Lessons from Using the I-Corps Methodology to Understand Cyber Threat Intelligence Sharing

Josiah Dykstra, Matt Fante, Paul Donahue,
Dawn Varva, Linda Wilk, Amanda Johnson
U.S. Department of Defense

Abstract

Cybersecurity researchers and practitioners continually propose products and services to secure and protect against cyber threats. Even when backed by solid cybersecurity science, these offerings are sometimes misaligned with customers' practical needs. The Innovation Corps (I-Corps) methodology attempts to help innovators, researchers, and practitioners maximize their success through deliberate customer discovery. The National Security Agency (NSA) has adopted I-Corps for internal innovation and optimization. In February 2019, NSA Cybersecurity Operations embarked on a study using this methodology to explore cyber threat intelligence sharing. Information sharing is a foundational practice in cybersecurity. The NSA also shares cyber indicators with authorized partners, and sought to understand how partners consumed and valued the information to better tailor it to their needs. After 60 customer discovery problem interviews with over 20 partners, six primary themes emerged. We describe our experiences using the I-Corps methodology to study and optimize internal processes, and lessons learned from applying it to information sharing. These insights may inform future applications of I-Corps to other areas of cybersecurity research, practice, and commercialization.

1 Introduction

Cyber attackers pursue many targets using the same tools, techniques, and infrastructure. As a result, community and commercial sharing of threat intelligence and indicators has become commonplace. Gartner defines threat intelligence as "evidence-based knowledge, including context, mechanisms, indicators, implications and actionable advice, about an existing or emerging menace or hazard to assets that can be used to inform decisions regarding the subject's response to that menace or hazard" [17]. For cyber threats, threat intelligence includes suspicious or known-bad email addresses, URLs, IP addresses, malware signatures, and behavior.

There are numerous well-known challenges with sharing cyber threat intelligence (CTI) [15]. These challenges range from protecting privacy, proprietary, or classified information [12] to interoperability and technical exchange formats [14]. Some CTI requires human intervention to avoid business disruption, increasing the cost to deployment and protection. CTI feeds are notoriously large and noisy [6]. One under-studied problem is discerning the value of shared intelligence. The consumers of CTI rarely provide feedback to the provider about the utility of an individual indicator. Consumers may also have difficulty assessing the security outcomes and value of shared CTI from a threat feed that could cost $150,000 per year [16].

Other research has primarily explored technical production aspects of CTI sharing. One study interviewing ten experts found that the primary factors affecting shared CTI related to limitations with integrating and consolidating CTI from different sources while also ensuring the data's usefulness [20]. As future research, the researchers suggested investigating threat intelligence use and impact. Many studies have identified quality issues as a barrier to effective CTI sharing, including relevance, timeliness, accuracy, comparability, coherence, and clarity [23]. The corollary is that consumer value and feedback are rarely captured. Platforms for sharing and managing threat feeds have continued to evolve, and some have suggested that the problem is shifting from creating such systems to generating value from the information [9].

The United States government has a role in sharing CTI. The Cybersecurity Information Sharing Act of 2015 (CISA) requires various federal government departments and agencies to develop procedures which promote voluntary sharing of CTI with federal and non-federal entities [1]. Among the examples given in CISA are the Department of Homeland Security (DHS) Automated Indicator Sharing (AIS) initiative [11] and Department of Energy's Cybersecurity Risk Information Sharing Program (CRISP) [10]. The National Security Agency (NSA) is authorized to share classified and unclassified cyber threat intelligence with authorized partners who defend their own networks and who may also share with their customers. The NSA shares with both "first party" U.S. government partners such as DHS [18] and "second

party" intelligence community partners (Australia, Canada, New Zealand, and United Kingdom) [5]. One important consideration for NSA's sharing is the equity decision between cyber defense and protecting sensitive sources and methods [22].

