CSIA RFI Responses, January 15, 2019

____

Update to the 2016 Federal Cybersecurity Research and Development Strategic Plan RFI Responses

DISCLAIMER: The RFI public responses received and posted do not represent the views and/or opinions of the U.S. Government, NSTC Subcommittee on Networking and Information Technology Research and Development (NITRD), NITRD National Coordination Office, and/or any other Federal agencies and/or government entities. We bear no responsibility for the accuracy, legality or content of all external links included in this document.

Request for Information Response: Update to the 2016 Federal Cybersecurity Research and Development Strategic Plan

This document was produced by Darktrace in response to the Request for Information published by the National Coordination Office (NCO) for Networking and Information Technology Research and Development (NITRD). It is intended to provide input for the 2019 update to the 2016 Federal Cybersecurity Research and Development Strategic Plan.

Introduction

The prefatory letter to the 2016 Federal Cybersecurity Research and Development Strategic Plan (the "plan") expresses a deep concern shared by leaders across both the public and private sectors. In the letter, John Holdren, then-Director of the Office of Science and Technology Policy, warns that, "While the United States is increasingly dependent upon cyberspace, cybersecurity has not kept pace with the increase in cyber threats." In the two and a half years since Holdren's remarks, his fears have proven highly prescient, as global attacks like WannaCry and Petya/NotPetya in 2017 demonstrated that even the most sophisticated companies are largely ill-prepared to counter an advanced cyber-threat.

However, the inability of traditional security tools to keep pace with today's evolving threat landscape is not merely a limitation of these particular tools; rather, it is a problem inherent to the fundamental approach upon which they are predicated. Because this approach relies on predefining threatening behavior based on past attacks, it remains, by its very nature, vulnerable to innovative attacks designed to behave unlike anything observed previously. More elusive still are insider threats, carried out by credentialed users who are familiar with the systems they manipulate and the information they exfiltrate. Governments are seriously imperiled by such insiders, whose efforts to leak state secrets and commit espionage frequently go undetected.

But while cyber security has become a yet greater challenge since the plan was published in 2016, the intervening years have also witnessed a long-overdue shift in the approach that many organizations employ to defend their networks. This approach, which leverages artificial intelligence to find subtle differences between typical and atypical behaviors, does not rely on rules and signatures to predefine what tomorrow's attack will look like. Instead, these AI security tools learn on the job to identify both the insider threats and the never-before-seen attacks that conventional tools miss, a capability which will ultimately allow them to realize the plan's vision of "a future where most [threat] detection will be done automatically at machine-to-machine speed." The U.S. government now has the opportunity to safeguard itself and its citizens by embracing this future, which entails promoting the development and adoption of autonomous cyber defenses.

Question Responses

Question 1: *What innovative, transformational technologies have the potential to greatly enhance the security, reliability, resiliency, and trustworthiness of the digital infrastructure, and to protect consumer privacy?*

Artificial intelligence and the Internet are the two defining features of the Information Age, and so it is little surprise that AI's application to the Internet has spawned countless transformative innovations. Fueled by the unprecedented wealth of data transmitted online, machine learning algorithms can assist and even surpass human professionals in performing increasingly complex tasks, such as countering sophisticated cyber-threats. Indeed, few applications of machine learning are currently as mature as cyber AI security tools, which intelligently detect these advanced threats by learning to differentiate between normal and abnormal behavior for each user and device in the particular networks they safeguard, drawing connections that incident responders miss. These AI defenses have come to protect companies across all industries and governments at every level.

This rapid and global transition to cyber AI has been catalyzed by the insufficiency of traditional security tools in the face of modern attacks, as evidenced by an ever-expanding list of major security breaches. Criminals are now launching never-before-seen exploits on a daily basis to bypass these traditional tools, whose reliance on past attacks to anticipate future ones prevents them from adapting as the cyber-threat landscape evolves. Further complicating matters is the fact that today's threats often originate within the network. In fact, almost three-quarters of reported cyber security incidents are attributed to insiders, while the best conventional security systems struggle to perceive the subtle behavioral anomalies that indicate when such insider threats are in progress. Conversely, the cyber AI approach consistently catches stealthy insiders, and its widespread adoption could deter them from even attempting to attack by rendering the risks prohibitive.

