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.

MP190034

January 15, 2019

Response of the MITRE Corporation to the NSTC RFI[1] on Cybersecurity R&D

For additional information about this response, please contact:

Duane Blackburn, S&T Policy Analyst
The MITRE Corporation
7596 Colshire Drive
McLean, VA 22102-7539

dblackburn@mitre.org
(434) 964-5023

[1] https://www.federalregister.gov/documents/2018/11/13/2018-24668/request-for-information-on-update-to-the-2016-federal-cybersecurity-research-and-development (Accessed January 15, 2019)

This page intentionally left blank.

Introduction

The MITRE Corporation is a not-for-profit company that works across government to tackle difficult problems that challenge the safety, stability, security, and well-being of our nation. We operate seven federally funded research and development centers (FFRDCs), as well as public-private partnerships, which gives us a unique vantage point across federal, state, and local governments, as well as industry and academia. MITRE works in the public interest to discover new possibilities, create unexpected opportunities, and lead by pioneering together for the public good to develop innovative ideas in areas such as artificial intelligence, intuitive data science, quantum information science, health informatics, policy and economic expertise, trustworthy autonomy, cyber threat sharing, and cyber resilience.

MITRE has longstanding and significant experience supporting multiple federal agencies on cybersecurity R&D and operations. For the last four years, we have had the honor of operating the nation's first cybersecurity FFRDC, the National Cybersecurity Center of Excellence, which fosters public-private collaboration to identify and solve today's real-world cybersecurity threats. We welcome the opportunity to respond to the National Science and Technology Council's (NSTC) Request for Information (RFI) supporting an update to the 2016 Federal Cybersecurity Research and Development Strategic Plan.

Because of our varied cybersecurity work, MITRE recognizes that there are some hard realities confronting the Government and the nation. These realities have evolved and, in some cases, become more urgent since the release of the 2016 Strategic Plan. We describe the most critical challenges below, as they shape our thinking in responding to the questions posed in the RFI.

First, the international threat environment has changed since 2016, and more countries are challenging U.S. national interests in cyberspace. Both peer and not-so-peer competitors are developing doctrines and operational concepts to employ offensive cybersecurity as instruments of power, integrated with their other national instrumentalities and resources. Technical threats have increased in sophistication, cyber "tradecraft" has become more intricate, and the barriers to entry for smaller powers to engage in malicious cyber activity have fallen. At the same time, our attack surface, represented by more connected infrastructures, has increased enormously – and continues to do so. In other words, the cybersecurity environment in which we operate has quickly become more challenging since 2016.

Second, a new cyber ecosystem is emerging around us, enabled by 5G Internet and IPv6. Developed to provide more speed and power, 5G will create a mostly seamless, low latency, direct connection between our connected devices and the global Internet. IPv6 will provide IP addresses for an unlimited world of connected devices, including those that power our nation's infrastructures. In turn, these infrastructures will be connected to, and interdependent on, each other. Connected cars will be powered by the smart grid. These cars will connect to smart roads, and the use of resources will be mediated by cloud-based analytics and artificial intelligence (AI). Much of the hardware on which we will depend in this new ecosystem will be built by foreign companies. We will need new security models and technologies to effectively manage our national security interests in global infrastructures built in this new ecosystem, that we did not build, do not own, and may not be able to govern.

Third, we must expand beyond the Federal Government-centric scope of the 2016 Strategic Plan. To truly make progress on this issue, we must create a *national* cybersecurity R&D community—comprising the Government, industry, academia, and non-profits—and leveraging resources beyond those owned

by the Government. The Government should build a national cybersecurity R&D community and governance structure that prioritizes problems and coordinates efforts throughout the broadest possible cybersecurity R&D community, explicitly recognizing the role and resources of both Government and non-government organizations to contribute to national cybersecurity R&D progress.

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?

