analysis or synthesis for actionable and predictive intelligence.

Challenge: Lack of standards for open source intelligence data taxes resources

The prevalence of non-integrated, non-standard content and delivery approaches from open source intelligence providers and subscription services burdened analysts, complicated correlation, and contributed to missed analytic opportunities.

Current state:

- Government, industry, and academic organizations in the CITP all reported challenges in efficiently collecting and integrating open source content into analytical products. Tradecraft required government analysts to meticulously record information about the source; a time consuming, manual process that was prone to errors. Some government organizations copied swaths of open source data and migrated it onto classified networks so analysts could safely access and analyze the data. This required duplicating data, which resulted in costly storage requirements and having to parse through unstructured open source data that was difficult to index and tag.
- Organizations in industry and academia also were inundated with data feeds that varied in format, making consumption and integration of information for further analysis difficult. Some entities in the CITP tackled this issue by developing initial concepts for standard formats and delivery of data in forms such as STIX, TAXII, and OpenIOC.

Best Practice #1: Repurposing search engine referral data

One organization was concerned with overseas competitors trying to duplicate a manufacturing process to replicate a product. The organization knew the production process, so they were able to gauge how far along competitors were in the process by utilizing Google Referral data. When the competitor was working on a particular stage of the manufacturing process, they used the Google search engine to learn as much as possible about that stage of the process. Since the manufacturing process is proprietary, and very few companies can afford the technology investment needed for production, websites owned by (or affiliated with) the participant typically came up in the Internet search engine's results. By aggregating and correlating the Google Referral data with the information they knew about their competitors, the organization was able to discern where in the manufacturing process its competitors were to anticipate what type of data was at the highest risk of being targeted for exfiltration.

Best Practice #2: Mind the gaps

Another entity wanted to ensure it was getting adequate coverage with its data gathering efforts. The organization created a data gathering plan, specifically detailing the types of information that it needed to collect in order to perform cyber intelligence analysis effectively. It then captured what data it was able to collect itself, and highlighted the remaining areas where it was missing coverage. Armed with these gaps, the organization provided specific data collection requirements to its third-party intelligence providers. When the intelligence providers sent these tailored products, the organization meta-tagged every one for indexing and solicited feedback from its consumers on the products. It utilized the consumer feedback to grade the products quality for timeliness and usefulness. Using the feedback and grades, the organization took this information to its intelligence providers to influence the type of reporting it would continue to receive. It also incorporated the feedback and grades into its yearly contract renewal discussions with the intelligence providers so the organization could make smart investments on whether to continue a relationship with a provider or seek other means to get pertinent intelligence. This process minimized the data gathering gaps, and ensured that their analysts didn't waste time on tasks they knew were covered by external providers. The organization also was able to wisely invest its data gathering resources.