Beyond discerning both innovative external attacks and insider threats, AI-equipped cyber defenses can now autonomously respond to the recent wave of fast-acting attacks targeting nations and businesses around the world. These attacks, including ransomware strains that complete encryption in under a minute, exploit the limitations of human incident responders: even when a machine-speed threat is detected, reacting before it's too late is a separate and difficult challenge. But in 2017, the first enterprise-grade autonomous response technology entered the marketplace, which contains fast-moving cyber-attacks within two seconds. By intelligently determining a proportionate intervention whenever a threat arises, this technology affords security teams the critical time needed to investigate and take action.

Despite profound innovations in the cyber security space, there is no silver bullet to stop motivated attackers from infiltrating their targeted networks, meaning that resiliency following inevitable network infections is imperative. However, for many organizations, the operational disruptions that occur in the aftermath of a breach are its most devastating consequence, from e-commerce sites that lose customers and sales to transportation services that shut down during rush hour. Herein lies the value of autonomous response tools: they allow for normal business operations to continue by surgically restricting just the threatening activity. This application of machine learning directly fulfills the 2016 plan's appeal for "resilient" digital infrastructures wherein "malicious activities result in minimal disruption of operations and minimal financial or security losses." At a time when mechanized cyber-threats strike at an increasingly breakneck pace, such resiliency can only be achieved with the help of equally fast-acting AI tools.

In addition to their direct utilizations for cyber security, AI technologies are shaping myriad other facets of the digital infrastructure, a shift which will ultimately benefit overall cyber hygiene. For instance, machine learning programs excel at spotting errors and irregularities within databases too vast or technical for humans to readily verify, a capability that is exponentially improving the trustworthiness of online services. Along similar lines, machine learning systems will help businesses and governments identify failing hardware, streamlining their digital processes while strengthening their security postures. Even cyber AI security tools themselves possess broader functionality than just threat detection and autonomous response. These tools furnish organizations with a comprehensive understanding of their networks, potentially enabling them to flag disgruntled employees, ensure compliance with internal online policies, or create a real-time cyber-risk profiles unique to their specific infrastructure.

With its diverse array of applications, machine learning is empowering human users to perform complex jobs with greater reliability and efficiency. Above all, these technologies share an intrinsic adaptiveness wholly unlike software that runs on inflexible programming. In 2016, the authors of the plan envisioned digital systems that "dynamically adapt to changing threats," and in 2019, that updated plan should reflect that machine learning has made such adaptable systems a reality.

***Question 5:** What changes to cybersecurity education and workforce development, at all levels of education, should be considered to prepare students, faculty, and the workforce in the next decade for emerging cybersecurity challenges, such as the implications of artificial intelligence, quantum computing, and the Internet of Things on cybersecurity?*

Artificial Intelligence

Today's rapidly shifting threat landscape precludes even industry experts from predicting tomorrow's specific security challenges. However, given the sophistication of present-day attacks, it is safe to assume that these future challenges must be overcome by cyber security tools that can learn on the job to keep pace with evolving threats. Such advanced tools, in turn, necessitate an educated workforce with the skills to build them.

The continued application of machine learning to cyber security places data science firmly at the heart of any future-oriented education in the field. Across all industries, the past decade has witnessed the universal recognition of the value of big data, the unchecked collection of data at every turn, and a subsequent realization that human beings lack the capacity to process this amount of data in meaningful ways. But unsupervised machine learning systems, by contrast, become increasingly useful with increasing information, information which they harness to find subtle connections between data points that their programmers did not anticipate. The highly subtle nature of modern cyber-attacks can therefore be best neutralized by these machine learning tools, and by the data scientists who develop them.

As automation dramatically shapes the job market, both data scientists and incident responders are now in higher demand than ever before. In fact, a 2016 analysis from ISACA estimated a global shortage of 2 million cyber security professionals by 2019. AI security systems have the potential to mitigate this shortage by automatically detecting and prioritizing anomalous behavior, allowing overworked teams to instantly triage emerging threats across large networks. To perfect similar AI solutions, data science should become a more prominent focus, both in the classroom and in the SOC.