In February 2019, leadership in NSA Cybersecurity Operations commissioned a team to use I-Corps and explore existing CTI sharing by the NSA and propose changes if necessary. This work describes the methodology and key findings from customers who receive CTI from the NSA. The paper is organized as follows. In Section 2, we present the I-Corps methodology. In Section 3, we describe how the NSA used I-Corps to study CTI sharing and the results from the study. Section 4 presents lessons learned from our experience for those wishing to consider our approach. Section 5 contains our conclusions.

2 Study Methodology

In this section we introduce I-Corps and provide an overview of the customer discovery process.

2.1 I-Corps

Innovation Corps (I-Corps) is a methodology developed by the National Science Foundation (NSF) based on Steve Blank's *Lean LaunchPad* course [13]. Lean LaunchPad is an approach to *lean startup*, a methodology for refining startup businesses and products through experimentation and rapid, iterative product design. Lean startup is sometimes compared with *design thinking*, a related approach to innovation. A key difference between these methodologies is where the product is introduced in the innovation cycle. In design thinking, the approach is to first establish the need for a product or service. The lean startup approach is to begin with a viable product, and make small, fast incremental changes to evolve the design using feedback from users.

Lean LaunchPad is based on the scientific method and has three parts: 1) the Business Model Canvas [21], to frame hypotheses; 2) Customer Discovery, to test those hypotheses in front of customers; and 3) agile engineering, for rapid and collaborative product development. Customer discovery is the portion described in this paper. Lean LaunchPad is now taught at over 50 universities across the U.S., and I-Corps is offered in 88 universities.

For NSF, I-Corps guides academics to transfer their research into successful commercialization through a disciplined process of customer discovery and experimentation. The basic premise of I-Corps is that entrepreneurs will be more successful if they align their products and services to customers' actual problems. These insights come from interviewing a range of potential customers.

Blank's original methodology designed for startup companies has been modified for other ends. In 2016, Blank and the U.S. Department of Defense (DoD) recognized this need in developing "Hacking for Defense" and adapted portions of Lean LaunchPad to focus on a mission rather than profit [8]. This adaptation included replacement of the Business Model Canvas with a Mission Model Canvas that is suited to users who aim to create value for beneficiaries (such as warfighters) rather than earn money. The DoD has also adopted I-Corps to accelerate the transition and commercialization of DoD-funded research [4]. The DoD solicits applications from current and recent DoD awardees on basic research topics to receive mentoring and funding to accelerate the transition and commercialization of the funded research.

The NSA adopted I-Corps in 2015 as one approach to innovation in a similar but distinct way from NSF and the DoD's uses [19]. A full-time team of NSA I-Corps staff train and coach internal project teams not towards commercialization, but for increased speed of deployment and impact to NSA missions. The goal is to help innovators and existing project owners to optimize their offerings for internal and external consumers. The NSA also used the I-Corps process for their Unfetter [2] and WALKOFF [3] projects.

2.2 Customer Discovery

I-Corps is predicated upon solution providers effectively understanding the practical problems of potential customers through semi-structured interviews. The methodology espouses that the customer discovery must be done by solution providers themselves and that direct learning cannot be outsourced to other investigators. Customer discovery is not a focus group, but allows a solution provider to validate their hypotheses about who their customers are and what actually matters to them. The primary output of this process is customer insights to help the team determine if iteration or pivoting (substantially changing their proposal) would lead to customer adoption. There are three stages to customer discovery described below: pre-planning, interviews, and analysis and insight.

Pre-planning. The first stage of preparing to engage with customers is pre-planning. The team begins by stating the assumed problem and value proposition of the innovation they wish to develop or improve for customers. Group brainstorming is used to define problems with specificity. Similar to the scientific method, the team must define hypotheses and assumptions about customer problems, processes, and needs. The team develops a series of open-ended questions to frame interviews with customers. Teams are given an ambitious goal to interview 100 individuals during the next stage, though the exact number is not prescribed. This number is intended to drive the team towards a full understanding of the field of customers and be able to correctly define one or more customer archetypes. In cybersecurity, these subjects may include end-users, managers, and security professionals, depending on the problem being solved. This stage should last a few days.

