



J U L Y
2 0 1 4

Surviving on a Diet of Poisoned Fruit *Reducing the National Security Risks of America's Cyber Dependencies*

By Richard J. Danzig



Center for a
New American
Security

Acknowledgements

This paper is about technology, conflict and insecurity. In contrast to the topic, it is impressive and affirming that many people helped me in ways that were warmly personal, generous and completely altruistic. In this effort to chart a path through a dark forest, I regularly got lost, banged into things I knew I didn't understand, and thought I understood things I didn't. Repeatedly, the people listed below put aside other important work they were doing and helped me. Of course, they bear no responsibility for errors of fact or judgment in this paper. Those are mine alone. But they tried to save me from them by talking to me, reading drafts, contributing ideas, criticizing my thinking, referring me to literature, reading new drafts and beginning the cycle again. As I progressed, regressed and progressed, I intensely appreciated their help. In retrospect, I appreciate it even more. These are people who care: care about accuracy in a product, care about helping someone (in this case me) who will put his name to something he knows less about than they do, and care about what the United States does about an issue whose importance and character they have struggled to understand and most of the rest of us need to.

I am deeply grateful for the friendship and help of the following:

Arati Prabhakar, Dan Kaufman and Norm Whitaker at DARPA contracted for this paper and then went beyond that and provided encouragement, ideas and active discussion of the results. Inspired by their leadership, Richard Guidorizzi, Mike Hsieh and Howie Shrobe in DARPA's Information Innovation Office and Paul Kozemchak in the Office of the Director provided valuable comments. Mike Walker (also a program manager in the Information Innovation Office) was extraordinarily generous in reading several drafts, suggesting additional readings and contributing exceptional insights. Mike's thinking particularly helped the discussion in Part II about searching for vulnerabilities. Craig Fields, a former director of DARPA and now chair of the Defense Science Board, was particularly encouraging and usefully critical of this paper. As noted in the endnotes to this paper, the views expressed are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. government.

Outside of DARPA, I especially want to recognize Jim Gosler (Johns Hopkins Applied Physics Laboratory), Ralph Langner (Langner Consulting, Munich, Germany), Peter Levin (Amida Technology), Dan Geer (In-Q-Tel) and Gary McGraw (Cigital Inc.). From the beginning, Jim made extraordinary efforts to help me understand the subject and his concerns about it. He sustained me with criticism and enthusiasm throughout. His insights particularly underlie the points about why we need to presume intrusion and the use of out-of-band responses as a means of protection. Ralph's precision, deep knowledge about and fine focus on cyber-physical attacks helped me a great deal. Peter came to this enterprise later than the others but immediately threw himself into making it better. However good it is or isn't, he did indeed make it better. One of Peter's contributions was introducing to me Dan Geer, who turned his immense energy, ingenuity and reflective qualities to offering some 90 comments on an early draft and nearly the same number for a later version. The errors and shortfalls that remain reflect only that there are limits to how much even Dan can educate me. Last on this list, but near the first chronologically, Gary, from the start of my efforts, directed some of his immense warmth, great exuberance and vast technical knowledge in my direction. This quintet – with Mike Walker, a sextet – sustained me through a difficult task. I note further that I would never have written this paper were it not for John Mallery (MIT) and David Mussington (Institute for Defense Analyses), who came to my house and urged me to undertake it. I am additionally grateful to

John for his subsequent comments on the document.