Advances in computer science, cryptography, computer architecture, electromagnetics, data analytics, and other areas are all being applied to cyber problems. Unfortunately, cybersecurity technologies are available to both offensive and defensive users. We continue to race to apply new technologies defensively before they can be exploited offensively. In addition, cyber defenders must protect more systems than ever before, including embedded systems, low-cost Internet of Things devices, industrial control systems, medical equipment, vehicles, and autonomous systems. Much cyber research centers on adaptation of other technologies (e.g., mathematics and data analytics) to cyber problems and adapting conventional enterprise cyber solutions to new systems (e.g., industrial control systems).

AI is the most promising, and potentially overhyped, of the new technology areas. Broadly, AI in cyber relates to three areas: classification, planning, and generation.

- *Classification/Detection*: Research has shown that machine learning (ML) approaches can reason over large numbers of dimensions and can be very effective at classification and detection problems. Companies such as Cylance are already marketing ML-based antivirus products, and others are developing analytics and firewall capabilities. ML is also being studied for continuous behavioral authentication and for large-scale extraction of cyber threat intelligence.
- *Planning:* AI is also an enabler of autonomy and decision support. Advances in reinforcement learning and online planning are being applied by DeepMind, OpenAI, and others to play complex computer games such as DoTa2 and Go. These technologies have great potential application to cyber response automation, automated testing, network defensive planning, and dynamic network reconfiguration (particularly when combined with new networking technologies, such as software defined networking).
- *Generation:* Studying the weaknesses of AI (i.e., Adversarial AI) will help us to understand the limitations of the technology, as well as to design AI algorithms that are robust against adversarial attacks. Advances in Generative Adversarial Networks are being studied for synthetic data generation (e.g., binaries, network flow data, etc.), software diversity, and automated code generation and patching.

The nation needs to conduct more R&D at the intersection of AI and cybersecurity to protect large critical infrastructures. This requires large-scale platforms capable of representing the vast amounts of data to be integrated into AI-enabled decisions. For example, China's City Brain project (a collaboration between the City of Hangzhou and tech giant Alibaba) represents the scale of platform needed to learn

about the cybersecurity vulnerabilities of AI-enabled systems and the data on which they depend. Platforms of similar scale should be made available for cybersecurity research into how AI can be used to improve cybersecurity and how AI-enabled systems must be defended.

Another significant research area is developing tools supporting software and hardware assurance. The software formal methods community continues to increase the scale that can be formally analyzed for correctness (i.e., from thousands to millions of lines of code), while also addressing composability and usability challenges. These technologies enable development of larger and more secure systems as well as new types of analysis of existing software (e.g., concolic analysis and other techniques showcased in the 2016 DARPA Cyber Grand Challenge). Other advances in formal analysis are being applied to communications and cryptographic protocols and hardware specification.

New demonstrations of power, timing, electromagnetic, and other side channel attacks (e.g., Rowhammer[2], Spectre[3], Meltdown[4], etc.) will continue to expand the thinking of cyber researchers and system designers. In the year since the Spectre and Meltdown vulnerabilities were first reported, we have seen an explosion of demonstrations of vulnerabilities as other researchers examine microarchitecture-level flaws in modern microprocessors (including security-specific hardware features such as Intel's SGX containers). This understanding is enabling computer architects to characterize these attack vectors and develop defenses.

Cryptography has always been an enabler of cybersecurity. Advances in partially homomorphic encryption and secure multiparty computation may enable practical computation on encrypted data, reducing the risk of stored data as well as vulnerability to side channel attack. Blockchain technologies show promise for data integrity, supply chain protection, identity and access management, and other applications. The related computation contracts area (e.g., Ethereum) may enable more trusted distributed computation at scale. Special purpose cryptographic protocols, such as identity-based encryption and attribute-based encryption, may offer improved usability and enable improved sharing of information.

Finally, practical advancements in lightweight cryptography protocols may enable the application of security to resource-constrained devices, such as IoT, cyber physical systems, and embedded devices.

