January 18, 2013

SEI Innovation Center Report: Cyber Intelligence Tradecraft Project

Summary of Key Findings

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

Software Engineering Institute | Carnegie Mellon

SEI Innovation Center Report: Cyber Intelligence Tradecraft Project Summary of Key Findings

SEI Innovation Center Report: Cyber Intelligence Tradecraft Project

Summary of Key Findings

Executive Summary

In February 2012, the Office of the Director of National Intelligence sponsored the Innovation Center at Carnegie Mellon University's Software Engineering Institute (SEI) to assess the state of cyber intelligence in the public and private sectors and identify common challenges to address through best practices and prototyping efforts. The assessment began by partnering with six U.S. government and 20 private organizations from multiple economic sectors to baseline the methodologies, technologies, processes, and training forming their cyber intelligence functions. The baseline data subsequently became part of a benchmarking process for evaluating participants in an analytic framework the SEI Innovation Center developed for performing cyber intelligence. The framework incorporates 35 factors in five functions: define the environment, data gathering, functional analysis, strategic analysis, and decision making/reporting and feedback.

Overall, organizations succeed in performing cyber intelligence by anticipating and prioritizing how to respond to potential cyber threats through routine cyber risk evaluations, maintaining threat actor awareness, and proactively reviewing and rating risk exposure based on threat. These organizations also balance protecting the network perimeter with looking beyond it for strategic insights and predictive analytics. However, general industry standards do not exist for organizations to use as a reference to enhance capabilities. This includes how to establish a cyber intelligence function within an organization, develop education and training requirements, and create and manage performance measures and return on investment figures.

This report highlights the key findings identified during the interview working sessions and analysis of the data up to this point. The intent of this report is to provide useful and timely feedback both to the participants and to the sponsor. While this is a representative cross section of information, there are still many insights that can be gleaned from continued analysis of this data and future research. This document is a summary of the key findings discovered thus far, and does not contain the supporting data set and accompanying analytic process.

Introduction

Cyber intelligence has grown from the halls of the U.S. government into a burgeoning industry providing tools and services in the private sector. As more and more public and private entities become involved with cyber intelligence, varying methodologies, technologies, processes, and training come to ill-define its operating environment. Recognizing the need to understand and improve this domain, the Office of the Director of National Intelligence (ODNI) sponsored the SEI Innovation Center to assess the state of cyber intelligence in the public and private sectors and identify common challenges to address through best practices and prototyping efforts. The following report describes the baseline and benchmarking approach and the resulting key findings that will shape how the SEI Innovation Center and its partners confront the domain's common challenges with technological, methodological, and educational solutions in the next phase of the Cyber Intelligence Tradecraft Project (CITP).

Participants

The SEI Innovation Center partnered with six U.S. government and 20 private organizations from multiple economic sectors, including defense contracting, energy, financial services, healthcare, higher education, information technology, intelligence providers, legal, non-profits, and retail. The participant list represents government organizations with a cyber intelligence mission, Fortune 100 companies, and a collection of niche companies. The individuals representing these organizations also come with diverse backgrounds and professional experiences. They include military, intelligence, and information security backgrounds with job titles such as chief information security officer (CISO), vice president of threat management, director of information security, information architect, intelligence analyst, and network/security analyst.

Cyber Intelligence Definition and Analytical Framework

With a wide range of participants conducting varying levels of cyber intelligence, the SEI Innovation Center instituted a common definition of cyber intelligence to use throughout the study. The definition attempts to incorporate U.S. government and private sector descriptions:

Cyber intelligence is the acquisition and analysis of information to identify, track, and predict cyber capabilities, intentions, and activities to offer courses of action that enhance decision making.

The SEI Innovation Center also developed an analytical framework to guide the study's baseline and benchmarking efforts. The framework originates from a version of the traditional intelligence cycle because this cycle represents a repeatable process. However, the utility of this model becomes limited when applied to cyber intelligence. The traditional intelligence cycle[1] is depicted as a linear flow that does not emphasize the inter-related nature of processes within the cycle, or how the intelligence cycle connects to other highly relevant functions. In the cyber intelligence environment, technology has a stronger role, which is reflected in the analytical framework by integrating technology into every component and emphasizing the analytical activities by breaking them into two components.

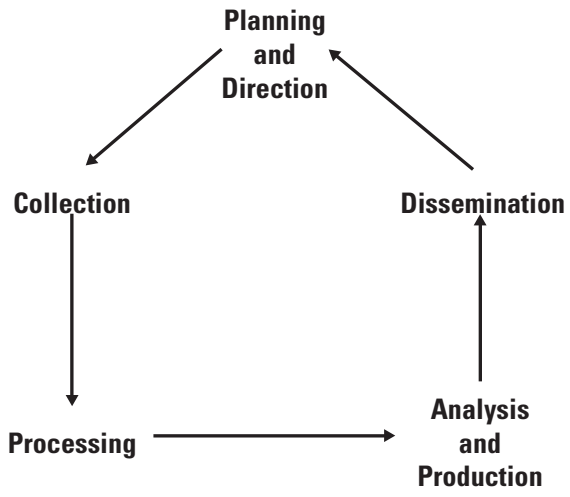

The SEI Innovation Center emphasized the unique cyber intelligence characteristics by creating a five-function approach that more accurately captures the inter-dependencies and outside influences on the cyber intelligence process. The first function, defining the environment, establishes the scope of the cyber intelligence effort, which influences how the second function, data gathering, explores information sources, collects the information, and aggregates it to support the third and fourth functions. Functional analysis, the third function, uses data gathering to primarily answer the "what" and "how" of cyber threats. The fourth function, strategic analysis, adds big picture implications to functional analysis, such as the "who" and "why" of cyber threats. The analysis is then disseminated to a fifth function, the decision maker. This function includes the reporting mechanisms and feedback channels that allow decision makers to influence strategic analysis, and make decisions that shape the organization's environment, causing the cycle to repeat. However, the framework does indicate that involving the decision makers is not always necessary, as strategic analysis can influence understanding of the cyber environment on its own.