Interviews. The second stage is customer interviews. This

stage consumes the majority of the team's time over the course of several weeks as interviews are arranged and conducted in parallel by members of the team. I-Corps emphasizes the importance of one-on-one in-person engagements. The team takes care to frame the interview as many customers are accustomed to sales presentations and demos and not someone simply listening to learn as they seek to understand the customer at a deep level. The interviewer covers the key questions developed during pre-planning but allow the conversation to flow. Given the opportunity to share their work and frustrations, customers may reveal unexpected insights to the interviewer. The primary goals of the interview are to understand the customer and his or her problems, and to validate the interviewer's hypotheses and assumptions without offering possible solutions. The interview stage is complete when the team can predict with consistency what they expect similar customers to say during an interview.

Analysis and Insight. The third stage is analysis of the qualitative data from the interviews. The key objective in this step is thematic analysis across customers, and developing customer archetypes and segments. The output of this stage are customer problem statements suggestive of success metrics from the customer perspective. For example: "Customers cannot immediately utilize cyber threat intelligence because of technical format differences, resulting in a delay in protecting their networks and distributing to their customer base." Potential success metrics for this problem could include reduced latency and CTI efficacy.

By delivering a problem statement and initial metrics of success, the team reaches the problem validation milestone and they can begin to explore solutions, having already validated the market for a solution.

The search for a solution begins with ideation. Here teams take a structured approach to divergently consider a wide range of possible solutions or improvements to address customers' problems identified during analysis. They then converge on a short list of the most promising ideas for subsequent testing via a series of minimum viable products. We consider the team to have reached solution validation when a group of early adopters have produced a measurable mission impact (e.g. analyst time saved) with a solution prototype. At this point, teams decide if the solution is worth building (and ultimately scaling) within the established corporate architecture based on early evidence of mission impact.

3 Study Results

In this section, we describe how we applied the I-Corps methodology to study NSA CTI sharing. Organizational leaders gave the CTI Sharing I-Corps Team eight weeks for the task. The team comprised nine cross-organizational subject matter experts who devoted 50% of their work time to the project. A senior steering group of three executives met with the team weekly to help ensure that the team had the resources and knowledge required. Two of the NSA's I-Corps mentors conducted a one-day training session about the I-Corps methodology at the kickoff and offered weekly coaching sessions with the team throughout the project.

Pre-planning. The team began by defining the problem as follows: "Reimagine how the NSA Cybersecurity Enterprise shares information with our customers for optimal cybersecurity outcomes." The team consolidated a list of known customers who receive CTI from the NSA. Give the limitation of time, they selected a subset of customers and divided into two teams, one focused on First Party partners and one on Second Party partners. Interviewees included front-line network defenders, network operators, managers, liaison officers, and integrees from other U.S. government departments and agencies, DoD, and other counterparts in partner countries. The team developed nine open questions listed in Table 1 to solicit feedback from customers about their experiences and challenges in using CTI shared with them.

Interviews. In total, 60 interviews were conducted with consumers of NSA CTI from more than 20 customer organizations. Finding the right contacts, connecting with the individual, and scheduling the interviews was the most time-consuming task, ranging from one or two days to a few weeks.

For most interviews, one or two team members met in person with the interviewee for 30 minutes. A small number of interviews were conducted by video conference or phone, and in some cases, interviews were conducted in groups. Soon after the interview, the team documented notes in a shared digital repository.

Analysis and Insight. The CTI Sharing Team analyzed interview notes individually and as a group. Six themes emerged across the customers who were studied:

1. Customers cannot immediately use NSA classified information, resulting in a delay in protecting networks and disseminating to their customer base.

2. Customers lack holistic awareness of NSA products and services, and therefore experience inconsistent adoption of NSA cybersecurity information and under-utilization of available information.

3. Customers experience delays in intelligence report dissemination, resulting in networks running at risk or diminishing the effectiveness of the information.

4. Customers lack a clear understanding of the technical context surrounding events impacting their ability to fully and effectively mitigate vulnerabilities.