A variety of computing and information technologies are creating additional options for cyber defenders. Hardware and operating system security features (e.g., SGX, TrustZone, virtualization

---

[2] Yoongu Kim, Ross Daly, et al., "Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors," http://users.ece.cmu.edu/~yoonguk/papers/kim-isca14.pdf
Mark Seaborn; Thomas Dullien, "Exploiting the DRAM rowhammer bug to gain kernel privileges." googleprojectzero.blogspot.com (March 9, 2015). https://googleprojectzero.blogspot.com/2015/03/exploiting-dram-rowhammer-bug-to-gain.html

[3] Kocher, Paul; Genkin, Daniel; Gruss, Daniel; Haas, Werner; Hamburg, Mike; Lipp, Moritz; Mangard, Stefan; Prescher, Thomas; Schwarz, Michael; Yarom, Yuval (2018). "Spectre Attacks: Exploiting Speculative Execution" https://spectreattack.com/spectre.pdf

[4] Moritz Lipp, Michael Schwarz, et al., "Meltdown: Reading Kernel Memory from User Space," Usenix Security Conference 2018.

support at the hardware level, and trustlets and Isolated User Mode Processes[5], etc. at the operating system level) will enable system software and application designers to create more secure designs. Hardware advances, such as quantum computing, neuromorphic computing, and AI processors (e.g., Google's Tensor Processing Unit), may make large-scale cyber problems more tractable. Finally, networking technologies, including programmable data plane processors and support for software defined networking, may enable more advanced packet analysis, dynamic network partitioning and reconfiguration, and improved distributed analytics.

Question 2: What progress has been made against the goals of the 2016 Federal Cybersecurity R&D Strategic Plan? Are there mature private-sector solutions that address the deficiencies raised in the 2016 Strategic Plan? What areas of research or topics of the 2016 Strategic Plan no longer need to be prioritized for federally funded basic research?

What progress has been made against the goals of the 2016 Federal Cybersecurity R&D Strategic Plan?

Less progress has been made against the goals within the 2016 Strategic Plan than the R&D community may have expected. Based on MITREs knowledge from working with our sponsors and interactions with the private sector, we offer the following specific insights regarding progress in each of the five recommendations promoted in the 2016 R&D strategy:

Recommendation 1 – Prioritize basic and long-term research in Federal cybersecurity R&D

We have not seen evidence that such prioritization has significantly affected federal research plans, individually or collectively, since the 2016 Strategic Plan was published. While many Government agencies have developed cybersecurity strategies since 2016, they appear to have been developed in isolation rather than as part of a larger whole-of-government strategy. Such cross-agency collaboration and information exchange can be critical to maximizing return on investment (ROI) on Federal R&D.[6] It is possible that some prioritization has taken place within the allocation of resources and is reflected in budgets and budget guidance provided the agencies. However, changes in emphasis to Federal cyber R&D programs based on such budget and guidance have not been apparent.

Similarly, private sector R&D has focused on market needs. There is no known overarching strategy that enables efforts from both the private and public sectors to work together to address national needs. As such, we also cannot identify the critical investment gaps in the nation's overall AI R&D ecosystem.

[5] Microsoft, Isolated User Mode (IUM) Processes, https://docs.microsoft.com/en-us/windows/desktop/procthread/isolated-user-mode--ium--processes

[6] Note the ongoing ROI efforts supporting Cross-Agency Priority Goal 14 within the President's Management Agenda.

Recommendation 2 – Lower barriers and strengthen incentives for public and private organizations that would broaden participation in cybersecurity R&D

Public and private organization participation in cybersecurity R&D has broadened since the 2016 Strategic Plan was published but is driven by greater awareness of cybersecurity issues and potential ROI than by government-issued incentives. A bigger-picture concern is an awareness that we don't have enough cybersecurity professionals to properly perform the work[7], if new incentives are developed.

Recommendation 3 – Assess barriers and identify incentives that could accelerate the transition of evidence-validated effective and efficient cybersecurity research results into adopted technologies, especially for emerging technologies and threats

We are not aware of new incentives made broadly available to accelerate evidence-validated R&D into adopted technology put in place since the Strategic Plan was published. Evidence-validated R&D is difficult because of a lack of agreed on cybersecurity metrics. However, a positive note is the continued efforts of the National Institute of Standards and Technology's National Cybersecurity Center of Excellence (NCCoE) to develop and encourage the adoption of practical and affordable reference architectures applicable to many commercial "verticals." This approach promises to lower economic and other barriers to adoption of stronger cybersecurity profiles by U. S. enterprises.