The Internet of Things

As the Internet of Things permeates an ever-growing number of physical appliances, protecting these appliances in the digital realm is becoming ever more essential. And yet, referencing the 2014 NSTAC projection that a staggering 26 to 50 billion IoT devices would be deployed by 2020, the plan concludes that "cybersecurity will be a daunting challenge at this unprecedented scale." Moreover, the plan speculates that "risks to the safety of IoT [devices], if unaddressed, will impede and block their adoption by society, preventing the full benefits of this technology from being realized for national priorities." In light of the transformative power of these devices, which tend to be designed with little to no emphasis on cyber security, declining to adopt them because of safety concerns would be a significant missed opportunity. On the other hand, continuing to buy IoT technologies in spite of their perils, as global consumers have largely done to this point, will jeopardize a majority of networks around the world.

From Internet-connected thermostats to IoT fish tanks, cyber-criminals have begun to target these vulnerable attack vectors to gain access to such networks. Indeed, endpoint security tools are rarely designed to cover IoT devices, whose integration with the physical environment means that many organizations are unaware of their existence. The solution to these obstacles is multifaceted. On a governmental level, the United States should establish basic security standards and frameworks for producers of IoT devices, in addition to working alongside tech companies to ensure that their ubiquitous IoT products do not compromise the American public's privacy or safety. Of course, given the fast-changing nature of cyber-threats and the limitations of endpoint tools in covering the IoT, organizations must secure their networks at large to secure their IoT devices, utilizing cyber AI defenses which neutralize attacks that have already infiltrated the perimeter.

Question 6: *What other research and development strategies, plans, or activities, domestic or in other countries, should inform the U.S. Federal cybersecurity R&D strategic plan?*

Earlier this year, the World Economic Forum named "cyber-attack" the number one human-caused threat to civilization, ahead of even war and terrorism. Yet in this era of coordinated cyber-crime and digital warfare between rival states, most governments still spend a minute fraction of their budgets on cyber defense compared to traditional forms of security. From the recent infiltration of the American power grid to the overseas threat actors who exfiltrated research from Western universities, modern attackers have proven that the Internet has no borders, while its potential for real-world harm has no obvious limits. Combatting such borderless, limitless threats requires an international, concerted response — one which cares little about the distinctions between domestic and global, public and private.

The aforementioned cyber AI technologies embody the successful integration of these public and private spheres. The foremost cyber AI system was created by a collaborative team of government intelligence experts, university mathematicians, and private sector specialists, whose complementary areas of expertise sparked the ongoing paradigm shift in cyber security. This kind of collaboration should be a primary characteristic of federal R&D in the field, both to militate against the damage of present cyber-attacks and to reap the rewards of a safer Internet. Such rewards, including increased consumer confidence, stronger intellectual property protections, and sturdier government infrastructure, will pay dividends for firms and states alike, incentivizing greater cooperation in the cyber security arena.

Indeed, there exists overwhelming consensus among industry experts that "the public and private sectors can both benefit from working together on cybersecurity initiatives," according to a 2017 report from the Penn Wharton Public Policy Initiative. In particular, the report argued that, because "the private sector controls much of the critical infrastructure that is vulnerable to cyberthreats," it has gained valuable insights into the nature of these threats and the risks they pose. The public sector, meanwhile, "is better positioned [than private companies] to investigate and prosecute cyber criminals" with its ability to "collect foreign intelligence, collaborate with other international agencies, and gain access to critical information regarding potential threats." Despite their obvious advantages, however, the effectiveness of public-private partnerships is currently hindered by the "complex regulatory and legal landscape surrounding cybersecurity."

In a fast-digitizing world, efforts to safeguard the United States in cyberspace must be streamlined if the nation is to retain its defensive advantage, which entails expediting the incorporation of pioneering cyber security technologies from the private sector. In particular, the U.S. government must prioritize the most promising of these technologies to date: cyber AI defenses. And whereas the legislative process moves deliberately by design, keeping pace with ever-changing cyber-threats compels the government to act swiftly and efficiently to facilitate a thriving digital economy.