Testimony of
Robert K. Knake
Senior Research Scientist, Global Resilience Institute
Northeastern University
&
Senior Fellow for Cybersecurity Policy
The Council on Foreign Relations

Before

The U.S. House of Representatives
Subcommittee on Cybersecurity and Infrastructure Protection
of the Committee on Homeland Security

"Preparing for the Future: An Assessment of Emerging Cyber Threats"

Tuesday, October 22, 2019
2:00pm

Room 310
Cannon Office Building

Introduction

Thank you Chairman Richmond, Ranking Member Katko, and members of the Committee for the opportunity to testify on this important matter. While other witnesses will focus on how the capabilities of specific threat actors may change and evolve, I would like to focus my remarks on how the technology landscape may change in the next five years and what that may mean for emerging cyber threats. Before I begin, let me be clear that the views I represent here are my own and do not represent my employers or any supporters of my work.

Looking back over the past decade, there are reasons to be hopeful for a secure cyber future. When my co-author Richard Clarke and I wrote *Cyber War: The Next Threat to National Security and What to Do About It* a decade ago, we predicted a dire future in cyberspace. Early trends then indicated to us that our adversaries would develop sophisticated cyber offensive capabilities and would use these capabilities to undermine our dominance of conventional military domains. We predicted correctly that North Korea would emerge, somewhat surprisingly, as a capable adversary in the cyber domain and highlighted China's ongoing campaign of economic espionage on behalf of its national champion companies. We of course failed to predict many of the key events that are top of mind today like Russia's use of the Internet to interfere in elections and sow dissent; however, in my view, our greatest error was our failure to see the technology trends that have allowed the defensive community to be able to manage the threat posed by even the most determined nation state adversaries.

In *Cyber War*, we concluded that private companies could not defend themselves against determined adversaries because cyberspace as a domain favors the attacker. Conventional wisdom at the time was that an attacker had all the advantages. An attacker only needed to find one vulnerable system to succeed whereas the Chief Information Security Officer (CISO) at a large enterprise had to defend thousands or hundreds of thousands of systems. This asymmetry was often captured as the idea that "the attacker only needs to compromise one vulnerable system; the defender needs to be perfect."

The good news is that technology trends and new doctrine for cybersecurity have dramatically changed the terrain of cyberspace. Companies at the leading edge of cybersecurity have been able to manage the threat from even the most sophisticated actors. If these trends continue and if policy is put into place to correctly align incentives, it is possible that in five years we may view cybersecurity broadly as a manageable problem. The bad news is that emerging technologies may once again favor the attacker, erasing the defensive gains of the past decade. In my remarks below, I will review the "good news" of the last decade and how these trends can be accelerated and adoption of better cybersecurity practices encouraged by Congress. I then will discuss the "bad news" of how emerging technology trends like artificial intelligence, the Internet of things and 5G, and quantum computing could favor the offense. I then provide some thoughts for how Congress can promote wider adoption of cybersecurity practices that are on the cutting edge today and shape the future of technology so that defenders are not left at a disadvantage tomorrow. Finally, I conclude with a brief review of the projects I am working on today that may help us build a more resilient cyber future.

The Good News: Cybersecurity is Possible

There is an old joke in cybersecurity, attributed to Dmitri Alperovitch, now the Chief Technology Officer (CTO) of the cybersecurity firm Crowdstrike. The joke, retold in many formulations, is always along the lines of "there are two types of companies: those that have been hacked and know it and those that have been hacked and don't know it." That may have been true a decade ago, but today there are three types of companies: those that have been hacked and know it, those that have been hacked and don't know it, and those that are actively and successfully managing the risk.