Recommendation 4 – Expand the diversity of expertise in the cybersecurity research community

The research community has expanded through the introduction of more cybersecurity programs nationally. For example, new university programs have been established at the University of Tulsa, and Georgia Tech has expanded its research program in recent years. The Academic Affiliates Council, associated with the National Cybersecurity Center of Excellence, represents a mechanism by which more than a dozen academic institutions have been invited to participate in the development of NCCoE reference architectures that can be adopted nationwide. What's needed, however, is a national cybersecurity R&D community that makes known national R&D priorities, and knits together the efforts of research institutions. Such a community could do much to harness a more diverse workforce nationally.

Recommendation 5 – Expand diversity in the cybersecurity workplace

Expanding diversity in the cybersecurity workforce in the public and private sectors remains a significant challenge, particularly in light of the cybersecurity talent shortage the nation continues to experience. Expanding diversity could benefit from HUB-Zone-like programs that encourage the develop of cybersecurity businesses in economically disadvantaged urban and rural communities, which play a strong role in Federal cybersecurity programs. In addition, more support can be given to historically black colleges and universities to establish cybersecurity programs, building on efforts already underway at institutions such as Morgan State University.

[7] See https://www.herjavecgroup.com/wp-content/uploads/2018/07/HG-and-CV-The-Cybersecurity-Jobs-Report-2017.pdf.

Are there mature private-sector solutions that address the deficiencies raised in the 2016 Strategic Plan?

The cybersecurity industry is coalescing around the NIST CSF[8] and MITRE ATT&CK[9] to make cybersecurity more quantifiable and to more clearly organize the cybersecurity resources and processes available to an enterprise. Neither framework alone creates cybersecurity; however, the two frameworks are being used to help organizations design, enhance, and "measure"[10] cybersecurity program effectiveness.

What areas of research or topics of the 2016 Strategic Plan no longer need to be prioritized for federally funded basic research?

None of the issues identified in 2016 have been overcome, and all remain important. Relative positions within a ranking of priorities will continue to evolve as capabilities mature and new threats are identified.

Question 3: What areas of research or topics of the 2016 Strategic Plan should continue to be a priority for federally funded research and require continued Federal R&D investments?

and

Question 4: What challenges or objectives not included in the 2016 Strategic Plan should be strategic priorities for federally funded R&D in cybersecurity? Discuss what new capabilities would be desired, what objectives should guide such research, and why those capabilities and objectives should be strategic priorities.

The major research thrusts identified in the 2016 Strategic Plan cover topics that should be a priority for continued R&D investment. As mentioned in the *Introduction* of this paper however, MITRE feels that for the nation to truly make progress, we must view cybersecurity R&D from a *national* perspective. Updating a Federal-Government-only plan will not reduce duplication of efforts in the private-sector, or sponsor research on critical gaps.

Non-governmental cybersecurity R&D investments are significant[11], which must be understood to meet this national-level need. The Government should build a national cybersecurity R&D community and governance structure that allocates problems and coordinates efforts throughout the broadest possible

[8] https://www.nist.gov/cyberframework
[9] https://attack.mitre.org/
[10] We use the term "measure" in a broad sense, not a pure quantifiable manner, but somewhere in between quantitative and qualitative measurements.
[11] For example: Venture Capital firms invested $7.6 billion in cybersecurity companies in 2017, up from $3.8 billion in 2016. https://cdn2.hubspot.net/hubfs/2199909/Feats%20Q1%202018/VC%20Funding.pdf

cybersecurity R&D community, recognizing explicitly the role and resources of both Government and non-government organizations to contribute to national cybersecurity R&D progress.

MITRE's ongoing efforts at interfacing between Government and industry have already identified three such gaps (discussed below). A dedicated national-level collaborative effort would help identify additional gaps and enable development of a priority list.