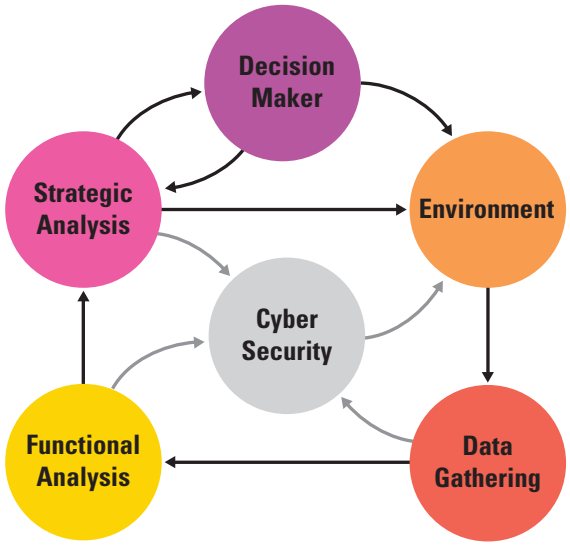

It is important to note that the analytical framework does not exist to address an organization's cyber security practices. Cyber intelligence is a critical component of cyber security and the two functions are inter-related; however, this study focuses on the cyber intelligence process. Throughout this process, data is shared with cyber security or network defense professionals, but as the illustration demonstrates, the cyber intelligence process operates independently and does not necessarily need to support a cyber security/network defense mission. Cyber intelligence can support a variety of business and government functions including strategic business decisions, national policy, international negotiations, military applications, or strategic communications.

[1] 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/analytic_culture_report.pdf. Last accessed 4 June, 2012.

Baseline and Benchmarking Approach

To capture the necessary methodologies, technologies, processes, and training that form the participants' cyber intelligence functions, the SEI Innovation Center employed an iterative process to develop a discussion guide that both parties talked through during a baseline session. The baseline session was usually conducted with a three-person cross-functional team with intelligence and software engineering backgrounds. The teams conducted three- to -four hour face-to-face interviews and explored the organizations' capabilities and activities in cyber intelligence. While the team used the discussion guide to ensure complete coverage of all topics, the session was intentionally not scripted to remove any biases the investigating team may have had and to provide an open forum for new ideas.

Upon completion of all the baseline sessions, the SEI Innovation Center reviewed, compiled, and organized the raw notes into a unified set for ease of analysis. From these unified notes, one reviewer evaluated the baseline results, selected the most appropriate assessment ranking, and provided the rationale drawn from the notes. A second team member assisted both in creating the benchmark data and serving as a peer reviewer. Once the peer review was complete, the entire team conducted a group review which ensured overall consistency for the evaluation while adding to the team knowledge base for the researchers who did not attend the actual baseline session.

The assessment criteria emerged and evolved from the baseline sessions themselves and the lessons the team learned through the process. The team developed a five-point ordinal scale to evaluate the criteria; the values used were "++", "+", "0", "-", "--" and "N/A" for "not applicable." The team chose an ordinal scale because it offered more flexibility when being applied to the tremendous variety of respondents and sizes of the participating organizations.

From the final combined data set, the team developed the baseline, then benchmarked each of the contributing data sets from each participating organization against that data. As with all ordinal data sets, numerical or interval analysis is inappropriate and consequently limited some analytical approaches. However, the data did yield interesting results and trends, and aided the identification of common challenge areas across the public and private sectors.

Each section below begins with a graphical depiction of the benchmark data shown in a stack column chart. Along the bottom of the X-axis are the benchmark criteria, and the Y-axis captures the number of respondents. Each graph shows the number of participants that were rated "0" or higher as being above the X-axis, and respondents that were rated below "0" are drawn below the X-axis. Overlaid on top of this data is the median curve.

