NEW AMERICA

CYBERSECURITY INITIATIVE

PAULO SHAKARIAN

THE ENEMY HAS A VOICE

Understanding Threats to Inform Smart Investment in Cyber Defense

FEBRUARY 2017

About the Authors

Paulo Shakarian, PhD, is the CEO and co-founder of Cyber Reconnaissance, Inc. (CYR3CON™). He is also a Fulton Entrepreneurial Professor at Arizona State University and New America Cybersecurity Initiative fellow. He has authored several books about cyber security including Cambridge's *Darkweb Cyber Threat Intelligence Mining* and Elsevier's *Introduction to Cyber-Warfare*. His work on cyber threat intelligence and artificial intelligence has been featured in *Forbes,* the *New Yorker, Slate,* the *Economist, Business Insider, TechCrunch*, and BBC. In 2016, the company he leads was named a semi-finalist in the Cisco Innovation Grand Challenge and selected for the NSF Innovation Corps program. He also named a "KDD Rising Star" by Microsoft Research and is a recipient of the Air Force Young Investigator Award. Previously, Shakarian was a Major in the U.S. Army where he was Defense Advanced Research Projects Agency (DARPA) Service Chief's fellow, faculty at West Point, and served two combat tours in Iraq, earning a Bronze Star and the Army Commendation Medal for Valor.

About New America

New America is committed to renewing American politics, prosperity, and purpose in the Digital Age. We generate big ideas, bridge the gap between technology and policy, and curate broad public conversation. We combine the best of a policy research institute, technology laboratory, public forum, media platform, and a venture capital fund for ideas. We are a distinctive community of thinkers, writers, researchers, technologists, and community activists who believe deeply in the possibility of American renewal.

Find out more at **newamerica.org/our-story**.

About the Cybersecurity Initiative

The goal of New America's Cybersecurity Initiative is to bring New America's focus on big ideas, bringing together technology and policy, and public engagement to the cybersecurity conversation. In doing so, the Initiative provides a look at issues from fresh perspectives, an emphasis on cross-disciplinary collaboration, a commitment to quality research and events, and dedication to diversity in all its guises. A collaboration between New America's Open Technology Institute and International Security program, our work explores important cybersecurity policy questions at all levels of government and policy making, from the state and local to national and international. Examining issues from the vulnerabilities equities process in governments and the importance of cybersecurity policy at the state and local level to the potential of strong and stable international regimes to promote better cybersecurity, New America's Cybersecurity Initiative seeks to address issues others can't or don't and create impact at scale.

Our work is made possible through the generous support of the William and Flora Hewlett Foundation, the Arizona State University, Microsoft Corporation, Symantec Inc., The Home Depot, Endgame Inc., and Facebook.

Find out more at **newamerica.org/cybersecurity-initiative**.

Contents

INTRODUCTION

In late July 2016, flight status screens in airports in Hanoi and Saigon broadcast derogatory messages toward Vietnam and the Philippines instead of the normal flight information. This cyber-attack, allegedly conducted by the notorious 1937CN Chinese hacking group, led to mass disruption to numerous travelers in the Asia-Pacific region that day.[1] However, this was not a random event, but rather the latest in a series of escalating back-and-forth cyber-attacks between China and Vietnam. Understanding such ongoing conflicts, as well as the capabilities of groups like 1937CN can help organizations better brace themselves when most at risk. For much of the past, cybersecurity measures have focused on looking internally at the vulnerabilities of an enterprise network. While this will continue to remain important, we will not obtain substantial improvement in cybersecurity of our infrastructure until we adopt an approach that is focused on the adversary. In short, the enemy has a voice in what happens—and we should expect attackers to adapt, innovate, and leverage community resources. Quality information on these enemies, also known as cyber threat intelligence (CTI), helps defenders better understand what vulnerabilities their likely adversaries will seek to exploit.

Taking a threat-focused approach to cybersecurity seems like a natural and sensible thing to do for organizations from small and medium enterprises to massive government agencies. In understanding the nature of the threat they face—that is to say who might be interested in breaching their security and why—they are able to craft better informed and data-driven security policies and maximize the return on their cybersecurity investments by identifying specific pressure points and crafting solutions that produce outsized impact. This reality makes ensuring a thriving market for CTI directly relevant to policymakers tasked with crafting policies that promote better national cybersecurity. This report is designed to help policymakers better understand what CTI is and how they can leverage it to help achieve public policy goals.