I would also like particularly to express my gratitude for ideas, criticism and encouragement from Ross Anderson (Cambridge University, U.K.), Michael Assante (Idaho National Laboratory), Chris Betz (Microsoft), Pierce Corden (American Association for the Advancement of Science), Sadie Creese (Oxford University, U.K.), Bill Crowell (Alsop Louie Partners), Bruce deGrazia (GHS Advisors), David Ferbrache (Ministry of Defence, U.K.), Mike Frantzen (Kudu Dynamics), Marc Gordon (American Express), Dan Guido (Trail of Bits), Melissa Hathaway (Kennedy School of Government), Bruce Held (Department of Energy), Michael Hopmeier (Unconventional Concepts), Martin Howard (U.K. Ministry of Defence), Paul Kaminski (Defense Science Board), Chris Kirchhoff (Joint Chiefs of Staff), Robert Knapp (The Evergreen State College), Frank Kramer (Atlantic Council), Irv Lachow (MITRE), Carl E. Landwehr (Cyber Security Policy and Research Institute), Martin Libicki (RAND Corporation), Tim Maurer (New America Foundation), Anne Neuberger (NSA), Joel Molinoff (CBS Broadcasting), Joe Nye (Kennedy School of Government), Andy Ozment (National Security Council and then Department of Homeland Security), Alan Paller (SANS Institute), Scott Proctor (eBay), Jonathan Reiber (Department of Defense), David Robinson (Robinson+Yu), Andrew Ruef (University of Maryland), Ron Sanders (Booz Allen), Jamie Saunders (Foreign and Commonwealth Office, U.K.), Lara Schmidt (RAND Corporation), Fred Schneider (Cornell University), Ari Schwartz (Department of Commerce), George Scott (Department of the Navy), John Stewart (Cisco), Michael Sulmeyer (Department of Defense), Tomas Vagoun (National Coordination Office for Networking and Information Technology R&D), Brandon Wales (Department of Homeland Security), Peter Weinberger (Google), Rosemary Wenchel (Department of Homeland Security), Tom Wheeler (then at Core Capital Partners, now chair of the FCC) and Mudge Zatko (Google).

Small group discussions at Boeing (with John Toomey, Jeff Trauberman, Leo Brooks, Greg Deiter, Kevin Meehan, Jeremy Bayer and Fred Schvien), Endgame (with Nate Fick and Niloofar Razi Howe), Mandiant (with Christopher Glyer, Marshall Heilman, Kevin Gronberg, Jen Weedon, and Doug Wilson), ManTech (with Bill Varner, Ken Silva and Eric Eifert) and Microsoft (with Chris Betz, David Burt, Jerry Bryant, Bruce Dang, Scott Cul, Steve Lipner, Tim Rains, Paul Nicholas and Matt Thomlinson) were very valuable in addition to written comments on the text from many of these individuals. Nate Fick contributed a second reading as well and suggested collocating the proposed new federally funded research and development center on the West Coast with a private sector board.

This paper was prepared for publication by the Center for a New American Security. At that admirable institution, I am grateful to Ben FitzGerald, Shawn Brimley and Bob Work for endorsing publication, to Ben FitzGerald, Jim Miller and Bridge Colby for several times reading and improving Recommendation 5, to Jacob Stokes for valuable research on the history of Internet security and to Dafna Rand and Liz Fontaine for their preparation of the manuscript. Evan Waranowski and Mary Gladstone particularly earned my gratitude for substantial and careful work checking and standardizing all the footnotes and Mary for her remarkable care in editing the entire manuscript.

I reiterate that all of the above contributed as individuals, not as representatives of their organizations. Of course, the views described here and any errors that remain are my responsibilities, not theirs.

The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. government.

TABLE OF CONTENTS

I. Executive Summary	5
II. Introduction	8
III. Why Information Systems Are Vulnerable	9
IV. Cyberdefense and Its Limits	13
V. Three Other Causes of Cyber Insecurity	17
VI. Creating a Minimal Shared Cybersecurity Standard and Some Recommendations	19
Appendix: A Note on Deficiencies of Data and Metrics	53

J U L Y 2 0 1 4


Surviving on a Diet of Poisoned Fruit

Reducing the National Security Risks of America's Cyber Dependencies

By Richard J. Danzig

About the Author

The Honorable Richard J. Danzig is a board member at the Center for a New American Security; the vice chair of the Board of Trustees of The RAND Corporation; a member of the Defense Policy Board and The President's Intelligence Advisory Board; and a director of Saffron Hill Ventures.



SURVIVING ON A DIET OF POISONED FRUIT: REDUCING THE NATIONAL SECURITY RISKS OF AMERICA'S CYBER DEPENDENCIES

By Richard J. Danzig

J U L Y 2 0 1 4

Surviving on a Diet of Poisoned Fruit
Reducing the National Security Risks of America's Cyber Dependencies



I. EXECUTIVE SUMMARY

By Richard J. Danzig

Digital technologies, commonly referred to as cyber systems, are a security paradox: Even as they grant unprecedented powers, they also make users less secure. Their communicative capabilities enable collaboration and networking, but in so doing they open doors to intrusion. Their concentration of data and manipulative power vastly improves the efficiency and scale of operations, but this concentration in turn exponentially increases the amount that can be stolen or subverted by a successful attack. The complexity of their hardware and software creates great capability, but this complexity spawns vulnerabilities and lowers the visibility of intrusions. Cyber systems' responsiveness to instruction makes them invaluable flexible; but it also permits small changes in a component's design or direction to degrade or subvert system behavior. These systems' empowerment of users to retrieve and manipulate data democratizes capabilities, but this great benefit removes safeguards present in systems that require hierarchies of human approvals. In sum, cyber systems nourish us, but at the same time they weaken and poison us.