In *The Fifth Domain*, Clarke and I conclude that the greatest advance in cybersecurity over the last decade was not a technology but a white paper. In "Intelligence Driven Security" a group of researchers and practitioners at Lockheed Martin presented the processes they had developed for detecting and disrupting adversary activity along the "Cyber Kill Chain". Published in 2011, the paper showed how defenders could take the advantage away from adversaries by breaking down the process by which an adversary attempted to achieve an objective on a network and building a security program around each of those steps. Unlike in conventional thinking on cybersecurity where a network compromise is considered a failure, the Kill Chain methodology sees that as only one step in the chain. Before an adversary can exploit an initial host on a network, they must engage in reconnaissance of the target, weaponize what they have learned into a package capable of compromising the target and deliver it. After they have achieved the initial exploitation, they then need to gain administrative rights, move laterally across the network to find their target, and then carryout out their intended action. That action might be to exfiltrate data off the network or to destroy operational systems. Whatever their goal, it is not simply to compromise a single system.

The concept of the kill chain has evolved and expanded since first published. MITRE Corporation has developed the ATT&CK Matrix to further breakdown the steps that happen after initial compromise into 22 discrete steps. However you breakdown the attackers progression, the key takeaway should be that detecting and stopping them is possible. Whether the adversary needs to go through seven steps or 22, they have to successfully avoid detection at each stage; defenders only need to detect them at any one stage. Once the adversary is on the defender's system, the defender should have the advantage. Gaining that advantage requires knowing the topology of your system better than the adversary and being able to detect anomalous behavior within it. This ability to detect and respond rapidly is what Crowdstrike and other companies have specialized in. Endpoint Detection and Response (EDR) has been the technical capability that has enabled "threat hunting" along the kill chain to occur at scale within enterprises. Managed Detection and Response companies are rapidly bringing these capabilities to the middle market.

Beyond detection and response, newer technologies have the potential to remove large swaths of risk. When properly deployed and managed with security in mind, cloud computing, containerization, and software defined networking, to name just three emerging technologies, can provide real advantages to defenders. Virtualization can allow new computing environments to be spun up and down for a specific purpose so rapidly that gaining a foothold in one of these new environments does an adversary no good because the environment itself does not persist. These

technologies can also allow for deception campaigns on a massive scale to create new opportunities for detection and to increase the work factor of adversaries.

All this adds up to the potential to make our country, our companies, and ourselves resilient to cyber attacks. Through the adoption of secure by default technologies we should be able to make it so that almost all attacks "bounce off" and that we can "bounce back" when attacks do succeed. From a policy perspective, what is needed now are the incentives and requirements to promote the adoption of these techniques and the technologies beyond the small handful of companies that are deploying them in a holistic way today. And of course, this transition needs to occur at a faster rate than adversaries can adopt new technologies that defeat them.

The Bad News: Technology Changes Could Erase These Gains

Just as we may be turning a corner on security, the technology landscape may change in ways that are not evolutionary but revolutionary. By that I mean that the technology coming online is not about the continuation of current trends or even the acceleration of trends but whole new classes of technology. Artificial intelligence, quantum computing, and 5G and the Internet of things may not intrinsically favor attackers over defenders but the offense is likely to adopt technologies that can give them an advantage faster than defenders and their targets are likely to adopt new technologies in ways that open up new swaths of vulnerabilities. I would like to now discuss three such technologies: 1) artificial intelligence; 2) 5G and the Internet of Things; and 3) quantum computing.

Artificial Intelligence

Arguably, artificial intelligence up until now has been a technology that has favored the defense. Many of the gains discussed above in the last decade are do to artificial intelligence applications within cybersecurity. For instance, the ability of advanced endpoint protection programs to identify never before seen malware using machine learning has made the work of adversaries much more difficult. The bad news is that as a the state-of-the-art in artificial intelligence advances, attackers are likely to use it in ways that will upend the basis of today's security architectures.