In this report, I start by discussing the general concept of CTI and how this powerful concept can reduce “offensive dominant” nature of cybersecurity and describe various types of such information. Then, to make the ideas a bit more concrete we examine how such information can provide insight into malicious hacker communities—in particular those on the deep and dark web. I then outline some challenges with cyber threat intelligence going forward and propose policy ideas that can help lead to improved access to such information across a variety of organizations.

THREAT INTELLIGENCE VS. THE "OFFENSE DOMINANT" CYBER PARADIGM

Cybersecurity is often referred to as offense dominant, meaning that the domain generally favors the attacker.[2, 3] The reasoning behind this is simple: a successful defense must block all pathways to a system while a successful attack requires only one. As the old hacker adage goes: "the defender must always be right—the attacker only needs to be right once."

This notion of an offense dominant cybersecurity stems directly from "best practices" in the field. These methods primarily rely on technical measures to improve defense. Traditionally these have included variations on patch management, firewall usage, intrusion detection, and anti-virus. However, an adversary particularly keen on gaining access to a system can study such defenses with the goal of finding the gaps. These actions are not limited to nation states or large criminal enterprises. The community of malicious hackers is a key perpetrator for these activities. While important, technical defense measures alone are unlikely to halt attackers. The offense will have the advantage in this case.

A threat-focused approach, while sensible in principle, can be difficult to implement in practice. While recent events have prompted discussions of the threat posed by major adversaries, many of these discussions remain superficial and not useful to actual network defenders. A useful threat-focused approach necessitates gathering information—or intelligence—on adversaries, and obtaining valuable threat information is a difficult endeavor. Obtaining high-grain intelligence on the activities of malicious hackers is a manpower-intensive task requiring many analysts. As a result, firms that provide this as a service today have yearly price tags into the seven digits.[4] This inherently precludes medium and small-cap companies from leveraging this valuable information. By the connected nature of our industries, infrastructure, and government ignoring the cybersecurity needs of the mid-tier invites peril to all. The companies at the mid-tier are often the very ones that provide utility infrastructure, fill niche supply-chain needs, and provided important out-sourced contracting support.

However, having intelligence on the adversary shifts this paradigm. By gaining insights on the adversary's behavior, we can better address the offense-dominant problem inherent in cybersecurity. This can be done in several ways. For instance, simply understanding how a peer organization was compromised is valuable information—provided the hacker did not attack your organization at the same time. Establishing a honeypot—a dummy system online with purposeful vulnerabilities to attract would-be attackers—can be viewed as a more proactive version of information sharing. Threat intelligence can also be found in earlier stages of an attack. For instance, the musings of hacking collectives on social media can provide early indications of major campaigns—such as those performed by Anonymous—especially if your organization is politically sensitive in nature. In the example in the beginning of the paper, we also see how highly contentious political situations can lead to cyber war—and knowing the cyber-political atmospherics can drive important security decisions and resource expenditure. Public-private information sharing is yet another venue emerging in various parts of the country—though tensions of confidentiality on both sides lead to cautious progress. Finally, later in this paper we will look at obtaining information from the deep and dark web, where we can get a glimpse of the malicious hacker ecosystem—understanding what malware and exploits are actively being sold, what hacking services are being requested, and how the community evolves over time.

> **Highly contentious political situations can lead to cyber war—and knowing the cyber-political atmospherics can drive important security decisions and resource expenditure.**

However, this is not to say that gathering intelligence on cyber threat actors is a trivial or easy matter. By their very nature, these threat actors must be as nimble and asymmetric as the technologies they both infiltrate and weaponized. The domain of cyberspace is unique in that impact can be more easily made by non-state actors including criminal enterprises and terror groups—and the barriers to entry for rogue nation states is much lower and easier to hide than with the development of conventional weaponry. Further, the line between suppliers and users of cyber weaponry is often blurred. Still, despite these aspects of the domain, cyber threat actors remain goal-driven, resource-constrained, and liable to leverage available community resources to carry out their operations. As such, a variety of sources of information (as previously discussed) are necessary to paint a picture of the threat.