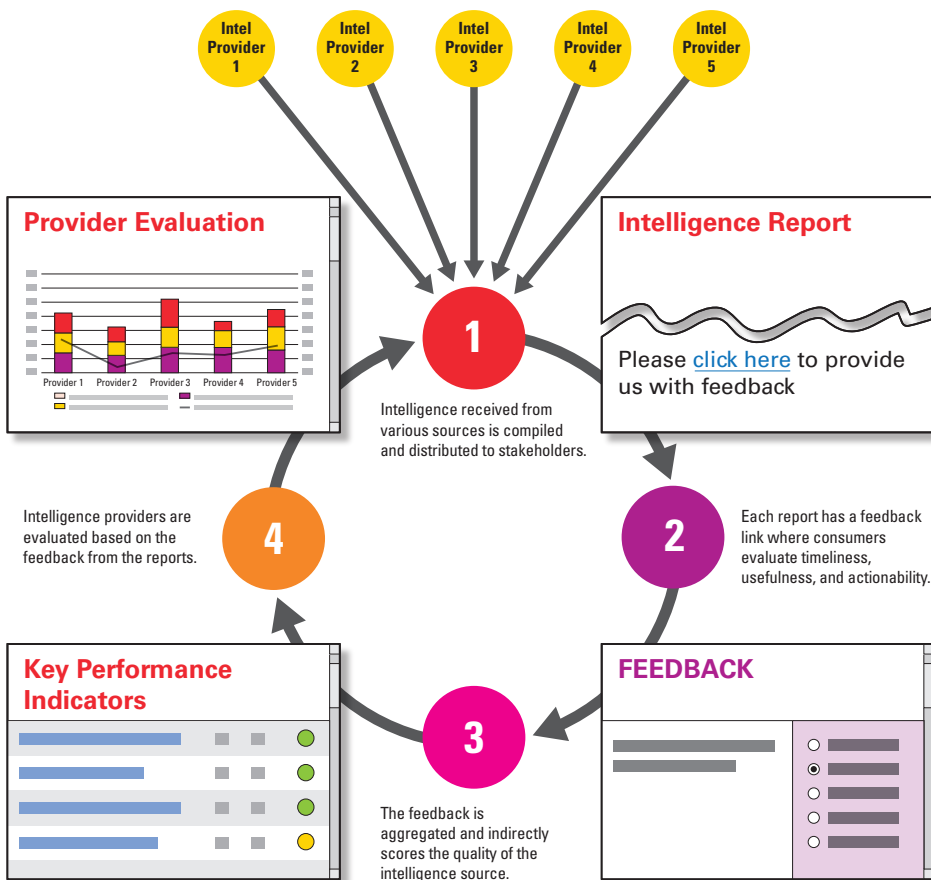


Figure 10 – Mind the Gaps Process

Functional Analysis




Figure 11 – Functional Analysis – CITP Baseline

Organizations produced quality functional analysis when a workflow existed to extract pertinent data from internal and external feeds, typically for the purpose of supporting cyber security by informing consumers of the technical complexities of the cyber threat. The process began with analysts taking technical information collected during data gathering and applying analytic tools and human resources to isolate potential threats. This information became intelligence as analysts validated its threat potential using personal and industry expertise, organizational threat priorities, present day situational awareness, and historical references. Analysts provided this intelligence verbally or through written means to internal strategic analysts and stakeholders responsible for cyber security or strategic decision making. However, the specific methods of collecting, analyzing, and communicating this analysis varied significantly among organizations in the CITP.

Challenge: Adopting a common cyber lexicon and tradecraft

The lack of a common lexicon and tradecraft was an impediment to the credibility of cyber threat data, which hampered analysis, attribution, and action.

Current state:

- **Lexicon:** During the CITP, organizations were asked to define key terms such as “cyber,” “intelligence,” “threat,” and “attack.” The definitions provided varied significantly within industries and across economic sectors. Even among more established cyber-related disciplines, such as cyber security, the vocabulary in use also carried different meanings depending on whether it was being offered by an entry-level analyst or manager.

- **Generic terminology:** Within government, industry, and academia, measures exist to prevent unwarranted and unlawful disclosures of identifiable information, such as IP addresses and company names. Organizations protected these details when sharing threat information by referring to them with generic terms, such as “IP 1” and “Company B.” While this ensured non-attribution, it inhibited other organizations from performing adequate functional, historical, or trend analysis to assess the threat’s impact to their enterprise. The process to request additional information on the IP or company also dissuaded analysts from engaging in these types of analysis because it was time consuming and usually resulted in no additional substantive information, especially within the government.
- **Tradecraft:** Many government organizations have adopted the intelligence community standard of consistently caveating threat analysis with estimative language and source validation based on the quality of the sources, reporting history, and independent verification of corroborating sources. Numerous individuals with varying levels of this skillset have transitioned to cyber intelligence roles in industry and academia, but the practice of assessing credibility remains largely absent. The numerous analytical products reviewed for the CITP either did not contain estimative or source validation language, or relied on the third-party intelligence service providing the information to do the necessary credibility assessment.

Challenge: Filtering critical cyber threats out of an abundance of data

Organizations struggled to accurately focus analytical efforts on critical threats because they could not adequately filter out data that once analyzed ended up being classified as low to moderate threats.

Current state:

- Many functional analysts were inundated with potential threat information that their job responsibilities required them to analyze, only to determine most of it posed a low to moderate threat to the organization. These time-consuming activities diminished their accuracy in identifying critical threats and devoting the necessary resources to analyzing them. To reduce the burden of personnel doing this work, some entities outsourced it to third-party intelligence service provider.
- In government, as well as a small set of organizations in industry, robust policy restrictions filtered out low level threats. Restricting the ability to open an executable, limiting the use of commonly exploited software, prohibiting USB storage devices, and impeding access to websites associated with scams and malware made it very difficult for low sophistication hackers (recreational, or “script kiddies”) to affect these networks. This freed up resources to focus on more sophisticated threats.