5. Customers expend time and resources to manually adopt signatures and work through format issues, resulting in implementation delays.

6. Customers lack a clear understanding of the attribution surrounding events, impacting their ability to effectively

1. Tell us a bit about your organization and your role. What are your goals? What information do you need to accomplish your goals?
2. What cyber threat information do you currently receive from the NSA?
3. How would you characterize the effectiveness of the information you receive from the NSA?
4. Please provide an example or describe how NSA information has provided a benefit to you or your team?
5. What frustrations or concerns do you have about how NSA information integrates into your workflow?
6. What would help you better accomplish your job? What are your challenges?
7. What would you need to better support your goals?
8. Who else should we talk to?
9. May we follow up with further questions, if needed?

Table 1: Guiding questions used during semi-structured customer interviews.

mitigate vulnerabilities and anticipate the adversaries' next steps.

The team did not differentiate problems specific to the First Party and Second Party groups. Other than differences related to security classification, customers receive the same CTI. The NSA's CTI sharing teams regard partners equally, but also recognize that some customers represent greater cybersecurity outcomes by protecting large or high-value networks and sharing with further downstream customers.

Given the time alloted, the Study Team focused on providing recommendations for the first two customer problems. These two problems were the most prevalent across the interviewees, and therefore offered the most potential value if addressed. The others were left for future work. Problem 1 was unsurprising and widely suspected by the team, but direct evidence from customers provided critical evidence and validation. That Problem 2 was among the very highest customer-generated issues was surprising for the team and for leadership. The team had hypothesized that the lack of feedback on individual indicators was a result of lack of awareness about their utility, but customers reported both an unawareness of available data and how or where to provide feedback. This insight was possible because of the customer discovery process.

For Problem 1, the team brainstormed more than 20 solutions related to the issue of customer challenges in using classified information. The I-Corps mentors guided the group in several sessions to generate this list. Next, the team grouped the solutions into common themes that emerged as people, process, policy, training, technology, and funding. Upon review, they curated three proposed solutions of greatest potential for addressing the problem. In a follow-up survey with the original interviewees, customers validated that the proposed solutions could improve sharing. The proposed solutions were:

1. Distribute guidance inside NSA about the criticality of increasing the amount of CTI at the unclassified level.
2. Reinforce guidance to evaluate equities between using CTI to protect customers and the need to protect sensitive sources and methods.
3. Establish an NSA working group to evaluate internal equities review processes and identify areas for improvements.

For Problem 2, the team brainstormed more than 20 solutions related to the lack of awareness about NSA products and services. These proposals spanned the domains of people, process, policy, training, and technology. Upon review, the team curated three proposals for immediate action:

1. Appoint an NSA outreach team to create a preliminary customer knowledge repository, such as web portal or catalog.
2. Engage with user interface solution providers inside NSA already working towards personalized and tailored customer service.
3. Create a comprehensive CTI portal with resources for customers, including contacts, federated queries, training, and multi-classification CTI.

The CTI Study Team also produced several general recommendations about CTI sharing as a result of the study. For example, they strongly emphasized that organizational

metrics and success criteria should be focused on cybersecurity outcomes of the partners, not simply on the volume of sharing. The team also urged leadership to review and emphasize clear and concise policies for releasing threat intelligence to the broadest audience by default. Finally, they suggested revisiting customers after the implementation of their recommendations to assess the value of the changes.

4 Lessons Learned

This study was one of the NSA's first I-Corps experience with a cybersecurity problem and offered lessons for all those considering future I-Corps engagements. At the conclusion of the project, the NSA's I-Corps coaches conducted 15-minute retrospective reviews with each team member and steering group member. The process produced several lessons learned that may inform others considering the I-Corps methodology for cybersecurity topics.

First, in a large organization such as ours, it can be challenging to know how to select participants for an I-Corps team. Our criteria included people with a relationship to the study area (CTI sharing), organizational diversity, technical and demographic diversity, and open-mindedness. Team members were mostly invited by name, and the senior steering group vetted each participant. At the conclusion of the project, organizational leadership expressed satisfaction with the team, although some team members reported concerns about some areas of expertise in the group as a whole.