Classifying Cyber Threat Intelligence

Clearly all of this information is useful in different time horizons. Understanding a recent attack on a peer organization may drive rapid defense measures in an effort to identify existing compromise and even initiate remediation. On the other hand, understanding changes in the hacker black market can drive more strategic decision-making. Prioritizing purchase and patch decisions can minimize exposure to vulnerabilities upon which darkweb malicious hackers are focused.

To better get our arms around the variety of information that can be considered as “threat intelligence,” we suggest classifying such cyber practices in tiers. The first and most basic tier is the situational awareness (SA). This entails the understanding of one’s own enterprise, of peer organizations (i.e. shared through an ISAC), and information obtained from sensors used to gather malicious activities in the wild (i.e. honeypot information). The second tier is the simplest form of *proactive* intelligence—the identification of an imminent threat to an organization. A prime example of this type of intelligence is indicators of pending hacktivist activities gained from social media. A third tier is a bit more advanced and forward leaning—understanding a dynamic shift in enemy *capabilities*. This involves developing knowledge of what exploits and malware are being developed. The most strategic fourth tier of intelligence encompasses general knowledge about the malicious hacking communities. Intelligence in this tier includes information about market dynamics within these communities, the rise and fall of particular personalities and venues, the nature of the conversations that take place in the forums, and the overall evolution of these communities.

The intuition behind this tiered system **(See Figure 1)** is that, at the higher tiers (tiers 3 and 4) the information can lead to decisions with progressively more long-term consequences. For example, for intelligence in the SA or first tier, the primary action that can be taken is to identify a signature or block an IP address. When a malicious hacker observes a sufficient number of organizations on his target list taking such actions, he changes tactics. Hence, the result is very short term. Likewise, preparing for an imminent cyber campaign (tier two) can lead to actions that will cause cyber defenders to make adjustments that last for weeks at a time until, ultimately, the threat actor's campaign comes to an end. An example of a measure for this second level would be purchase of additional DDoS protection in preparation of a DDoS by a hacker collective. Tier three is where decisions start to become more long-term. Identifying ahead of time what software the adversary will develop malware and exploits against can lead to a variety of decisions. These decisions range from deciding to prioritize certain patches, discontinuing use of a piece of software, purchasing or developing software, and segregating certain computers from the rest of the network. Decisions based on fourth-tier intelligence are likewise strategic and could lead to decisions on what types of cyber threat intelligence to consume or where an organization places strategic investments in cybersecurity over the long term. Although decisions based on third- and fourth-tier intelligence do not provide short-term gains, they are certainly more cost effective in the long run. Making more strategic level decisions based on third- and fourth-tier intelligence can potentially obviate a large volume of short term adjustments based on the lower tiers.