Gap 1: Cyber Systems Engineering (CSE)

Decision makers charged with defending our nation's digital assets— whether from a small company or a major acquisition command—are struggling. They are tasked with allocating resources and funding to most effectively and efficiently improve their cybersecurity posture, but they lack the rigorous methods and tools to understand the complex interplay among their threat, their system, their mission, and possible mitigations. The objective of a CSE research thrust would be to provide the necessary framework/process/tools to assess how to best procure, design, and augment critical digital assets to achieve the maximum, quantitatively measured and verified, cybersecurity posture.

Also key to CSE is decision support, which includes complex cyber systems integration and a framework of appropriate actions that cyber professionals should take when their assets are under duress or attack. The cybersecurity technology market space is teeming with hardware and software tools that inform on intrusion detection activities (by signature or analyses) and offer proactive (or automated) intrusion prevention capabilities. However, these systems require intensive integration efforts that take into account operational impact, maintenance costs, and real return on value to the mission at hand. Additional aspects to be considered include the likelihood of a system taking action on non-events that waste time and resources and impact operations; events for which the response was not measured or matched to the threat, resulting in negative mission impact; or insider threats.

Even with a staff of professionals trained, these aspects are a challenge to manage in live operations when staff must understand how these challenges manifest themselves within dynamic computer and networked environments. The legal landscape is well engaged and framed by efforts such as the Tallinn Manual[12], which offers perhaps the most comprehensive analysis of how the law applies to cyber. The Tallinn Manual and other legal frameworks require continual engagement and expression across the cyber professional community, so that technical defense approaches, countermeasures, and evidentiary standards can become incorporated into decision-support guidance.

The threats are becoming dynamic, as well, and with the rise of zero-day attacks and advanced persistent threats in the cyber offense space, the reliance on tool integration is not enough. To execute defensive measurers quickly and resolutely, cyber professionals require strong decision-support options and authorities. Decision support actions should be derived from frameworks and tuned to specific missions, allowing staff to work with the reported results from technology tools and the inward flow of threat intelligence to lower signal-to-noise ratios, and then take appropriate to counter threats and attacks.

[12] https://www.law.georgetown.edu/international-law-journal/wp-content/uploads/sites/21/2018/05/48-3-The-Tallinn-Manual-2.0.pdf

Cybersecurity systems engineering must integrate these new capabilities into their designs, and cybersecurity operations should also include the integrated use of these capabilities.

Gap 2: Deception

As noted in the 2016 Strategic Plan, deterring an adversary is an important component of the future of cybersecurity. After all, the most successful defense is the one that causes an adversary to abandon the contemplated attack before its even launched. A critical technique to deter adversaries is cyber deception. Cyber deception can effectively increase an adversary's uncertainty, improve attribution, and enable better protection. Note that this research thrust would extend well beyond today's honeypot cyber deception, which is considered "deception for detection," and does not achieve any of the desired uncertainty, attribution, and protection effects. Broadly, there are two interconnected forms of cyber deception that warrant research: content deception and environmental deception.

Content deception employs tactics, technologies, and techniques to ensure that an adversary is unable to discern truth from fiction in any information exfiltrated from a victim system, thereby rendering that exfiltrated information useless. For example, suppose a defense contractor kept 10,000 slightly altered designs for a next-generation combat aircraft along with the real design. (AI could help create altered designs that are indistinguishable from the original.) Now suppose an adversary successfully executed a cyber-attack to steal the aircraft design and exfiltrated some or all the designs. Unable to discern the real design from the fake, that adversary would be left with useless information. Thus, content deception, especially if deployed routinely across the nation, has great potential to significantly deter would-be adversaries.

Environmental deception employs tactics, technologies, and techniques to create a replica digital system that is indistinguishable from the targeted digital system from an adversary's perspective. Once an adversary is placed into a deception environment, defenders can study the adversary's post-exploit behavior (e.g., how is privilege escalated, how is command and control established). This method is extremely powerful. It allows defenders to observe effective and ineffective defensive techniques, discover unexpected attack vectors, and improve attribution. Like content deception, environmental deception becomes exponentially more powerful when deployed nation-wide. Adversaries won't know if they're in a real system or in a deception environment, thereby increasing their risk and uncertainty.