Overall Findings

State of Cyber Intelligence

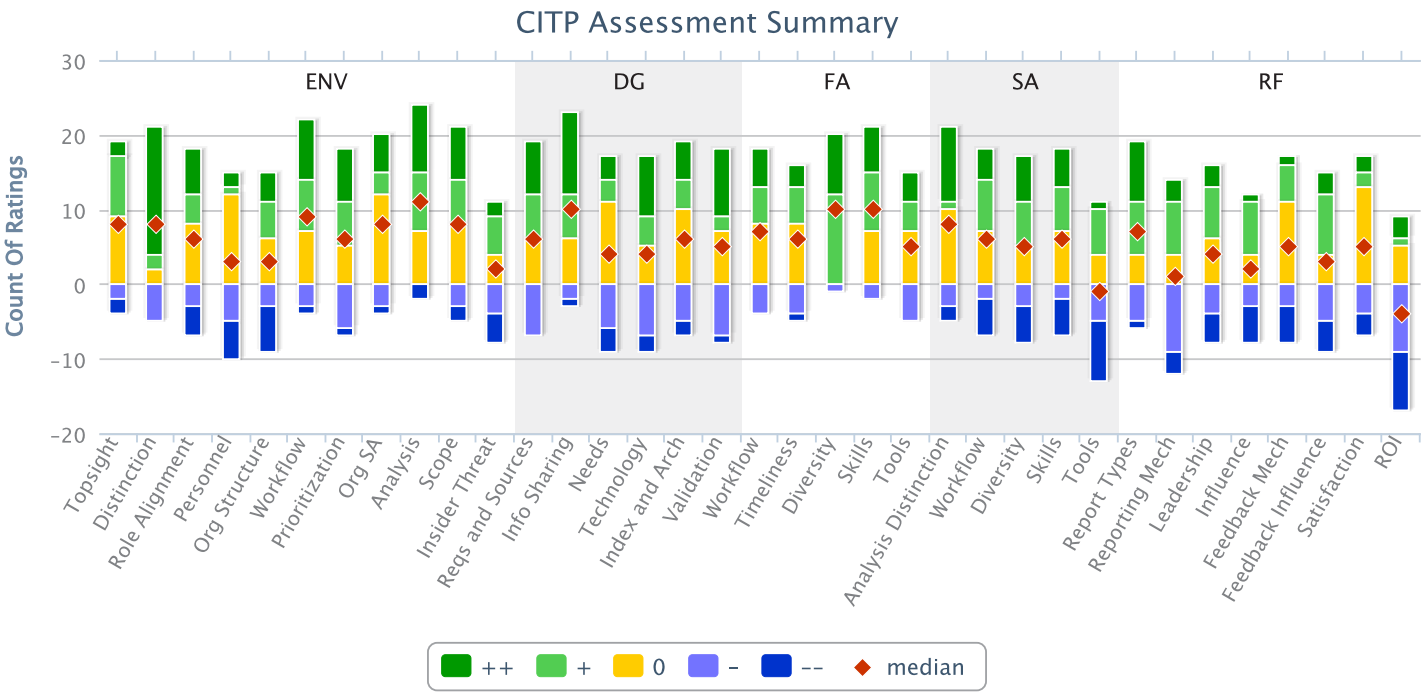

Overall Findings, continued

Most participants identify cyber intelligence and cyber security as two distinct and capable work functions that interact when necessary to best support the organization. In conducting cyber intelligence, participants generally performed well when trying to understand their internal and external cyber environment, gather threat data, and analyze technical threats ranging from malware to email phishing. These efforts effectively supported the organization's cyber security function, but did not adequately address strategic concepts, and thus failed to support decision making beyond cyber security. This exhibits the endemic problem among participants that their functional analysts cannot communicate to non-technical audiences. It also demonstrates how organizations struggle with sharing information within their own industry, with the exception of the financial services industry.

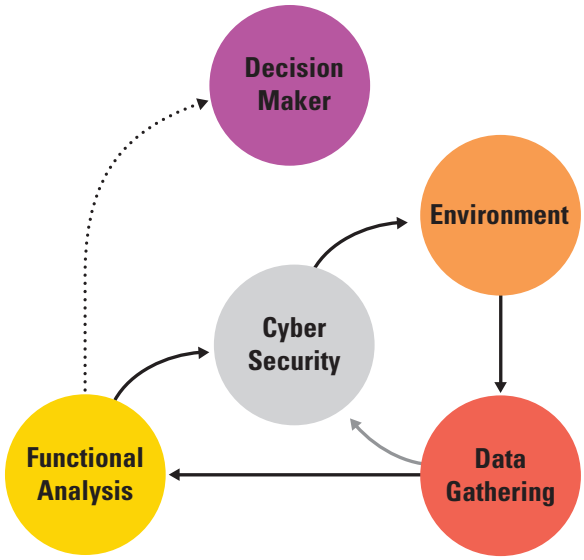

Challenge: Applying a strategic lens to cyber intelligence analysis

Despite having a wealth of data available, organizations are challenged in moving beyond functional analysis of low-level network data to a strategic analysis of threats and threat indicators.

Current state:

- Most organizations showed difficulty incorporating strategic intelligence analysis into existing security-focused processes. Correspondingly, organizations with poor or no strategic analysis functions had difficulty 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 to strategic analysis, from not having data indexed and stored for analysis, to not having resources to perform trend analysis and look at individual network events in a more strategic context. In many industries, organizational leadership does not fully understand or appreciate the complexities of cyber security. As a result, getting leadership to invest in tools or additional resources is hampered by the inability to link the benefits of a strategic cyber intelligence process to the organization's goals and objectives.
- Most organizations have mature processes that incorporate functional analysis, but only as a means to support cyber security (see below). The challenge in this model is communicating the importance and relevance of technical issues to decision makers in compelling terms that they will understand. Although security benefits from the functional analysis, the addition of strategic analysis was the most effective means of bridging the communication gap between cyber security and non-technical leadership.

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

The highest performing cyber intelligence processes actively shared (not just consumed) data in formal and informal information sharing arrangements.

Current state:

- Government participants demonstrated very good internal information sharing practices. Many of these organizations had codified processes that require sharing draft analysis products, network security data, or indications and warnings data to other departments within the organization. However, access to data from external government organizations was consistently cited as a challenge. Organizational culture is the largest roadblock to success in this space, as mature technology solutions are available to overcome classification and need-to-know restrictions on information sharing.
- The financial sector participants uniformly demonstrated robust information sharing arrangements both internally and externally (see Best Practice #1 below). Outside of this sector, information sharing for commercial and private sector participants varied widely. Organizations that generally scored below average had difficulty sharing data in a meaningful way, resulting in a reactive, patch-and-remediate cyber security posture. Like the government, the largest barrier to external information sharing is cultural; organizations expressed reluctance to share "sensitive" network data and intelligence indicators with competitors. Additionally, the more data that is shared, the better. Organizations that shared indicators of malicious activity, draft analytical reports, and contextual data surrounding malware or bad IPs cited the information sharing arrangement as a key component to helping them stay ahead of relevant cyber threats.

- The government and several private companies have attempted to facilitate information sharing arrangements with the private sector, with limited success. The private sector opportunities are limited, and have a financial cost to enter into the arrangement. Most of the government-sponsored arrangements tend to be redundant; they report the same data, report that same data 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: Information sharing in the financial sector

Of the organizations baselined, the financial sector participants generally exhibited the strongest information sharing culture, processes, and mechanisms. Internally, financial sector participants exhibited formal communication channels between cyber security experts, analysts, and the various business divisions within the company. Analysts produce a range of intelligence products, each one designed to meet the needs of internal stakeholders, from strategic summaries for executive leadership to company-wide products that educate the workforce on pertinent cyber threats. Strategic cyber intelligence analysts work closely with functional analysts to understand the scope and nature of cyber threats, which better allows them to communicate risks and impacts to business operations throughout the company.

Externally, the financial sector participants are very active in information sharing arrangements and have measurably benefitted from participation in the Financial Sector Information Sharing and Analysis Center (FS-ISAC). The financial sector organizations baselined in this study unanimously agreed that indicators and warnings issued by the FS-ISAC directly enhanced network security. Additionally, the FS-ISAC community facilitates analytical exchanges, allowing the organizations to better understand the capabilities and techniques of cyber actors targeting the financial sector. As an added benefit, the FS-ISAC has fostered an informal information sharing relationship with participating members, allowing security experts and analysts to informally share and collaborate on ideas. Of note, this robust sharing environment exists with little to no adverse effects despite the healthy competition in the marketplace between the information sharing partners.