Deepfakes have made headlines recently in the political world. For public figures who have thousands of hours of voice and video recordings available online, artificial intelligence can now be used to piece together snippets of them talking to literally put words in their mouths. Deepfakes are likely to come into play heavily in the 2020 election and defenses against them are lagging. Use of AI for deepfake detection made news over the summer but in this arms race, adversaries look to have an advantage, tweaking their tools and testing against deepfake detection technology until they can defeat it.

Initially, Deepfakes required large libraries of voice and video but as the technology improves, the amount of source data required is rapidly coming down. That will mean that many of the fundamental controls we have in place today to combat cybercrime may no longer be trusted. The cybersecurity community has worked hard to educate companies about the dangers of wire

transfer fraud – to train finance departments to be suspicious of emails from the CEO ordering them to wire funds on an emergency basis, for instance. But what if, instead of compromising the email system, adversaries compromise voice and video systems, and your boss in her natural speaking voice that you hear everyday, calls you to confirm that she does in fact need you to wire those funds right now? The ability to create deepfakes from smaller and smaller sets of source material will make that scenario possible for many companies in a short period of time. That will mean that the ultimate root of trust – believing what we see and hear – can no longer be trusted.

5G and the Internet of Things

Internet of Things (IOT) technology is rapidly being distributed within critical infrastructure and in homes and businesses in ways that appear to ignore the security lessons we learned over the last twenty years within enterprise systems. Coding practices are poor in the space, firmware is difficult to update, and systems are widely exposed to the public Internet. What's more, with the advent of 5G, massive, ubiquitous wireless connectivity will mean that many of these devices will be directly connected to the public Internet with no defense-in-depth built around them. Within the consumer market, we have seen a troubling trend of "set and forget" connected devices that, after being setup, are not monitored for security and do not receive updates to their software after problems are discovered. Unfortunately, this trend does not appear to be confined to the home IOT market. The same problem is occurring even within industrial control systems.

Quantum Computing

Far more than these other two technological shifts, quantum computing is likely to upend computer security because it will upend computing. A calculation that might take a classical computer several centuries to complete could be done by a quantum computer in the blink of an eye. Experimental systems today are showing a lot of promised toward achieving this kind of capability. Google may already have achieved what is known as "Quantum Supremacy", using a quantum computer to complete a mathematical equation faster than a conventional system could.

Quantum computing has the potential to be extremely disruptive to security, allowing encryption protocols to be defeated; whether quantum resistant encryption will be deployed ubiquitously and will prove to defeat quantum computing is an open question. The combination of artificial intelligence technology with quantum computing open some scary possibilities. More than anything else, government needs to ensure that the United States is a leader, not a follower, in the development of quantum computing.

The Ugly: Government Intervention in Necessary

For most of the last twenty years, US government policy across Administrations has largely been about getting out of the way and hoping that markets would solve cybersecurity problems on their own. Where government has intervened, intervention has been uneven and light touch. Today, I believe we are starting to recognize that markets alone will not solve our cybersecurity dilemma. I think it is fair to conclude that the industries that are doing the best at actively managing risk in cyberspace are also actively regulated: financial services and the defense

industrial base. Many of the approaches to security that are working today were pioneered in these sectors. Driving these innovations to other markets will require creating the right set of incentives and requirements. I have been pleased to see that more so than in any previous Administration, the current leadership of the Department of Homeland Security has recognized that regulation, smartly and carefully implemented, is necessary to drive the level of security required for our nation. The Department's cybersecurity strategy is explicit on this point. In the IOT space, DHS should lead efforts to regulate the security of IOT devices in the sectors that it regulates including chemicals, pipelines, and the maritime industry.

I believe that the Internet of Things Cybersecurity Improvement Act would be a good first step toward improving IOT security. The Act would set standards that sellers of IOT technology to the Federal government would need to meet as well as establish disclosure requirements when manufacturers discover vulnerabilities. The approach uses government's massive purchasing power to improve security more broadly. Companies that develop technologies on a "build once, sell everywhere" model will likely meet the governments requirement for all their commercial offerings rather than just for those sold to government. These requirements, once set, could then be adopted to regulate the use of IOT in critical infrastructure sectors.