Gap 3: Multiparty Computing

The 2016 Strategic Plan discusses "Cryptographic mechanisms to data" as a research priority that enables data operation and modification without decryption (homomorphic encryption). MITRE recommends adding a related effort to allow multiple parties to jointly perform computations across disparate data sets while keeping the underlying data sets private (so called, multiparty computing or MPC). The potential value of MPC can be illustrated when applied to the challenges of privacy-preserving data mining and secure intelligence sharing, as described below.

- *Privacy-Preserving Data Mining.* Consider several hospitals that wish to compute some statistics on a joint data set of patient records. The hospitals might be trying to compute statistics on the

success rate of a particular type of chemotherapy for breast cancer patients. Or, the hospitals might be trying to compute statistics related to a drug's side effects. Obtaining the right statistics is paramount in medical science research, but several laws (e.g., HIPAA) prevent hospitals from sharing valuable patient data with anyone. MPC has great potential to solve this problem. The hospitals would be able to simply run a protocol to compute the required statistics without revealing the patient records to each other, because MPC enables privacy-preserving data mining. And as data becomes more central to national priorities and daily activities, the importance of privacy-preserving data mining will go up significantly.

- *Secure Intelligence Sharing*. Suppose the intelligence agencies of two nations seek to pool information about spies from a third nation. Each agency is willing to share its valuable intelligence only if the other country has valuable information about one of its targets. This presents a catch 22: how does one agency ascertain whether the other has valuable information without either of them being the *first* to share? MPC would represent an attractive solution in this case, as well.

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?

Defining the cybersecurity workforce challenge

"Cybersecurity" is a term fraught with multiple definitions. Similar to the term "cloud," cybersecurity refers to a great many interrelated sub-disciplines, including: combined hardware and software systems (firewalls, IDS/IPS/SIEM, IoT), pure software systems (tools, coding practices, code review), networking (SW-defined networks and security, physical networks, wireless, routers/switches), defense (analysis, operations/monitoring, SNOCs), offense (red/blue/purple teaming), policy (development, application to organizations), audit, vulnerability research, product research/certification, and many others. Each of these could be studied and expressed in great detail, but more important is how they interrelate and impact the career trajectories and market of available cybersecurity professionals.

Because of the breadth of skill needed to address today's cybersecurity challenges, an adequate cybersecurity workforce must include many different types of professionals with skills in: networking, software and systems, auditing, policy (to include standards and guidance), monitoring and analysis, and hacking (dark web operators, botnet developers, script kiddies, vulnerability scanners, social engineers, and malicious administrators/insiders). Hackers, while not typically employed, educated, or trained as such, are the result of the available technology. They perceive how to benefit from taking advantage of the situation and usually develop and hone their own skills and determine their own technical "career" trajectory. Hackers can range from "white" (hacking for the "greater good"), "black" (operating in a criminal manner), or "gray" (good intentioned activities using questionable methods and ethics). Efforts to hire hackers into the Federal Government and industry are often attempted but can meet with

disastrous results[13]. It should be noted that technically minded cybersecurity professionals likely know some hackers or have or themselves experimented with hacking somewhere on the spectrum.

Some cybersecurity skill paths allow for practitioners to smoothly transition from one discipline path to another, and some do not. It is hard for professionals to find a clear developmental path for entering the cybersecurity field and easily transitioning into high-demand subdisciplines. To become credible in some of these fields requires massive amounts of personal time, talent, and investment. For example, an individual with a strong networking background across the LAN/WAN domains can find a great home in the cyber subdisciplines of firewalls, IDS/IPSs, and the like. This is because those technologies operate on the same principles the networker already knows. The same could be said for the software developer or for any other complex technical discipline. Cybersecurity, for them, is just branching out.