Best Practice #2: Aligning functional and strategic cyber intelligence resources

High performing cyber intelligence programs use a mix of functional analysis and strategic analysis. In three organizations that were baselined, 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 understanding adversary intentions and capabilities makes it infeasible for any one person to fully understand the landscape. For this reason, successful cyber intelligence programs adopt 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 strategic analysts and functional experts allowed the analysts to more effectively understand complex technical details. This, in turn, allowed the analysts a better understanding of the threats and risks, benefitting their ability to communicate these concepts to leadership. The SEI Innovation 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 providers that offered technical expertise, or participating in an online collaboration portal where the expertise of the participants was difficult to verify.

An additional benefit of this best practice was the ability for analysts to communicate data gathering requirements to the people that have access to the collection tools. The functional experts typically had the ability to adjust data gathering tools or resources so that analysts could receive the data that they needed. One organization surveyed had analysts sitting next to functional experts that were responsible for a data gathering tool. As the analyst received new requirements, or wanted to pursue interesting data, the analyst could ask the functional expert to collect different data and receive the information almost instantly.

Challenges and Best Practices by Function of the Analytical Framework: Define the Environment

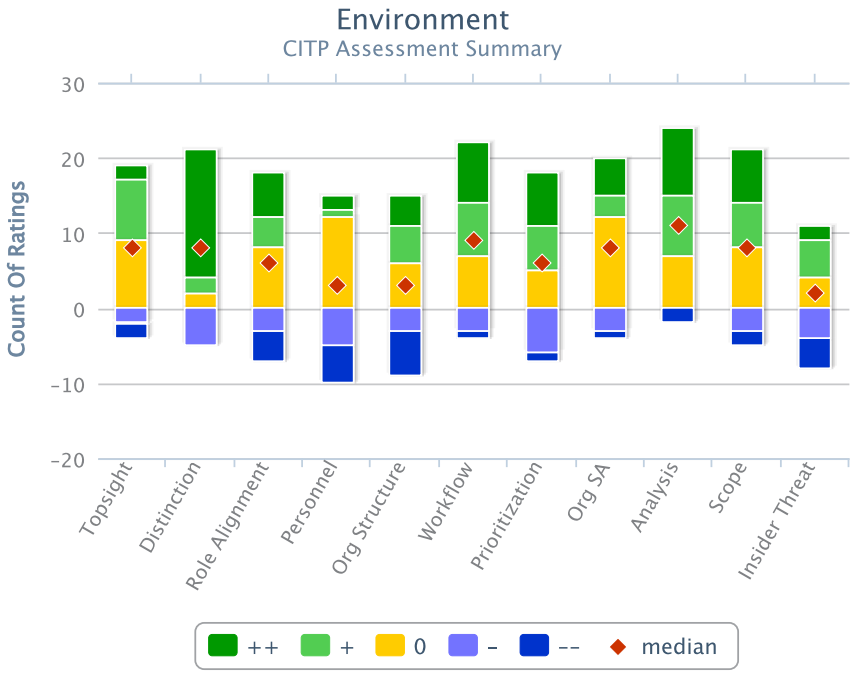

Establishing the scope of the cyber intelligence effort occurs by understanding the organization's internal and external environment. Internally, this begins with the organization correctly determining where their cyber intelligence function should exist, which currently is a challenge for most participants. However, some overcome this issue by prioritizing threats to best align resources and ensuring to incorporate a good mix of functional and strategic analysts within the cyber intelligence function. They then utilize this mix to study the organization's global cyber presence, what infrastructure is accessible through the Internet, and to identify what data needs to be collected to maintain network situational awareness. Externally, understanding the organization's environment means examining what entities are capable of affecting the network. Participants generally limit these entities to threat actors, vulnerabilities, intrusion or network attack vectors, and the tactics, techniques, procedures, and tools used by the identified threat actors. They do not tend to gauge the threat emanating from their software supply chain, but in some instances do track external factors affecting the different services or products they provide using open source monitoring. By investing the time and energy to define their environment, organizations significantly improve their data gathering efforts, resulting in a more efficient and effective cyber intelligence process.

Challenge: Understanding threats to the software supply chain

The unknown provenance of software complicates the ability to define the cyber environment.

Current state:

- Software development is a critical component of the networked world. Businesses, government, and individuals are completely reliant on software to perform day-to-day jobs. From software in commercial enterprises, industrial control systems, and military technology, to operating systems found on personal computers, error-free and reliable software is a necessity. When buying software, or having it coded for a specific purpose, the customer is generally unaware of who is performing the actual coding (much of software coding is out-sourced), how reliable the code is, or to what extent it has been error tested by the developers. This puts commercial and government customers in the position of having to accept supply chain risks when contracting for software development, exposing the customer to potential security compromises that could cost them proprietary information, R&D resources, business models, and future profits.
- Many commercial and government entities do little to no vetting of software for security and counterintelligence purposes prior to acquisition. Although some organizations in the CITP study stated they vet software vendors to ensure acquisition of the best available product on the market for their enterprise, they do 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 is organizationally situated can affect its focus, performance, and effectiveness.

Current state:

- To fully leverage the benefits of a cyber intelligence program, it should be organizationally situated such that it can access leadership to inform strategic decision making, and communicate to network security offices to collaborate on products and inform security policy. In practice, nearly every organization baselined housed its cyber intelligence function in a different spot. From risk management, to security operations, to threat intelligence, to network management, private sector organizations inconsistently located cyber intelligence personnel, which likely contributed to the varied performance. Organizations 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. Organizations that house 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 is forming the relationship with network security, so that data is shared and intelligence products are informed by the technical expertise of the security staff.
- In the government, this was less of a problem although variances were observed. For the most part, the cyber intelligence component was located where financial resources were able to sustain it. In one case, that was in a geographically focused unit. In other cases, cyber intelligence was interspersed throughout multiple divisions of the same organization, augmenting other analysts and providing a cyber component to their specific area of expertise.

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

Cyber intelligence programs that incorporate the overarching goals of the organization into their cyber environment saw benefits in structuring data gathering requirements and the scope and focus of their analytical efforts. One commercial organization has incorporated cyber security into its business culture. This resulted in an extra emphasis placed on the cyber intelligence component as being a mechanism to identify potential threats that may impact this organization. Cyber intelligence analysts were kept appraised 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 to the organization. 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 that are of interest to the organization. Because the organization's leadership prioritized security across all products, cyber intelligence is ingrained with product development from the conceptual phase to its public release.