The first part of this paper illuminates this intertwining. The second part surveys the evolution of strategies to achieve greater cybersecurity. Disadvantaged by early design choices that paid little attention to security, these strategies provide some needed protection, especially when applied collectively as a coordinated "defense in depth." But they do not and never can assure comprehensive protection; these strategies are typically costly, and users will commonly choose to buy less security than they could obtain because of the operational, financial or convenience costs of obtaining that security.

Three other factors, discussed in Section V, amplify cyber insecurity. First, the cyber domain is an area of conflict. Cyberspace is adversarial, contested territory. Our adversaries (including criminals, malevolent groups and opposing states) co-evolve

with us. The resulting ecosystem is not static or stable. Second, the speed of cyber dissemination and change outpaces our recognition of problems and adoption of individual and societal safeguards to respond to them. Protective actions are likely to continue to lag behind security needs. Third, in cyberspace America confronts greater-than-customary limits to U.S. government power because of the global proliferation of cyber capabilities, cyber attackers' ability to remain outside the United States even while operating within the country's systems and our likely inability, over the long term, to avoid technological surprise. Two-thirds of a century of technological dominance in national security matters has left the United States intuitively ill-prepared for technology competitions that it probably will not continue to dominate and in which there is a high likelihood of surprise.

What then is to be done? The concluding part of this paper does not attempt to recapitulate or evaluate efforts now extensively debated or in progress. It focuses instead on recommending initiatives that deserve fresh attention from U.S. government decision-makers. These include:

1. Articulate a national security standard defining what it is imperative to protect in cyberspace. The suggested standard is: *"The United States cannot allow the insecurity of our cyber systems to reach a point where weaknesses in those systems would likely render the United States unwilling to make a decision or unable to act on a decision fundamental to our national security."* A more stringent standard may later be in order, but this standard can now secure a consensus, illuminate *the minimum* that the United States needs to do and therefore provide an anvil against which the nation can hammer out programs and priorities.
2. Pursue a strategy that self-consciously sacrifices some cyber benefits in order to ensure greater security for key systems on which security depends. Methods for pursuing this strategy include stripping down systems so they do less but have fewer vulnerabilities; integrating humans and other out-of-band (i.e., non-cyber) factors so the nation is not solely dependent on digital systems; integrating diverse and redundant cyber alternatives; and making investments for graceful degradation. Determining the trade-offs between operational loss and security gain through abnegating choices will require and reward the development of a new breed of civilian policymakers, managers and military officers able to understand both domains.
3. Recognize that some private-sector systems fall within the national security standard. Use persuasion, federal acquisition policies, subsidy and regulation to apply the abnegating approach to these systems. While doing this, reflect an appreciation of the rapidity of cyber change by focusing on required ends while avoiding specification of means. Refrain from regulating systems that are not critical.
4. Bolster cyber strategic stability between the United States and other major nation-states by seeking agreement on cyber constraints and confidence-building measures. As an early initiative of this kind, focus on buttressing the fragile norm of not using cyber as a means of physical attack between China, Russia and the United States.
5. Evaluate degradation in the sought-after certainties of mutually assured destruction (MAD) as a result of uncertainties inherent in cyber foundations for nuclear command, control and attack warning. If we are moving to a regime of mutually unassured destruction (MUD), suggest to China and Russia that we are all becoming less secure. Then pursue agreements that all parties refrain from cyber intrusions into nuclear command, control and warning systems.
6. Map the adversarial ecosystem of cyberspace in anthropological detail with the aim of increasing

our understanding of our adversaries *and our own* incentives and methods of operation.