Figure 1 | Tiers of cyber threat intelligence

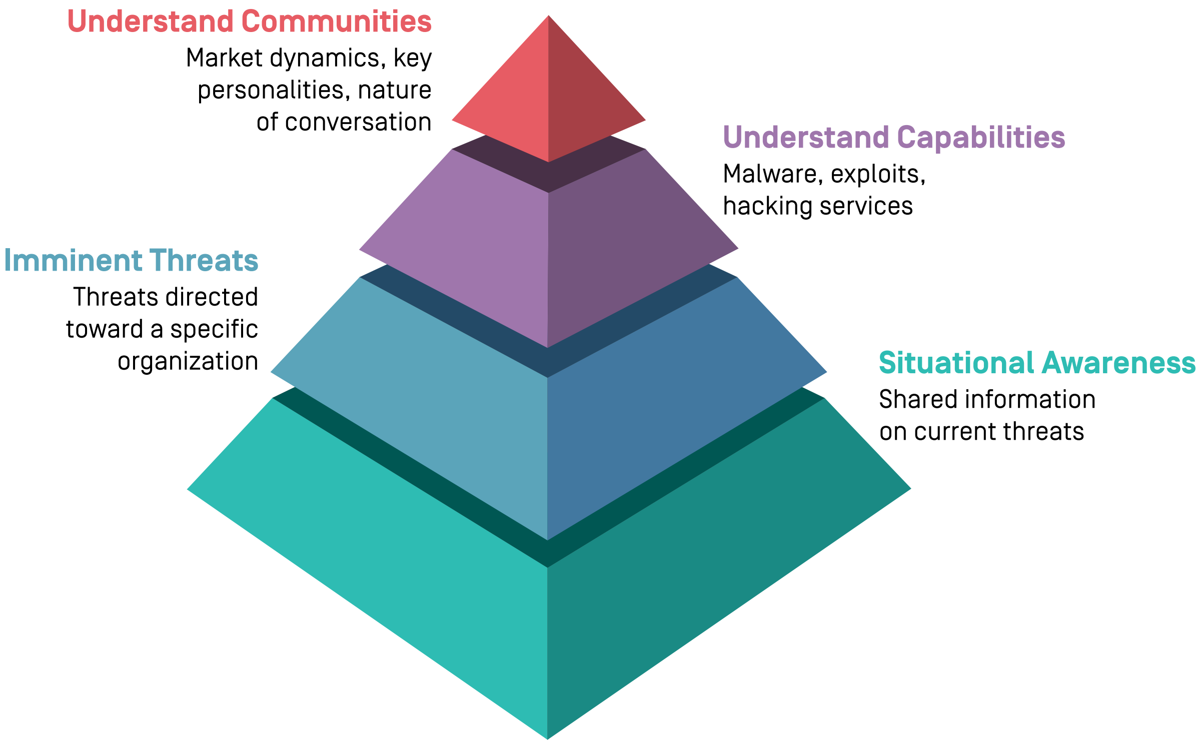
Understand Communities
Market dynamics, key personalities, nature of conversation
Understand Capabilities
Malware, exploits, hacking services
Imminent Threats
Threats directed toward a specific organization
Situational Awareness
Shared information on current threats

Moving toward the lower tier of this paradigm, we are faced with alerts that necessitate more immediate action. In fact, the information is so imminent at these tiers that they often can lead network defenders to identify systems that have *already* been breached. However, the expense incurred by an organization as the result of a breach only increases with time. For an organization to take action to stop, detect, or mitigate the effects of a breach, actions must be taken swiftly as the costs only mount with time. However, even pre-breach actions can be costly—for instance a power utility patching a critical vulnerability in an industrial control system could potentially lead to outages—hence directly affecting customers. In the future, as we become dependent on technologies such as the Internet of Things (IoT) and autonomous vehicles, breaches will be able to directly affect the everyday safety and well-being of large populations.

At the lower end of the paradigm, information sharing is also of high importance. To this end, there exist various Information Sharing and Analysis Centers (ISAC) exist for various verticals. The various ISACs provide a natural trusted community to share such information, and in some cases information relating to cybercrime is already shared in this manner. There also exist organizations for public-private sharing. For instance, the non-profit Arizona Cyber Threat Response Alliance (ACTRA) facilitates sharing of law enforcement data with cybersecurity professionals working in the critical infrastructure sector—and similar organizations are copying the ACTRA model in other states. However, in any sharing scheme has trade-offs. One of ACTRA's strengths is that itself is not a government organization—and hence can gain trust by industry members with relative ease. However, for the same reason, ACTRA will never have the same level of clout within the government as an organization like the National Cybersecurity and Communications Integration Center (NCCIC). Perhaps the way forward is a mix—for example ACTRA has a bi-directional sharing agreement with NCCIC—which is primarily used for the sharing of threat assessments and threat advisories.

While enhanced sharing of information is certainly important, the reactive nature of this strategy suggests that it is not a comprehensive solution. Hackers who specialize in finding exploits and building malware platforms continue to improve their craft, especially with regards to how stealthily their malware infects and operates. For example, a study from Symantec found that, on average, zero-days exist "in the wild" for over 300 days before identification.[5] Likewise, in 2016, malware platforms were known to persist on a target system for a median of 146 days before discovery.[6] Hence, we include more proactive forms of cyber threat intelligence as well in the paradigm.