Best Practice #1: Comprehensive workflow to identify cyber threats and inform customers

Based on established standard operating procedure policies, the cyber-focused intelligence operations entity of an information technology organization described how it used a comprehensive functional analysis workflow to identify legitimate cyber threats and inform customers of these threats in a timely fashion. Data initially is identified as a potential threat when automated tools pull information from the organization’s network and security sensors per a prioritization model that incorporates data gathering needs, analyst expertise, and the parameters of an internally developed threat scoring system. Once the data reaches a specific threat threshold, it is placed in an email folder. A senior analyst responsible for monitoring this folder then reviews the potential threat data and assigns it to another analyst.

The assigned analyst uses multiple resources, including previous intelligence reporting, additional data feeds, personal expertise, and open source research to address the threat’s technical and cyber security components in a formal security alert. Per a predetermined timeline, the analyst works to produce an initial security alert with an 80 percent solution that internal and external customers can use to protect their enterprise against the threat. He or she has 90 minutes to produce an alert on a critical threat, six hours for a high threat, and 24 hours for a low threat. After the initial alert is disseminated, it becomes a living document placed in a common email folder for all analysts within the cyber-focused intelligence operations entity to edit and update with the goal of reaching the 100 percent solution. Each updated version of the security alert is automatically sent to customers via email, showing the

entire history of how the alert has changed over time. The security alert also is incorporated or serves as the basis for other formal products produced by the intelligence operations entity.

Best Practice #2: Producing scripts to automate the filtration of known threat data

Functional analysts at a government organization stated that they leveraged relevant environmental factors and intelligence requirements provided by strategic analysts to write scripts for automating the distribution of network activity into threat categories that functional analysts could choose to access according to threat criticality. Over the years, they have written so many of these threat scripts that many low to moderate and routine threats are automatically filtered out of the network activity. Eliminating much of this “noise” provided the functional analysts a smaller data set from which to investigate potential new threats. This resulted in more timely and accurate functional analysis being provided to strategic analysts, decision makers, and consumers.

Strategic Analysis




Figure 12 – Strategic Analysis – CITP Baseline

Strategic analysis added perspective, context, and depth to functional analysis, and incorporated modus operandi and trends to provide the “who” and “why” of cyber threats. It was ultimately rooted in technical data, but incorporated information outside traditional technical feeds—including internal resources such as physical security, business intelligence, and insider threat, and external feeds covering global cyber threat trends, geopolitical issues, and social networking. The resulting strategic analysis populated threat actor profiles, provided global situational awareness, and informed decision makers of the strategic implications cyber threats posed to organizations, industries, economies, and countries. Performing such analysis required a unique mix of technical and intelligence skills that organizations continue to debate on how to acquire, nurture, and lead.

Challenge: No industry standard for cyber intelligence education and training

The cyber intelligence workforce was a heterogeneous mix of technical experts and non-technical intelligence analysts, neither completely familiar with the nuances and complexity of the other.

Current state:

- Every organization in the CITP employed some combination of trying to teach technical experts intelligence tradecraft or to teach all-source intelligence analysts fundamentals of network technology. Across government and industry, there was no clear, definitive standard for the skills and competencies required for cyber intelligence professionals. The executive director for technology risk of an organization stated that if such a standard were adopted, getting his staff trained and certified would be a top priority.

- The organizations devoting the most resources to bridging the gap between analyst and security professional resided in the government. Depending on the agency, government cyber intelligence analysts spent anywhere from six weeks to 18 months being immersed in training in intelligence tradecraft, analyst tools, networking fundamentals, courses on legal and organizational policies, operational implementation of intelligence, and effective writing. Not coincidentally, most of the successful organizations in industry have built their success on hiring former government and military intelligence professionals.
- Many of the organizations in the CITP claimed they prefer to train an analyst to understand the technical aspects of cyber security than to train an information technology (IT) practitioner how to do intelligence analysis. It should be noted that despite voicing this opinion, the actual composition of their analytic staff all had IT backgrounds. When asked what an ideal candidate looks like, proficiency in the cyber environment was the top requirement.