Second, the CTI Study Team required more time than anticipated in pre-planning. Rather than a few days, defining scope and execution required almost two weeks. We hypothesize that this may stem from building the team with diverse backgrounds from our large enterprise, who then needed time to orient and baseline across the team. The additional time was an acceptable cost to having a diverse team. Similar dynamics could occur in other new groups experiencing normal team formation, and leaders should account for this possibility.

Third, the CTI Study Team reported a desire for more upfront I-Corps training. Our I-Corps office offers a five-day training class to those wishing to learn the methodology. In the interest of efficiency for the eight-week CTI Sharing I-Corps, the coaches and steering group decided to conduct training continually throughout the project, starting with a one-day I-Corps orientation. The belief was that just-in-time training would improve efficiency. In retrospect, team members expressed a desire to have received more instruction about the entire I-Corps process in the beginning.

Fourth, we discovered that direct supervisor support for members of the I-Corps team was essential to success. Our I-Corps coaches advised in the beginning that leadership support was an essential prerequisite for success. Organizational senior leadership was unwavering in their support, with messaging to the entire workforce. Direct supervisors of the CTI Study Team members were not directly consulted in the planning or execution of the project. As a result, some team members did not feel empowered to use 50% of their work time on the project. In a large organization with several levels of hierarchy, it can be especially important to seek support from direct supervisors of team participants. Other researchers have suggested that managers must be active participants in design thinking, particularly because unexpected findings during the process can generate defensiveness and fear in participants [7]. In addition to the engagement and leadership from our senior steering group and I-Corps coaches, we may seek an increased role for direct supervisors.

We were pleasantly surprised that customers were willing to honestly offer positive and negative feedback about their experience with our CTI sharing practices. We expected that customers might be uncomfortable offering negative comments to the team in face-to-face interviews. We identified only one partner for whom this seemed to occur and who provided entirely positive comments during the interview despite consistently negative feedback to others offline.

Finally, the NSA has found that I-Corps is a fruitful methodology for some problems but not others. We recognize that I-Corps is one of many techniques for innovation in cybersecurity. When evaluating whether I-Corps is an appropriate approach, there are several factors to consider. This approach requires sufficient time and resources to effectively engage with customers, typically six to eight weeks. In addition to localized solutions, I-Corps is especially suited to problems for which external customers can offer insight. Because the evaluation of information sharing depends largely on customers, I-Corps was an appropriate and worthwhile approach.

5 Conclusions

In this paper, we discussed the I-Corps methodology for customer discovery as applied to cybersecurity. To illustrate the use of I-Corps, we described how we were able to identify key customer problems and potential solutions for improving cyber threat intelligence sharing. This methodology generalizes for the study and optimization of many areas of cybersecurity.

In the future, we look forward to continued experimentation with I-Corps to increase team satisfaction, efficiency, and cybersecurity outcomes. The solutions we proposed related to information sharing have not yet been tested to evaluate how effective they are for customers, though changes to CTI sharing are already underway. We plan to report these findings in the future as NSA continues to use I-Corps for mission innovation and optimization.

Finally, we intend to further study information sharing, and in particular the cost of CTI production and the value of shared cyber threat intelligence. There is a widely held belief in the security community that information sharing improves security posture and produces greater defensive agility [15]. Research is needed to evaluate the validity of these claims and deliver threat intelligence the protects users and networks.

6 Acknowledgements

We thank the anonymous reviewers for their comments and suggestions. We especially thank the anonymous team members and contributors to the study.

References

[1] Senate Resolution 754, Cybersecurity Information Sharing Act of 2015 (CISA). https://www.congress.gov/bill/114th-congress/senate-bill/754.

[2] Unfetter Project. https://nsacyber.github.io/unfetter/.

[3] WALKOFF. https://nsacyber.github.io/WALKOFF/.