Many individuals who want to enter the cybersecurity field find that they do not have the requisite technical background to make an immediate impact. This often leads them to seek certifications (e.g., CISSP) in hopes that they can later adapt their professional career into cyber. Many technical training firms and higher-education schoolhouses offer courses that range from boot camps to full semesters of instruction that lead individuals to prep for and obtain certifications. Once Certified, these individuals are often ill-equipped to make the critical technical impact leading to sharp engineering decisions, as they are textbook trained and have not worked in production environments across technical teams. As entry-level trained and certified staff, they may find themselves in policy enforcement and systems-hardening roles, while others will break into the management field, leading technically-minded engineering teams.

This often leads organizations to a "systems" security culture, rather than an "assets" security culture. In the former, staff are hardening systems and devices but are not thinking holistically about the broad range of where threats, vulnerabilities, and attacks can occur. Such staff find that they are not well connected to the information and assets that the larger organization values but are instead focused on generic security applications. When an organization's executives are not engaged in a conversation about critical assets and information, as well as the appropriate level of investment and attention to protect those assets and information, then the default security posture will lack a true risk-managed perspective. While risk management and how it corresponds to asset protection exists in published industry and government standards and instructions, new junior and even seasoned mid-level staff might find that they are not engaged in a risk-management process. It would be useful to first consider the protection of what is most important (the assets themselves, whatever they are – physical or data) as well as true management and acceptance of risk. This must be driven from senior leadership down to the staff.

A national strategy for educating the cyber workforce pipeline would need to set objectives on:

- How we should protect assets
- How we handle defensive and offensive operations
- What groups most need the training

Research is required to consider which aspects of cyber require solutions at higher levels, and then work down from there. In a market-based economy, the Government has a role to play in stimulating growth,

[13] https://www.nytimes.com/2010/11/14/magazine/14Hacker-t.html (accessed January 14, 2019)

innovation, and ideas. We need to assess where policy and incentives can be engineered to generate the needed cyber workforce pipeline. Cyber demands creativity and resilience across all related disciplines.

Addressing the cybersecurity skills gap

With the right target disciplines in mind, there are several recommendations to consider for conducting research and closing the cyber skills gap: [14]

- Cybersecurity education needs to start earlier, with individuals properly understanding key concepts well before they graduate high school.[15]
- Offer new opportunities to provide access to qualified candidates (e.g., changing degree requirements for new hires), and focus job postings on core requirements.
- Apply NIST's National Initiative for Cybersecurity Education (NICE) Cybersecurity Workforce Framework[16] where it applies to cyber roles.
- Investigate novel hiring and training practices, including looking at prioritized technical and professional skill paths, and where cyber can be included.
- Consider apprenticeship programs and commit to employee development.
- Seek productive partnerships and programs to maximize impact.
- Broaden the understanding and application of cyber at all levels and expand educational opportunities across the board.
- Boost the CyberCorps Scholarship for Service program[17] to meet specific R&D goals and initiatives where gaps exist, and work with academia to help incentivize and shape their course offerings to prepare the next generation of cyber professionals.

Quantum is a game changer

A pending technological breakthrough in quantum computing would herald a new age in which computers are capable of exponentially higher computational power than they are today. Because of techniques used in modern cryptography, a highly functional quantum computer could instantly render cybersecurity defenses ineffective. Worse, the United States is not in the lead for delivering on this game-changing technology; countries such as China [18] are making heavy investments and progress in the technology.

[14] https://assets.aspeninstitute.org/content/uploads/2018/11/Aspen-Cybersecurity-Group-Principles-for-Growing-and-Sustaining-the-Nations-Cybersecurity-Workforce-1.pdf?_ga=2.73996562.1859401697.1543460706-160395908.1543460706 (accessed January 14, 2019)
[15] https://www.washingtonpost.com/pb/news/powerpost/paloma/the-cybersecurity-202/2019/01/11/the-cybersecurity-202-high-schoolers-must-start-training-for-security-jobs-to-fill-the-talent-gap-professors-say (accessed January 14, 2019)
[16] https://www.nist.gov/itl/applied-cybersecurity/nice/resources/nice-cybersecurity-workforce-framework (accessed January 14, 2019)
[17] https://www.sfs.opm.gov/
[18] https://www.nytimes.com/2018/12/03/technology/quantum-encryption.html (accessed January 14, 2019)

There is research being conducted, some leveraging quantum computing for cyber defense, that could mitigate this risk. Updating Government systems and infrastructure to deal with the cybersecurity risk in a quantum-enabled world will take as long as 10 to 20 years[19] to complete. It is therefore of critical national importance to research, refine, and implement these changes.