Best Practice #2: Modeling threats to shape resource allocation

Cyber security resources are limited. Organizations that attempt to broadly protect their data from all cyber threats tend to inefficiently invest these resources, and are slower to adapt to the changing trends and techniques of cyber threats. Organizations with more effective cyber intelligence programs implement a tiered threat model that helps determine the severity and priority of threats and potential targets of threat actors. Organizations were found to be more agile and could respond more appropriately to threats when the threats were ranked and prioritized on a regular basis. In the financial sector, organizations employ several variations of a threat matrix. A simplified version of this is represented below:

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

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

Data Gathering

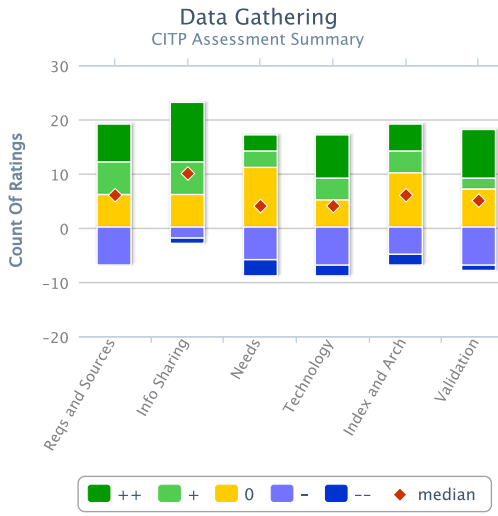

Data gathering is influenced by the organization's unique cyber environment and consists of identifying data sources, collecting the data, and aggregating it to support future analysis or to address basic cyber security issues. Effective data gathering consists of both internal (e.g., netflow, logs, user demographics) and external sources (e.g., third-party intelligence providers, open source news, social media), and focuses collection on pertinent threats and strategic needs of the organization as identified in the "define the environment" process. Without clearly defining the environment, data collection becomes disorganized, and organizations are apt to collect more data than can be usefully processed, or not have adequate data to conduct meaningful analysis on critical cyber threats.

Challenge: Data hoarding

Organizations know they need data for functional and strategic cyber intelligence analysis; however, the lack of planning and ineffective use of technology results in collecting and storing far more data than they can currently process.

Current state:

- Organizations are inundated with data. Some organizations baselined for this study collected so much data that they simply discard it without looking at it. Other organizations save data, but do not use it effectively and it continues to idly accumulate in their servers. Other organizations collected relevant data from multiple sources, but failed to correlate it.
- When acquiring open source information, many companies operate in an ad-hoc environment of threat intelligence subscriptions, industry colleagues, and a collection of personally-selected open source websites and RSS feeds. Many of the companies that subscribe to threat intelligence services found this relationship frustrating to manage, as they must specifically tell the services what issues/keywords they want data on. As the companies' intelligence needs evolve, there is latency in communicating the new needs to the service, and getting new intelligence information. The services' inconsistent sharing of keyword-specific information and lack of timeliness in doing so are ongoing issues, so the companies instead tackle open source data collection in any way they can. In many instances, this involved establishing traditional RSS feeds using applications such as Google Reader. The problem with this approach involves having to manually sift through hundreds of articles for relevant data, which inhibits consistent and accurate information collection. Furthermore, this data is notoriously difficult to correlate with other data sources (network data, social media, chat rooms, geopolitical news sites) and complicates 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 burdens analysts, complicates correlation, and contributes to missed analytic opportunities.

Current state:

- Government and private sector organizations alike reported challenges in efficiently collecting and integrating open source content into analytical products. Tradecraft requires government analysts to meticulously record information about the source; a time consuming, manual process that is prone to errors. Some government organizations copy swaths of open source data and migrate it onto classified networks so analysts can safely access and analyze the data. This requires the government organization to duplicate data that already exists, resulting in costly storage requirements, and unstructured open source data that is difficult to index and tag for analysts to sort through.
- Non-government organizations are inundated with data feeds that vary in format, so consumption and integration of information for further analysis is difficult. Some are tackling this issue by developing initial concepts for standard formats and delivery of data such as STIX, TAXII, and OpenIOC.