7. Use the model of voluntary reporting of near-miss incidents in aviation to establish a data collection consortium that will illuminate the character and magnitude of cyber attacks against the U.S. private sector. Use this enterprise as well to help develop common terminology and metrics about cybersecurity.
8. Establish a federally funded research and development center focused on providing an elite cyber workforce for the federal government. Hire that workforce by cyber competition rather than traditional credentials, and promote, train, retain and assign (including to the private sector) that workforce by standards different from those currently used in federal hiring.

II. INTRODUCTION

What are the risks for U.S. national security from our increasing dependence on the greatest technological development of recent decades: digital information systems? How should the United States diminish the possibilities that computer memory, logic and communication networks can be disrupted and usurped? How can the nation reduce the consequences of cybersecurity failures to the extent they occur?²

These questions were first intensely asked about the information technology (IT) systems supporting military combat capabilities.³ In recent years, concerns broadened to focus on the vulnerability of civilian transport, communication and electric power systems on which many military systems depend. These vulnerabilities have traditional security implications because military technologies and operations depend on information technology. Beyond that, it is now evident that intellectual property and commercially strategic information stored on IT systems are being accessed and exfiltrated, perhaps to a degree that affects America's economic position.⁴

The Stuxnet computer worm and the Iranian cyber responses to it pose another set of issues. While denial of service attacks, spear phishing and website defacements demonstrate that cyber attacks can be destructive to cyber systems, Stuxnet and its kin show how actors that penetrate cyber systems can use them to attack physical targets.⁵ This development poses the question of whether we are entering an age in which attacks through cyberspace may achieve significant and dramatic national consequences, for example by physically damaging or disabling important parts of our power grid,⁶ financial systems or other parts of society.

Understanding the danger of these threats requires comprehension of the confusing complex of

"We are staking our future on a resource that we have not yet learned to protect."

GEORGE TENET, DIRECTOR OF THE
CENTRAL INTELLIGENCE AGENCY,
APRIL 6, 1998¹

technologies, the varied types of attackers, the range of consequences of attacks and possible prophylactic or mitigating actions. This paper aims to improve the U.S. government's responses by increasing policymakers' understanding of cyber insecurity and by recommending a number of actions that are now inadequately featured on the national security agenda.

Section III is designed to deepen the lay reader's understanding about why information technology security weaknesses exist and are likely to persist. Section IV offers a brief summary of defensive efforts that have been undertaken to mitigate these weaknesses, as well as the limits of these efforts. Section V highlights three factors that constrain U.S. government efforts to improve cybersecurity. Section VII proposes a minimal national security standard for cybersecurity and then offers nine specific recommendations for U.S. government actions beyond those already occurring.

III. WHY INFORMATION SYSTEMS ARE VULNERABLE

The beginning of wisdom about cyber systems is to understand that vulnerability is inherent in the technology. Information technologies offer a Faustian bargain: The capabilities that make these systems attractive make them risky. Cyber systems create serious security problems because they concentrate information and control and because the complexity, communicative power and interactive capabilities that enable them unavoidably create vulnerabilities. The low visibility of IT operations and the speed of change within the field intensify these difficulties. Though good cybersecurity practices can mitigate sources of insecurity and poor practices can amplify them, there are no practices that can completely eliminate them.

Concentration of Information and Manipulative Power

Cyber systems are designed to facilitate the storage and manipulation of vast quantities of information. Previously, for example, teams of workers physically opened and closed valves on pipelines, credits and debits were recorded at banks by hand, and precious information was stored under lock and key. Information technology allows these things to be done electronically and therefore much more efficiently and quickly, by a single direction in a large number of nearly simultaneous instances and without some aspects of human error. But it also inherently empowers anyone, including unauthorized actors, who can gain access to this capability.

The beneficial use of computers amplifies the abilities of nefarious individuals, groups and states to destroy pipes, manipulate financial records or abscond with valued information. The greater the scale and benefit to the legitimate user, the greater the disruption, diversion or destruction that the usurper can achieve. Edward Snowden

stole an estimated 1.7 million documents⁷ – and before him, PFC Manning, some 700,000 documents⁸ – from U.S. government classified stored data.⁹ Penetration of one business (Target) resulted in the theft of credit or other personal information in numbers equal to about 40 percent of the American adult population.¹⁰ Many have previously committed thefts and engaged in espionage, but never before had these acts yielded such vast troves of data. Modern information technologies empower and incentivize¹¹ subversion at scale.