Fundamentally, however, I believe that setting requirements is insufficient. We need to make device makers responsible for the full lifecycle of security by making them liable for harm caused by their devices. I recognize that this notion is a radical departure from how we have approached liability within the information technology realm thus far but now that these devices are making their way into national security systems and life safety systems, I think it is critical that we create incentive structures that truly value security. In the next section, I discuss one effort we have undertaken at the Global Resilience Institute to create a model for liability for cybersecurity.

Beyond, IOT, the leadership of the Cybersecurity and Infrastructure Security Agency (CISA) has made election security the agencies number one priority. CISA will need to build on its current efforts to counter-election interference to play a role in combating the proliferation of deepfakes in the political realm and for enterprise security. Crucial to this effort will be building strong, operational partnerships with social media companies that go well beyond today's arm length interactions. Steps must be taken to breakdown the reluctance by Facebook, Google, Twitter, and other social media companies to truly partner with government on this problem.

For quantum computing and artificial intelligence, government's role should be less about managing the cybersecurity implications and more focused on ensuring that the United States competes and wins in these technologies. I tend to be skeptical of analogies to arms races or calls for Apollo Programs or Manhattan projects, but on the basic science in these fields, those kinds of approaches are warranted. Both China and Russia have made gaining an advantage in AI a national priority. China has also done that on quantum. I believe our market based approach to technology development comes with real advantages but in the development of these core capabilities, I worry that a race that is the Chinese state vs. Silicon Valley is one that Silicon Valley will lose. We need a national effort to ensure that US technology leadership continues into the next decade.

Each of these lines of effort will take at least half a decade to produce meaningful results - thus it is crucial that the efforts begin now.

What We are Doing at GRI

The challenges we face are large, but they are not insurmountable. While much work remains to be done, let me take this opportunity to highlight four efforts underway at the Global Resilience Institute that may contribute to improving our national cyber resilience over the next five years.

Creating a National Transportation Safety Board for Cyber Incidents

Resilience is a concept that we have talked a lot about in the field of cybersecurity but it a far better developed idea in other fields like emergency management and psychology. One of the key components of resilience I have taken away from studying the concept in these other fields is the importance of adapting following a bad outcome. Learning from disasters or even from so-called "near misses" is critical to the development of resilience. To this end, as far back as 1991 practitioners in the field have suggested that government should develop the equivalent of a National Transportation Safety Board (NTSB) for cybersecurity incidents, a "Cyber NTSB". Given that this idea was first suggested three decades ago but has yet to reach fruition, we are planning a workshop, sponsored by the National Science Foundation, to develop a prototype process for how such an organization would operate. We plan to hold the workshop in the spring of 2020.

Building a High Assurance Network for Collaborative Defense

Critical to building resilience is creating a model for Collaborative Defense. The "partnership" that has been the central tenet of our national cybersecurity policy for two decades needs to evolve to real-time, operational collaboration. In order for that to happen, we need collaboration platforms where the members of this partnership can trust each other. Government needs to be able to trust that the intelligence it shares will be protected and only shared appropriately and securely. But private companies need the same degree of assurance when they share with government and with each other. Today, the platforms on which we collaborate, internet connected, general purpose computers, are not trustworthy. Moreover, we often do not know whether we can trust our partners that are using those computers.

When I testified before this committee two years ago, I discussed early thinking about how to develop such a network. Today I am pleased to say that, working with our partners at the Advanced Cybersecurity Center and with a generous grant from a private foundation, we have developed a prototype network. This network takes advantage of the trends in computing that have dramatically lowered cost: in-expensive computing at endpoints and cloud computing to provide immense computing power for analytics and other services. For about $300 a year, we can provide a high assurance endpoint that can only be accessed by specified users to connect to a secured, private network for threat collaboration. This model provides the basis for addressing the issue of trust in the users and trust in the systems by replicating at far lower costs many of the design criteria of the classified networks used by government today.