> **In the future, as we become dependent on technologies such as the Internet of Things (IoT) and autonomous vehicles, breaches will be able to directly affect the everyday safety and well-being of large populations.**

The upper tiers of the classification system require less urgent action. For instance, identifying a new breed of malware that just was put for sale on the darkweb likely indicates that few (if any) hackers have employed it yet. Further, the decisions made based on such information are often decisions that must be made in the first place. Prioritizing patches is already common practice in every company with any kind of security focus—adding an element of what the adversary is focused on in such a prioritization can allow certain attacks to be obviated. Purchase of software is another such decision. In addition to normal price, compatibility, and usability criteria, acquisition personnel can also consider where the threat is focused and avoid purchasing software for which bad actors are currently developing exploits.

While decisions made from the upper tiers of cyber threat intelligence can be highly useful and lead an organization to entirely avoid certain attacks, it is also the most expensive. Simply put, advanced information on what malicious hackers are up to is difficult to come by. Further, the communities of these actors is growing. As a result, firms specializing in these have relatively large requirements for employees with highly specialized skills such intelligence analysis, linguistic ability, and counter-intelligence—in addition to general cybersecurity know-how. The expensive prices have largely kept this type of information with large firms whose market cap puts them toward the top of the Fortune 500. Well outside of this range are a variety of very important companies such as utilities, infrastructure, and logistics. Breaches against such firms, whose compromise has the ability to affect larger organizations can and do occur.

UNDERSTANDING PROACTIVE CYBER THREAT INTELLIGENCE

Now that we have established some basic concepts of cyber threat intelligence, we shall provide some example use-cases of how such information can be effective. We shall examine one of the source of intelligence that would be considered at the "upper layer" of the model of the last section: information gained on malicious hacker communities from the darkweb—portions of the Internet accessible only though certain secure protocols such as Tor or i2p.

Let's start with a quick example. In February 2015, Microsoft identified a vulnerability in the Windows operating system. At the time Microsoft disclosed the flaw, it was unknown if the vulnerability could be exploited. In other words while the flaw existed, it was questionable if it could actually be used by a piece of malware to compromise a system. This is interesting to note as there were over 15,000 vulnerabilities released in 2015.[7] Hence, in practice, enterprise network defenders must prioritize which vulnerabilities they patch. In April, we observed a seller on the darkweb advertising an exploit for about $10,000.[8] There was still no word from the security community until July when researchers from FireEye identified a variant of the Dyre banking Trojan that leveraged the exploit in the wild.

Now this is not to say that FireEye was late to the game. They may have very well found one of the first uses of the exploit. In fact, it is very likely that the exploit on sale on the darkweb in fact preceded any widespread use. However, it seems in this case that the malicious hackers were working on figuring out the exploit at the same time as "white hat" security professionals. However, there is an important distinction: while the white hats would likely have moved on, knowing the vulnerability was able to be patched, the malicious hackers would continue to focus on it—as they realize that a difficult-to-exploit vulnerability may go unpatched due to prioritization.

Clearly, this type of information about the goings-on of malicious hacker communities can be informative. However, it is difficult to obtain and, as a result, can often be expensive and difficult to share. Further, the explosive growth of Tor and related services will only exasperate this problem in the future. The number of sites on Tor has more than doubled in 2016.[9]

Further, the information on the darkweb represents just one source of information. We highlight it here as it provides a concrete example of how the information is useful. In understanding the ecosystem of cyberattackers, information across all the layers of our model are necessary. For instance, not every exploit developed will be sold on a darkweb market. An exploit developed by a nation-state may stay secret for years. On the other end of the spectrum, we often see groups such as Anonymous use louder channels to announce their intentions, as evidenced by their use of social media to recruit individuals to aide in denial of service attacks. It is generally regarded that a solid cyber threat intelligence program should consider multiple sources[10] —and this mirrors best-practices in traditional forms of intelligence. However, the resources and manpower involved in collecting, fusing, analyzing, and sharing such information generally precludes companies with a market capitalization below $1B from setting up a serious cyber threat intelligence program. With the interconnected nature of business today, this creates vulnerabilities as these firm include critical infrastructure and have also been habitually used by malicious hackers as a “launching pad” when planning operations against a larger target.

UNDERSTANDING STAKEHOLDER INTERESTS

In order to encourage the use of cyber threat intelligence to promote better security amongst mid-tier companies, we must recognize the stakeholders, their level of understanding, and what their interests are.