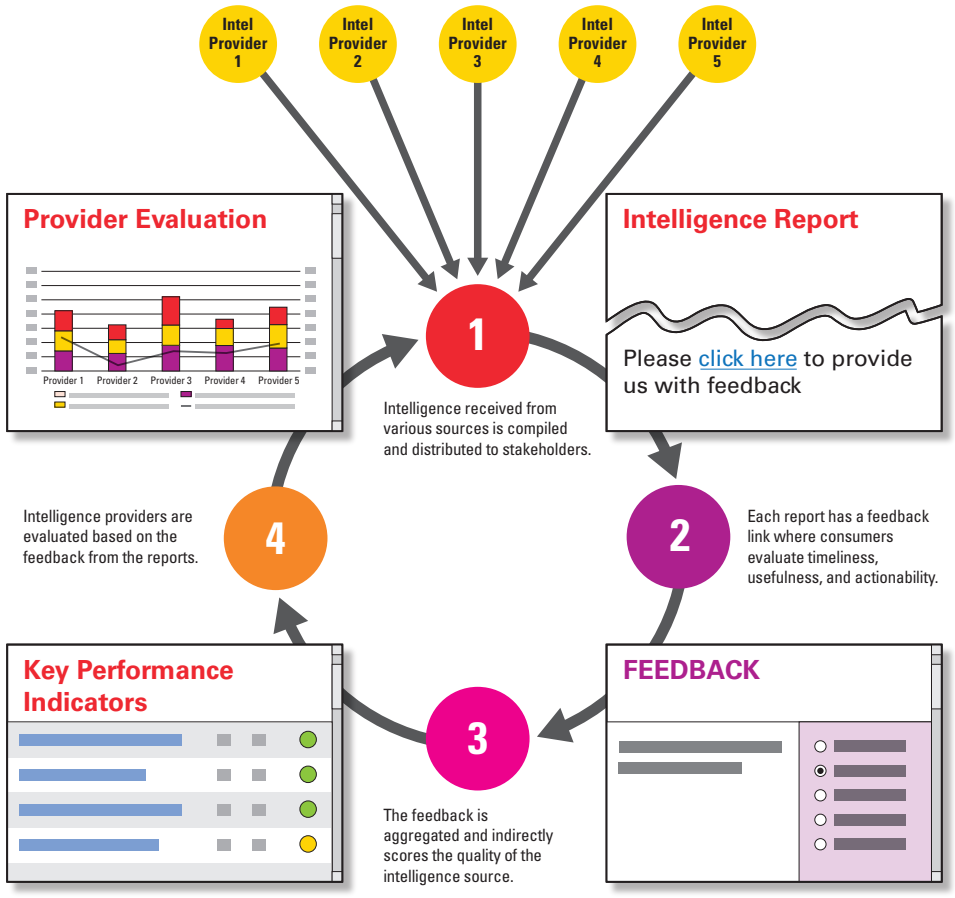

Best Practice #1: Repurposing search engine referral data

One participant was concerned with overseas competitors trying to duplicate a manufacturing process so that they could 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 Internet search engines to learn as much about that stage of the process as they could. Because the manufacturing process was proprietary, and very few companies could 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 and could anticipate what type of data was at the highest risk of being targeted for exfiltration.

Best Practice #2: Mind the gaps

One participant in the study wanted to ensure their organization 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 do effective cyber intelligence analysis. It then captured what data it was effectively able to collect itself, and then highlighted the remaining areas where it was missing coverage. Armed with these gaps, the organization worked with multiple third-party intelligence providers, and was able to provide specific data collection requirements. As intelligence services provided products to this organization, every report would be meta-tagged for indexing, and feedback from consumers on the reports collected. Using this consumer feedback, the organization was able to grade the quality of the reports it was getting for timeliness, actionability, and usefulness. Using this feedback, it worked with its vendors and influenced the type of reporting it would receive. Every year when these contracts were up for renewal, the organization was able to make smart investments on whether to continue a relationship with an intelligence provider, or seek out other means of getting that data. This process minimized the data gathering gaps, and ensured that their analysts didn't waste time on data gathering tasks that they knew were covered by external providers. It ensured that resources were smartly invested in data gathering, and that the organization didn't invest more money in obtaining data than necessary.

Functional Analysis

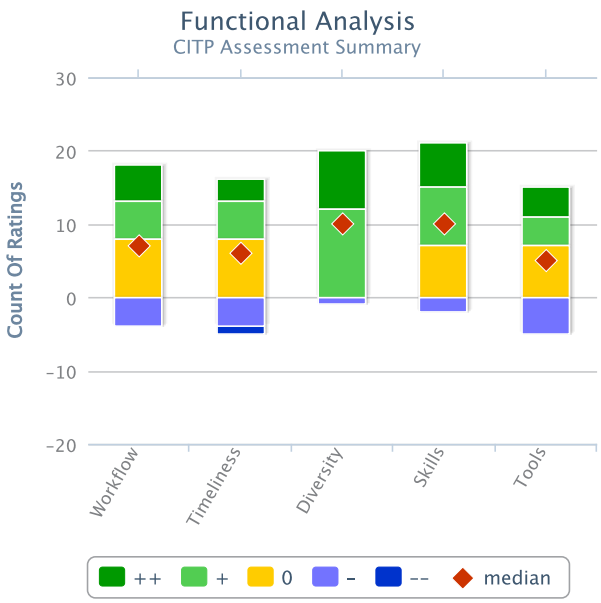

Organizations produce actionable, functional analysis when a workflow exists to extract pertinent data from internal and external feeds for the purpose of informing strategic analysts, decision makers, and consumers about the "what" and "how" of a cyber threat to detect, deter, and defeat it. The process begins with analysts taking data involving malware or email phishing schemes that was collected during data gathering and applying analytical tools and human resources to isolate potential threats. The data becomes intelligence when analysts evaluate its threat potential against source validation and multiple sources, including personal and industry expertise, organizational threat priorities, present day situational awareness, and historical references. Analysts then provide this intelligence verbally or through written means to internal strategic analysts, decision makers, and external consumers responsible for network defense or cyber security.

Challenge: Adopting a common cyber lexicon and tradecraft

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

Current state:

- Lexicon: During the baseline sessions, participants 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 if it was being provided by an entry-level analyst or manager.
- Generic terminology: Within the public and private sectors, measures exist to prevent unwarranted and unlawful disclosures of identifiable information, such as IP addresses and company names. Organizations protect these details when sharing threat information by referring to them with generic terms, such as "IP 1" and "Company B." While this ensures non-attribution, it inhibits 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 dissuades analysts from engaging in these types of analysis because it is time consuming and usually results in no additional substantive information, especially within the U.S. government.
- Tradecraft: Many U.S. government agencies 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 or corroborating sources. Numerous individuals with varying levels of this skillset have transitioned to cyber intelligence roles in the private sector, but the practice of assessing credibility remains largely absent. The numerous analytical products reviewed for this study either did not contain estimative or source validation language, or relied on the vendor providing the information to do the necessary credibility assessment.

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

Organizations struggle to accurately focus analytical efforts on critical threats because they cannot adequately filter out data that once analyzed ends up being classified as low to moderate threats.

Current state:

- Many functional analysts are inundated with potential threat information that their job responsibilities require them to analyze, only to determine most of it poses a low to moderate threat to the organization. These time consuming activities diminish the organization's accuracy in identifying critical threats and devoting the necessary resources to analyze them. Because of the tax this takes on human resources, some organizations are working with information/intelligence vendors to automate the analysis of these low to moderate threats.
- In the government, as well as a small set of organizations in the private sector, robust policy restrictions are used to filter out the low level threats. Restricting the ability to open executables, limiting use of commonly exploited software, prohibiting USB storage devices, and impeding access to websites that are associated with scams and malware make it very difficult for low sophistication hackers (recreational, or "script kiddies") to affect these networks. This frees 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 uses 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/cyber security components in a formal security alert. Per a predetermined timeline, the analyst works to produce an initial security alert with an 80% 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% 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 U.S. government organization leverage relevant environmental factors and intelligence requirements provided by strategic analysts to write scripts automating the distribution of network activity into threat categories that functional analysts can 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" provides the functional analysts a smaller data set from which to investigate potential new threats. This results in more timely and accurate functional analysis being provided to strategic analysts, decision makers, and consumers.