The pipeline for quantum physicists is insufficient to meet the nation's need to meet this challenge. Much of this need can be addressed with advanced engineering. Therefore, it is recommended that the United States advance the science of quantum engineering, which will enable a workforce capable of delivering on quantum cyber techniques to shore up commercial and government critical systems to prepare for the coming digital encryption environment.

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?

Few countries have national cybersecurity R&D plans from which the U.S. plan can benefit. The 2011 U.K. Cyber Security Strategy calls on the Government to work with industry to undertake "research and development to keep producing innovative solutions," but does not describe the mechanisms for doing so. Perhaps the best approach would be to explore the lessons learned by the United States itself in the post-WWII era to build national R&D strategies for nuclear energy and aerospace technology. These strategies created whole-of-nation R&D communities capable of propelling the United States into a leading position for both technologies.

The private sector is driving evolution of the global information technology infrastructure by conducting important R&D, engineering, and product development. The private sector is applying AI to a wide range of applications, from health care to transportation, and from energy production to retail. New cloud architectures and infrastructures, many of which are larger in scope and scale than those deployed by the public sector, are being deployed worldwide.

The adoption of IPv6 will provide almost limitless opportunities for the interconnection of devices resident in every infrastructure; 5G Internet will create high-speed, direct interconnectivity without intervening networks, and will result in global infrastructures that pose new security challenges that traditional approaches cannot solve. Cloud-based analytics are being used already to optimize business infrastructures; cloud-based information technology infrastructures will be used to support critical infrastructures, while AI will be used to help manage those infrastructures. Indeed, AI is already being used by businesses (retail, transportation, and supply chain management) in ways that affect the nation's transportation and energy infrastructures, even as we await the direct application of AI to the management of those infrastructures.

[19] https://www.washingtonpost.com/news/innovations/wp/2015/05/11/quantum-computing-is-about-to-overturn-cybersecuritys-balance-of-power/?noredirect=on&utm_term=.5f1f04b68a54 (accessed January 14, 2019)

As a result of all these changes, the Federal Government should be linked more closely to the R&D activities driving the development of the larger IT domain, rather that awaiting the development of new IT technologies and then defining approaches to secure them,

The U.S. Federal R&D community should work in concert with international partners, including, the NATO cyber R&D community, which already works closely with the Army, Navy, and Air Force Research Laboratories and provides access to important cybersecurity R&D resources throughout much of Europe. Working with, and perhaps through, these Laboratories, the Federal R&D community can synchronize, and possibly modulate, the priorities and content of the NATO partners' R&D activities. The United States could also use the knowledge of those partners to help shape Federal R&D programs, particularly in cases in which foreign progress is already promising.

In addition, the Federal R&D community should coordinate more closely with key academic partners. As an example, MITRE manages an Academic Affiliates Council with which it coordinates R&D efforts and can offer this as a possible model for adoption by the Federal R&D community.

The Government should also collaborate with key industrial partners, particularly those with independent R&D budgets, to make progress more quickly. Coordination with industry, for example, will help Government influence the development of commercial products that are in use throughout the nation's critical infrastructure—the bulk of which is owned and operated by the private sector.

For the best possible outcome, the Government should work with academia, industry, and non-profit partners to build a whole-of-nation cybersecurity R&D community, charged with implementing a national cybersecurity R&D strategy. The participants would work together to set priorities and assign activities that address every sector of U.S. national security, the functions of a sovereign government, and our national economy.