Government

The government has a general interest in ensuring the security of its citizens and holding businesses accountable to have products and services used by the citizens. With regard to cybersecurity, critical questions in this regard include: “How safe is the customers' data?" "Is this technology product secure?" and "What risk transfer mechanisms (i.e. insurance) are in place?" However, at the same time, the government also has a responsibility to grow the economy. Hence, unreasonable restrictions against corporations can have adverse effects: costs can soar, innovation can be stifled, etc. For example, it may not be reasonable to expect a company with a $100M market cap to be immune to attacks from top hackers working for foreign governments, so, as with many issues, the government must work to strike a balance between establishing parameters to drive certain behaviors as well an incentives to encourage other behaviors.

Large Market Cap Companies

Large market cap companies have a significant interest in avoiding cyberattacks. As the popularity, size, and value of such firms make them targets for malicious hackers of all sorts, both the probability of experiencing frequent attacks and the consequences of successful attacks are both greater. However, even in such organizations where the ROI of solid cybersecurity is apparent, such companies also have competing interests. For example, the sharing of cyber threat information clearly benefits cybersecurity program, but may often prove difficult due to compliance and confidentiality issues. In a different way, corporate partnerships, acquisitions, and mergers can lead to interactions with entities with a lesser cybersecurity posture—effectively increasing the attack surface to malicious hackers. However, all too often, cybersecurity issues are overlooked in this context for the sake of short-term efficiency.

Mid and Small Market Cap Companies

In our interviews with CSOs from mid-tier companies, there is a broad recognition of the importance of both cybersecurity and threat intelligence. However, the resources within such companies dedicated to cybersecurity is small—as is the budget for such activities. Further, as the visibility and size of these companies make them less likely targets, the perceived threat is often lower—which in turn means the perceived ROI for cybersecurity expenses is also lower. CTI's purpose for these companies should be to help CSOs and other C-level executives understand this reality, thereby encouraging small and medium market cap companies to invest more efficiently.

Cyber Threat Intelligence Vendors

The cyber threat intelligence industry is growing and expected to be valued at $5.5B by 2020. The technology and analysis required for entry into such business endeavors necessitate people with highly specific combinations of skills. For instance, individuals well-versed in both cybersecurity and machine learning or both linguistic ability and hacker culture are rare yet highly valuable to such endeavors, but can come at great expense. Likewise, early adopters of such intelligence feeds and technology tend to be larger market-cap companies, as they are best positioned to understand the value proposition of such firms. However, such high prices may not be tenable in the long term as they may hinder such firms from entering a broader market (i.e. mid and small market cap companies).

POLICY RECOMMENDATIONS

In proposing policies, we focus on two things: demonstrating value and aligning interests. By better demonstrating value of cyber threat intelligence technologies and services, more companies will be likely to adopt the technologies—thereby allowing the avoidance of more cyberattacks. By aligning interests, we can identify ways to reduce the cost burden of expensive cyber threat intelligence offerings for those who are least able to afford it. First, we introduce ways to set parameters to encourage investment in cyber threat intelligence.

Add requirements to the Federal Acquisition Regulation (FAR) for threat intelligence. Currently, the FAR requires 15 basic cybersecurity measures. However, all of them deal with an introspective look to the company's own IT infrastructure and do not pose even simple requirements dealing with threat such as exchanging threat information (even external to government sharing) and maintaining current indicators of compromise (which to some extent can be enabled by open sources)—not to mention the more proactive sources of intelligence discussed in this paper. As cyber threats are the direct result of actions by an active adversary and as there exist current solutions to give some insight into an adversary behavior, it makes a great deal of sense to integrate this into the FAR. Similar actions can be taken for other standards such as HIPAA.

Accreditation standards for medium and small market cap companies. Accreditation standards for medium and small market cap companies can be based on compliance standards would serve as a way for larger companies to identify smaller firms who have solid cybersecurity practices. This would better align the incentives of medium and small companies for improved cybersecurity posture with the desire to do business with larger firms. Such checks on accreditation by larger firms would also serve as a mechanism to better handle cybersecurity liability—as it the requirement for certification would fall to the smaller firms.