Strategic Analysis

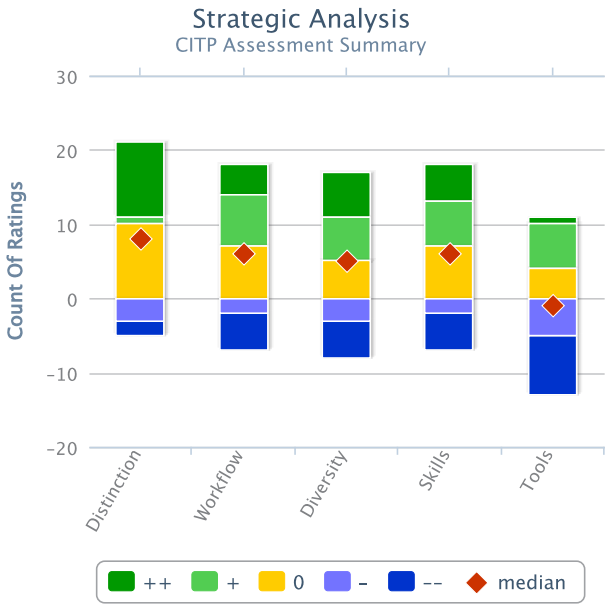

Strategic analysis adds perspective, context, and depth to functional analysis, and incorporates modus operandi and trends to provide the "who" and "why" of cyber threat analysis. Strategic cyber intelligence analysis is ultimately rooted in functional data, but incorporates information outside traditional technical feeds—including internal resources as physical security, business intelligence, and insider threat, and external feeds covering global cyber threat trends, geopolitical issues, and social networking sites. The resulting intelligence informs both technical and non-technical decision makers and intelligence consumers of the strategic implications these threats pose to the organization, industry, country, and economy.

Challenge: No industry standard for cyber intelligence education and training

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

Current state:

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

- The government dedicates a significant amount of resources to bridging the gap between analyst and security professional. Depending on the agency, government cyber intelligence analysts spend 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. A significant proportion of our participants have former government and military intelligence professionals within their teams and indicated a continuing desire to hire these individuals.
- Many of the organizations baselined claimed they prefer to train an analyst to understand the technical aspects of cyber security than to try and train an information technology person how to do intelligence analysis. It should be noted that despite voicing this opinion, the actual composition of their analytic staff all had technical backgrounds. When asked what an ideal candidate would look 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 must be dynamic enough to capture rapidly evolving tools, capabilities, and sophistication of adversaries.

Current state:

- Many of the intelligence methodologies observed in the government were developed in an era when intelligence analysts were focused on counting tanks, missiles, and airplanes held by hostile countries and predicting what the leaders of those countries planned to do with them. Applying these same processes, workflows, and tradecraft to the cyber domain is not always feasible. By the time a strategic-level product on an emerging threat makes it through the publication process, it's already out of date.
- Several process shortfalls compound the problem across government. One organization baselined mentioned strategic cyber intelligence products that took months to get published. In one organization, cyber intelligence analysis goes through the same production and review process as an analytic piece about topics with a much longer shelf life. A big part of the bottleneck is that cyber intelligence is sufficiently new that there is not a robust senior corps of cyber intelligence analysts. Since the process requires senior analysts to review all products, junior analysts are more familiar with the technology they are writing about than the senior reviewers. This delays the production process as complex, controversial, or unfamiliar topics are pushed to the end of the queue while easier, more mainstream reporting is prioritized.

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 evolve over time to ensure their network defenses are adequately prepared. One government agency has built profiles of adversaries including TTPs, malware used, tools, C2 infrastructure, names used, spear-phishing tactics, and common targets. Compiling this data has helped them to attribute new activity, and track the evolution of their adversaries. One commercial company extended the profile to include the motivation of the hackers, and how they make their money.

One commercial organization excelled in this area, even mapping the threat to potential sponsoring organizations. By doing open source research on the sponsoring organization, the company is able to narrow down the types of data that will likely 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 threat. As the motivations of the threat changed, this company was able to adapt its threat profile and identify new types of data that would be at risk. As organizations expand their business into overseas markets, this company was able to anticipate which threats this activity would potentially trigger, and incorporate these risks into the company's business strategy.

Best Practice #2: Global situational awareness

Cyber intelligence that looks beyond the organization's network perimeter provides strategic insights that feed predictive analysis in several participants. One organization uses a tool that provides visibility into the IP ranges of commercial partners. That way, when a vendor is compromised, the organization can take preventive measures and ensure that the malware doesn't spread into its networks, or that an attacker is not able to move laterally from the supplier's network into its network. Another company utilizes geo-political analysts to add context to what the cyber intelligence analysts are seeing. This company has an international supply chain, so the collaboration between the cyber analyst and the geo-political analyst often yields insights that better prepares the company's leadership when traveling overseas.

Another organization looks at what hackers are doing to other companies both inside its economic sector and around the world in areas where this company has major business interests. This organization looks at these external issues and attempts to determine if the issues could affect it in the near or long term. Examples included incidents at domestic services industries and international commerce entities. The organization then produces an "external breaches" slide for leadership that depicts these issues. Many of the events covered are selected because the organization has or might have business relationships with them; therefore, the threat could adversely affect them.