Challenge: Adapting traditional intelligence methodologies to the cyber landscape

Because technology changes so quickly, the process of producing cyber intelligence analysis had to be dynamic enough to capture rapidly evolving tools, capabilities, and sophistication of adversaries.

Current state:

- Many of the intelligence methodologies observed in government organizations were developed in an era when intelligence analysts focused on methodically counting tanks, missiles, and airplanes held by hostile countries to predict, over time, what their leaders planned to do with these weapons and transportation vehicles. Applying these same processes, workflows, and tradecraft to the cyber domain was not always feasible. By the time a strategic-level product on an emerging threat made it through the publication process, it was already out of date.
- Several process shortfalls compounded the problem across government. One organization in the CITP mentioned that strategic cyber intelligence products followed the same process as more traditional intelligence products, and took months to publish. A big part of the bottleneck was that since cyber intelligence is a relatively new discipline, no robust senior corps of cyber intelligence analysts exists in government. Thus, the senior analysts required to review all analytical products did not know as much about the subject matter as the junior analyst authors. This delayed the product being published, as the reviewers pushed complex, controversial, or unfamiliar cyber intelligence topics to the end of the queue to make way for ones on topics the reviewers were more comfortable with subject-wise.

Best Practice #1: Know your enemy

The highest performing cyber intelligence programs have built profiles of the top cyber threats, and tracked these actors as their tactics and tradecraft evolved over time to adequately prepare the organizations network defenses. One government organization in the CITP created profiles of adversaries that included TTPs, malware used, tools, C2 infrastructure, names used, spear-phishing tactics, and common targets. Compiling this data helped them to attribute new activity, and track the evolution of their adversaries. An organization from industry extended this type of a profile to include the motivation of hackers and how they made their money.

Separately, an industry entity excelled in this area by mapping threats to potential sponsoring organizations. Through open source research on the sponsoring organizations, the industry entity was able to narrow down the types of data likely to be targeted, and work with network security experts to create diversions, honey pots, and employ other defensive measure to try and get out in front of the threats. As the motivations of the threats changed, this organization adapted its threat profile to identify new types of at-risk data. Furthermore, when its different business units expanded their work into overseas markets, this organization was able to anticipate the threats this activity would trigger, and incorporated these risks into the business unit's overarching strategy.

Best Practice #2: Global situational awareness

Cyber intelligence that looked beyond the organization's network perimeter provided strategic insights that fed predictive analysis. One organization in the CITP used a tool that provided visibility into the IP ranges of commercial partners. That way, when a vendor was compromised, the organization could take preventive measures and ensure that the malware didn't spread into its networks, or that an attacker was not able to move laterally from the supplier's network into its network. Another entity utilized geopolitical analysts to add context to cyber intelligence analysis. This organization has an international supply chain, so the collaboration between the cyber and geopolitical analysts often yielded insights that better prepared the entity's leadership for traveling overseas.

Another organization in the CITP looked at what hackers were doing to other entities both inside its economic sector and around the world in areas where it has major business interests. This entity examined these external issues and attempted to determine if the issues affected it in the near or long term. Examples included incidents at domestic services industries and international commerce entities. The organization then produced an "external breaches" slide for leadership that depicted these issues. Analysts selected many of the events being covered because the organization has or might have business relationships with them; therefore, the threat could adversely affect this entity.

Decision Maker Reporting and Feedback




Figure 13 – Decision Maker Reporting and Feedback – CITP Baseline

When organizations overcame the challenge of communicating strategic analysis to leadership, their cyber intelligence programs became an integral part of strategic planning. Decision makers digested the information to provide feedback for shaping analytical efforts and to adjust the direction of the overarching organization. The CITP’s analytic framework reflects this approach, showing how these types of decision makers could continue the cyber intelligence process when analysts effectively demonstrated return on investment and carved out necessary communication channels.

Challenge: Communicating “cyber” to leadership

Decision makers removed from the cyber environment generally lacked technical backgrounds, and functional analysts generally lacked experience writing for non-technical audiences.

Current state:

- The technical complexities associated with cyber security were difficult for many organizational leaders to appreciate. For the majority of organizations in the CITP, leadership did not have a need (or desire) to understand the technical details of what was happening to their networks; they just wanted to know why it was important to the organization. At one government organization, a cyber intelligence analyst noted that because cyber security and cyber intelligence were relatively new areas, there was a dearth of senior leadership in the cyber intelligence field. This lack of a senior corps of cyber-savvy analysts meant there was a lack of mentorship to junior analysts, which only perpetuated the problem of poor communication between non-technical leadership and the cyber community.