However, in addition to setting parameters, policy to induce companies to better embrace proactive cyber threat intelligence can also include incentives. We list some of these incentives as follows.

Increased government investment in research showing the value of cyber threat intelligence. The concept here is quite simple. Government R&D efforts that demonstrate the value provided by threat intelligence technology will give independent credence to claims made by threat intelligence providers adopting or licensing such technologies. Use of technology from a government funded and evaluated program will carry weight and allow for companies of all sizes to critically evaluate various marketing claims, which tends to create a hype cycle around many cybersecurity offerings. A current example of this is IARAPA's CAUSE program. In this program, various cyber threat intelligence technologies will be evaluated for their ability to predict real-world cyberattacks[11]. As the information on those attacks is provided by actual cyber defenders, the results from this program could be compelling.

Business mentorship programs. The U.S. government has in the past promoted business mentorship of smaller companies by larger ones through mechanisms such as the SBIR/STTR programs. A similar vehicle could be created for cyber threat intelligence whereby larger firms who have the in-house expertise to understand such capabilities can best be employed. This could greatly help in demonstrating the value of such technologies to the smaller firms as well as reducing their manpower requirements, as the learning curve and time-to-employment will be reduced for such technology.

THE ENEMY HAS A VOICE

The policy recommendations provided in this paper are designed to speed the adoption of cyber threat intelligence technologies. We discussed how these technologies and services can alter the playing field of cybersecurity by allowing the defenders to obtain a better perspective of the threat. We examined various different types of threat intelligence and explored how an exemplar—darkweb-based intelligence—can be used to benefit enterprise network defense. However, the expenses associated with such technology and services makes adoption difficult for smaller firms, which, due to the interconnected nature of technology, is harmful to society as a whole. We assert that through carefully-crafted policies demonstrating the value of such technology and aligning the interests of key players, this technology can become more widely adopted, thereby leading significant progress in cybersecurirty—one that considers the voice of the enemy.

Notes

[1] Tao, B., & Grimm, A. (2016, Sep.). South China Sea Conflicts Spills into the Cyber Domain: China vs. Vietnam. www.cyr3con.com/blog

[2] William J. Lynn, “Defending a New Domain: The Pentagon’s Cyberstrategy,”Foreign Affairs 89, no. 5 [2008].

[3] Andrea Locatelli, “The Offense/Defense Imbalance in Cyberspace” ISPI, no. 203 [2013]. http://www.ispionline.it/sites/default/files/pubblicazioni/analysis_203_2013.pdf.

[4] Perloth, N. (2015, Sep. 13). “Intelligence Start-Up Goes Behind Enemy Lines to Get Ahead of Hackers.” New York Times.

[5] Leyla Bilge and Tudor Dumitras,“Before We Knew It: An Empirical Study of Zero-Day Attacks In The Real World,” CCS [2012].

[6] Mandiant Consulting, M-TRENDS [Milpitas, CA: Mandiant, 2016].

[7] See statistics provided by the firm Risk Based Security which specializes in cataloging software vulnerabilities, https://vulndb.cyberriskanalytics.com/#statistics

[8] J. Robertson, A. Diab, E. Marin, E. Nunes, V. Paliath, J. Shakarian, P. Shakarian, Darkweb Cyber Threat Intelligence Mining, Cambridge University Press, 2017.

[9] Tor Metrics. https://metrics.torproject.org/

[10] R. McMillan, K. Pratap, “Market Guide for Security Threat Intelligence Services,” Gartner, 2014.

[11] “Cyber-attack Automated Unconventional Sensor Environment (CAUSE)” IARPA, Office of the Director of National Intelligence. https://www.iarpa.gov/index.php/research-programs/cause

This report carries a Creative Commons Attribution 4.0 International license, which permits re-use of New America content when proper attribution is provided. This means you are free to share and adapt New America's work, or include our content in derivative works, under the following conditions:

- **Attribution.** You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

For the full legal code of this Creative Commons license, please visit **creativecommons.org**.

If you have any questions about citing or reusing New America content, please visit **www.newamerica.org**.

All photos in this report are supplied by, and licensed to, **shutterstock.com** unless otherwise stated. Photos from federal government sources are used under section 105 of the Copyright Act.