[4] I-Corps @ DoD Funding Announcement. https://basicresearch.defense.gov/News/Articles/News-Display/Article/1490285/i-corps-dod-funding-announcement/, 2018.

[5] National Security Agency. UKUSA Agreement Release 1940-1956. https://www.nsa.gov/news-features/declassified-documents/ukusa/, 2016.

[6] Omar Al-Ibrahim, Aziz Mohaisen, Charles A. Kamhoua, Kevin A. Kwiat, and Laurent Njilla. Beyond free riding: Quality of indicators for assessing participation in information sharing for threat intelligence. *CoRR*, abs/1702.00552, 2017.

[7] Christian Bason and Robert D. Austin. The right way to lead design thinking. *Harvard Business Review*, March–April 2019:82 – 91, 2019.

[8] Steve Blank. The mission model canvas - an adapted business model canvas for mission-driven organizations. https://steveblank.com/2016/02/23/, 2016.

[9] Sarah Brown, Joep Gommers, and Oscar Serrano. From Cyber Security Information Sharing to Threat Management. In *Proceedings of the 2nd ACM Workshop on Information Sharing and Collaborative Security*, WISCS '15, pages 43–49, New York, NY, USA, 2015. ACM.

[10] Department of Energy. Energy Sector Cybersecurity Preparedness. https://www.energy.gov/ceser/activities/cybersecurity-critical-energy-infrastructure/energy-sector-cybersecurity.

[11] Department of Homeland Security. Automated Indicator Sharing (AIS). https://www.dhs.gov/cisa/automated-indicator-sharing-ais.

[12] Gina Fisk, Calvin Ardi, Neale Pickett, John Heidemann, Mike Fisk, and Christos Papadopoulos. Privacy principles for sharing cyber security data. In *Proceedings of the 2015 IEEE Security and Privacy Workshops*, SPW '15, pages 193–197, Washington, DC, USA, 2015. IEEE Computer Society.

[13] National Science Foundation. I-Corps Resources. https://www.nsf.gov/news/special_reports/i-corps/resources.jsp.

[14] Scott E. Jasper. U.S. Cyber Threat Intelligence Sharing Frameworks. *International Journal of Intelligence and CounterIntelligence*, 30(1):53–65, 2017.

[15] Chris Johnson, Lee Badger, David Waltermire, Julie Snyder, and Clem Skorupka. NIST Special Publication 800-150 Guide to Cyber Threat Information Sharing, 2016.

[16] Craig Lawson, Ryan Benson, and Ruggero Contu. Market Guide for Security Threat Intelligence Products and Services (ID G00380381). Retrieved from Gartner database, 2019.

[17] Rob McMillan. Definition: Threat Intelligence. https://www.gartner.com/en/documents/2487216, 2013.

[18] Office of the Inspector General of the Intelligence Community. Joint report on the implementation of the cybersecurity information sharing act of 2015. https://oig.justice.gov/reports/2018/AUD-2017-005.pdf, 2017.

[19] Greg Otto. NSA, other feds using innovation to improve security. https://www.fedscoop.com/nsa-other-feds-using-innovation-to-improve-security/, 2016.

[20] Christian Sillaber, Clemens Sauerwein, Andrea Mussmann, and Ruth Breu. Data Quality Challenges and Future Research Directions in Threat Intelligence Sharing Practice. In *Proceedings of the 2016 ACM on Workshop on Information Sharing and Collaborative Security*, WISCS '16, pages 65–70, New York, NY, USA, 2016. ACM.

[21] Strategyzer. Business Model Canvas. https://www.strategyzer.com/canvas/business-model-canvas.

[22] B. Michael Thomas and Neal L. Ziring. Using classified intelligence to defend unclassified networks. In *Proceedings of the 2015 Hawaii International Conference on System Sciences*, HICSS '15, pages 2298–2307. IEEE, 2015.

[23] Wiem Tounsi and Helmi Rais. A survey on technical threat intelligence in the age of sophisticated cyber attacks. *Computers & Security*, 72:212 – 233, 2018.