Challenge: Difficulty capturing return on investment

Organizations typically used return on investment (ROI) calculations to justify the costs associated with business practices or infrastructure requirements. In cyber intelligence, coming up with ROI remains difficult.

Current state:

- Government organizations typically used performance measures that focused on quantity (e.g. number of reports generated), but not necessarily on quality or impact of intelligence. Analysts were encouraged to get feedback, but valuable feedback on intelligence products was limited and anecdotal. In industry, performance measures, particularly those that could demonstrate return on investment, were critically needed. Seasoned practitioners became well aware of the value proposition and the potential costs of not engaging in cyber intelligence, but defining real metrics that could be used to justify resource needs and ensure corporate support was very difficult. Some organizations had the ability to easily assign dollar values to protected assets; others used the cost of recovery from compromises. For many organizations, the measure of the value of cyber intelligence remained elusive.

Best Practice #1: Failure analysis

One organization in the CITP that was struggling to produce ROI metrics took the approach of looking at past events and capturing what the negative or potential effects of adversarial access to its data could have been for it. To do this, the organization looked at what information was publicly available from partners in its supply chain, and then went and looked at what data was targeted by hackers from its networks. The team was able to surmise what competitors knew about their events based on this analysis, and estimated what competitors could have done with this information had they wanted to disrupt the event. In some cases, when the organization discovered that data was being taken, it spent time and money to create diversions to confuse competitors. The cyber intelligence analysts were able to capture the costs associated with these activities, and essentially used them as “negative ROI” figures for leadership. This failure analysis sent the message that had more resources been used to protect this data and track the competition’s interest in it, the organization could have saved the money spent on creating the diversions.

Best Practice #2: Carving channels for communication

A robust reporting approach considered content appropriate and necessary for the audience and relevant to the organization, with thought for frequency, timing, and delivery. An organization in the CITP wanted to maximize the benefit of their cyber intelligence analysis by using it to support cyber security and senior leadership. To accomplish this, the company identified groups of decision makers that included senior leadership, risk managers, individual business units, and security staff. It then established communication channels via email distribution lists to provide these decision makers tailored analytical products, such as monthly cyber security tips newsletters and weekly senior leadership briefings. This prevented irrelevant information from appearing in leadership’s email folders, and created a culture where managers knew that if there was an email from the cyber intelligence program, it contained timely and pertinent information worthy of consumption.

The organization also utilized this effort to solicit feedback by including a link at the bottom of each product for recipients to comment on the utility of the intelligence. Although feedback was initially low, the mechanism was in place to receive comments and new information collection requirements.

Conclusion

This report discussed the SEI Emerging Technology Center's Cyber Intelligence Tradecraft Project (CITP). The purpose of the CITP was to study the state of the practice in cyber intelligence and advance the capabilities of organizations performing this work by elaborating on best practices and prototyping solutions to common challenges. It accomplished this by using the methodologies, technologies, processes, and training forming the cyber intelligence programs of 30 organizations to develop a baseline that the SEI Emerging Technology Center benchmarked against its cyber intelligence analytic framework.

The CITP's key findings indicated that organizations used a diverse array of approaches to perform cyber intelligence. They did not adhere to any universal standard for program development, data gathering, or analyst training. Instead, pockets of excellence enabled certain organizations in government, industry, and academia to successfully perform cyber intelligence by sharing information through venues like the Financial Services-ISAC and communicating return on investment using post-event failure analysis. The CITP key findings also show how organizations could filter data to identify threats by aligning data gathering with input from threat prioritization models and predict threats via the repurposing of search engine referral data. Overall, this report finds that any organization can excel at performing cyber intelligence when it balances the need to protect the network perimeter with the need to look beyond it for strategic insights.

Copyright 2013 Carnegie Mellon University

This material is based upon work funded and supported by ODNI under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center sponsored by the United States Department of Defense.

Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of ODNI or the United States Department of Defense.

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

This material has been approved for public release and unlimited distribution except as restricted below.

Internal use:* Permission to reproduce this material and to prepare derivative works from this material for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use:* This material may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other external and/or commercial use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

* These restrictions do not apply to U.S. government entities.

DM-0000620