In my view, the model we have developed should be adopted by the Department of Homeland Security to create what we have dubbed CInet for Critical Infrastructure Network. Using existing authorities, the Secretary of Homeland Security should establish a new safeguarding standard for Confidential information, the existing level below Secret in the classification schema. The standard should be built around the prototype we have developed which eliminates the most common paths to compromise (spear-phishing, credential compromise, and watering hole attacks) and prevents end users from unintentionally releasing information through a series of technical controls. Having vetted the concept with a handful of critical infrastructure companies, we believe that this model could fit into the current operating models within critical infrastructure security operating sectors. We also believe that by harnessing current best practices in the private sector for continuous monitoring of insider threats, the Secretary could also promulgate a different standard for granting of clearances at the Confidential level that would be better, faster, and cheaper. Then would come the hard part of convincing the intelligence community to target collection to provide relevant threat intelligence to participating companies and to downgrade it to the Confidential level.

Designing a Darknet for the Electric Grid

Many of the same technology trends that could provide attackers an advantage over the next five years can also be harnessed to increase security for critical infrastructure. Advances like software defined network (SDN), increased mobile bandwidth with 5G, and artificial intelligence can enable far higher degrees of assurance for critical infrastructure than can be attained today. This is the idea behind our Darknet project to create a separate network for the electric grid using "dark" or unlit fiber optic cables. GRI initially began work on this concept with a grant from a private foundation and is now partnering on it with Oak Ridge National Laboratory.

Developing an Insurance Regime that Promotes Better Security

Cyber insurance was supposed to help drive down risk. In theory, the insurance sector, in exchange for providing insurance coverage, would require companies to prove that the risk they underwrote was being managed. In practice, as the recent spate of ransomware attacks on city governments has demonstrated, cyber insurance is simply transferring the risk and enriching the criminal groups behind the attacks. Yet, in other sectors, insurance markets have proved remarkable mechanisms for encouraging risk reduction. Dr. Stephen E. Flynn, the director of Northeastern's Global Resilience Institute, and I have been developing a model for insurance that would promote risk reduction rather than just risk transference. Dr. Flynn, a retired Coast Guard officer, has posited that the regime put in place under the Oil Pollution Act of 1990 after the *Exxon. Valdez* oil spill could be ported over for data security. In other words, we should treat data spills like oil spills. Under that regime, ships entering US waters must provide proof in the form of a Certificate of Financial Responsibility that their owners or their guarantors in the insurance industry have the financial resources to cover the cost of cleaning up an oil spill should containment on their vessel fail. Notionally, owners of data could be required to take out insurance policies to cover the full societal cost should they fail to protect the data that they hold. In this thinking, Congress could establish a dollar figure per record and then require holders of personal data to obtain insurance to cover those loses. From there, market mechanisms would

take over to determine how to price risk. This model could also be adapted for critical infrastructure. For instance, if natural gas pipeline owners had to obtain private insurance to cover the costs of a disruption to service caused by malicious cyber activity, markets would likely require a far higher degree of assurance than would be required through a standard regulatory model. In the coming months, we will engage the insurance industry on further developing this concept.

Biography

Rob Knake is a Senior Research Scientist at Northeastern University's Global Resilience Institute and a Senior Fellow at the Council on Foreign Relations.

Rob served from 2011 to 2015 as Director for Cybersecurity Policy at the National Security Council. In this role, he was responsible for the development of Presidential policy on cybersecurity, and built and managed Federal processes for cyber incident response and vulnerability management.

He is co-author of *Cyber War: The Next Threat to National Security and What to Do About It* (HarperCollins 2010) and *The Fifth Domain: Defending Our Country, Our Companies, and Ourselves in the Age of Cyber Threats* (Penguin 2019).