Reporting and Feedback

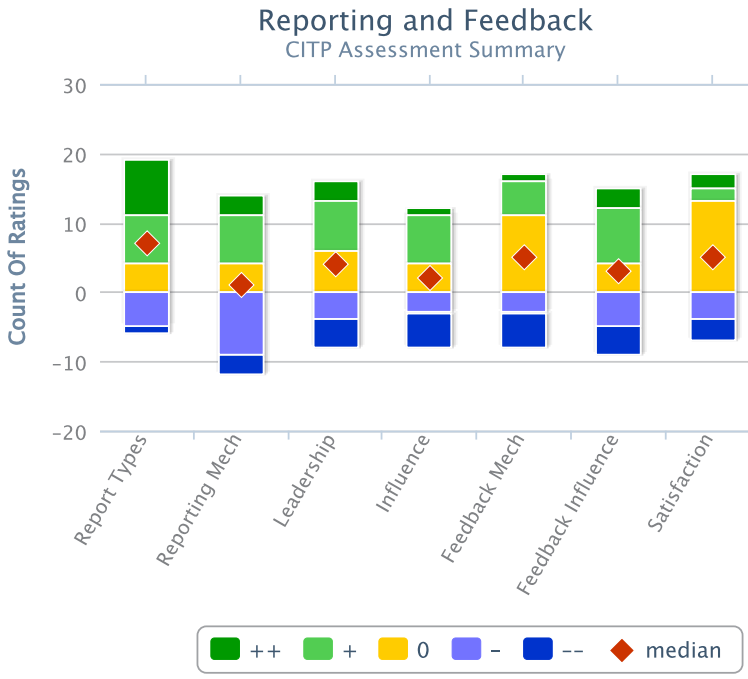

Strategic analysis is primarily intended for decision makers. The decision makers utilize this information to provide feedback to the cyber intelligence analysts, or adjust the direction of the organization. Both actions impact the cyber intelligence process, and redefine the environment, serving as a driving force to continue the cyber intelligence process through the analytical framework.

Challenge: Communicating "cyber" to leadership

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

Current state:

- The technical complexities associated with cyber security are difficult for many organizational leaders to appreciate. In the majority of organizations baselined, 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 are relatively new areas of focus, there is a dearth of senior leadership in the cyber intelligence field. This lack of a senior corps of cyber-savvy analysts means there's a lack of mentorship to junior analysts, which only perpetuates the problem of poor communication between non-technical leadership and the cyber community.

Challenge: Difficulty capturing return on investment

Organizations typically use return on investment 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 use performance measures that focus on quantity (e.g. number of reports generated), but not necessarily on quality or impact of intelligence. Analysts are encouraged to get feedback, but valuable feedback on intelligence products is limited and anecdotal. In the for-profit realm, performance measures, particularly those that can demonstrate return on investment, are critically needed. Seasoned practitioners are well aware of the value proposition and the potential costs of not engaging in cyber intelligence, but defining real metrics that can be used to justify resource needs and ensure corporate support is very difficult. Some organizations have the ability to easily assign dollar values to protected assets; others use the cost of recovery from compromise discovered too late. For many organizations, the measure of the value of cyber intelligence is elusive.

Best Practice #1: Failure analysis

One organization that was struggling to produce ROI metrics took the approach of looking at past events and capturing what the negative effects 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 is able to surmise what competitors knew about their events based on this analysis, and estimate 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 in an attempt 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. The message: had more resources been used to protect this data and track the competition's interest in this data, then we could have saved the money spent on creating diversions.

Best Practice #2: Carving channels for communication

A robust reporting approach considers content appropriate and necessary for the audience, relevant to the organization, with thought for frequency, timing, and delivery mechanism. An organization wanted to maximize the benefit of cyber intelligence analysis, and not only use it to support cyber security or senior leadership. First, the company identified groups of stakeholders, including senior leadership, risk managers, individual business units, security staff, and even the entire company workforce. Then, communication vehicles were established to reach out to each of these stakeholders. A monthly newsletter with tips on identifying spear-phishing, or safe Internet browsing habits targeted the general population. An email distribution list to the company's heads of risk management was created. The intelligence analysts were given a slide that was included in the weekly briefing to senior leadership. And the heads of each business unit were given their own email distribution list. From there, intelligence products were categorized by which stakeholders would benefit from this data, and intelligence was disseminated through the appropriate channel to reach each stakeholder. 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 intelligence analysts, it contained timely and pertinent information for their portion of the business.

This effort was coupled with proactive feedback, including a link at the bottom of each product that would allow recipients to comment on the utility of the reporting. Although feedback was initially slow to come in, the mechanism is in place and as managers identify new areas that they want the analysts to focus on, they can use the feedback channels to communicate these needs.

Conclusion

The goal of this study was to obtain a baseline understanding of the state of cyber intelligence. This report describes the results of the study organized around a notional cyber intelligence analytical framework. It defines cyber intelligence in the context of our study and identifies key gaps, challenges, and best practices in each of the component areas.

Overall, the study's findings indicate that the cyber intelligence analytical framework is a useful model for understanding the state of cyber intelligence. It ensures coverage of all functions consists of *define the environment, data gathering, functional analysis, strategic analysis* and *decision making; to include reporting and feedback*. For organizations to succeed in performing cyber intelligence, each component must be supported and amplified by accompanying processes, effective analytic tools, access to the right data, a competent workforce, and strong and visible organizational leadership offering support.

Organizations find it challenging to assimilate and analyze large quantities of poly-structured and unstructured data from both internal (e.g., proxy logs), and external sources (e.g., news feeds). In addition, a general lack of industry standards exists for organizations to consult to enhance capabilities—from establishing the cyber intelligence function, to education and training requirements, to developing and managing performance measures and return on investment figures. In this small, focused study, the findings suggest that many organization fall short in performing strategic analysis, producing reported analysis results, and soliciting useful feedback and incorporating that feedback into the process.

Organizations that routinely evaluate their cyber risks, maintain a keen awareness of threat actors, and proactively review and rate their risk exposure based on the threat (and not just vulnerability) are more likely to anticipate potential attacks and prioritize protection activities. Successful organizations also balance protecting the network perimeter with looking beyond it for strategic insights and predictive analytics.

Following this study, the SEI Innovation Center will publish a more detailed report on the state of cyber intelligence in summer 2013. The SEI Innovation Center also will continue to review the study's results and incorporate them into prioritized challenge areas already identified as having the greatest ability to enhance the state of cyber intelligence—visualization and analysis, data analytics, and training and education. Through Fall 2013, the SEI Innovation Center will explore these challenge areas through prototyping efforts validated by the participants.

This material is based upon work funded and supported by Office of the Director of National Intelligence 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.

DM-0000135