Information technologies offer a Faustian bargain: The capabilities that make these systems attractive make them risky.

Communicative Capacity

Because communications, like other information, can be expressed as binary bits, improvements in the speed, quantity and cost of digital systems produce great gains in communicative capacity.¹² Breakthroughs in packet switching (sending parts of messages in pieces by diverse routes rather than through central hubs), fiber optics and related technologies compound these gains. To continue the previous examples, these changes permit central processors and/or humans in the loop to receive signals indicating whether pipeline valves at great distances are open or closed and whether changes in pressure warrant different settings. They make it possible for banks with immense numbers of transactions to receive and send funds nearly instantly among themselves and between themselves and clients. They permit remote surveillance of precious goods. They empower owners and vendors to monitor and

modify software and hardware operations from a distance. Because of this communicative power the proliferation of computational capabilities has engendered the growth of networks.

But network vulnerabilities greatly expand computational vulnerabilities. Remote access creates opportunities for subversive entry. The global character of the World Wide Web assures that these opportunities are widely shared. As the director of the CIA put it in 1998, when confronting the shock of the news: "Think about it for a moment – we share the same network with our adversaries."¹³

One of the rewards, but one of the risks, of present information systems is that they tear down gates and eliminate gatekeepers.

Communicative capabilities combine with the accessibility of information and its standardization to permit the removal of intermediaries charged with controlling information and the processing of transactions. This development, which a Microsoft executive has called "disintermediation,"¹⁴ obviates the requirements for, among others, travel agents, hotel reservation desks and bank clerks before booking rooms and tickets or withdrawing money. Disintermediation has obvious benefits, but removing hierarchies of approval opens opportunities for malevolent actors. One of the rewards, but one of the risks, of present information systems is that they tear down gates and eliminate gatekeepers.

Moreover, the need to receive and respond to messages means that portions of the operating

system of the recipient computer have to be engaged for tasks such as directing incoming data, acting on them when they provide machine instructions and translating them into a form readable by humans. That engagement creates risks of manipulation. Destructive directions can be commingled with benign incoming material.¹⁵ Many users now recognize this when they avoid "phishing attacks" – emails with links that when clicked (i.e., executed) launch contaminating programs.¹⁶ More subtle forms of contamination through interaction may, however, be spawned by PDF files, photos and websites (including those arrived at by misdirection).¹⁷

Finally, the communicative capability of modern computers permits exfiltration of data without ever assuming the complications, risks and costs of having a person enter places where that data is stored. It also facilitates attackers' operating in many places simultaneously.

Size and Complexity

Large programs offer a great many places for a hostile actor to hide malevolent code.¹⁸ They also offer innumerable points of attack ("attack surfaces"). As the cost and size of computing hardware have plummeted, the density of hardware has increased. Some graphics processing systems, for example, utilize more than a billion transistors.¹⁹ The size of software programs has concomitantly expanded. Responding in large measure to hardware opportunities, the Linux operating system, for example, has grown from 176,000 lines of code when introduced 20 years ago to over 15 million lines of code in 2011.²⁰ There are reported to be some 8.6 million lines of code in the Pentagon's new Joint Strike Fighter aircraft.²¹ A third-party calculation estimated some 50 million lines of code in Microsoft's Vista operating system.²²

Large programs have large numbers of errors and points of weakness. And the concision and

precision of software programs creates weaknesses of such subtlety that they are hard to discern.²³ A layperson might think of the problem as akin to using English to write complex legislation – for example, the 4-million-word U.S. tax code. Loopholes inevitably occur. Though they can be reduced by more controlled and precise programming language²⁴ and by checking and rechecking the language, no language has been developed that offers broad functionality without vulnerability.²⁵ Attackers have an inherent advantage: They can succeed with single points of entry. Defenders must attempt to plug all holes.

It is a widely accepted rule of thumb that one error will be introduced for every thousand lines of well-written code that are embedded in a software system²⁶ and that the defect rate increases with the size of the project “with very large projects having four times as many errors per thousand lines of code as small projects.”²⁷ Errors are not equivalent to security risks and too much precision should not be read into these estimates.²⁸ But they are indicative of the magnitude of the problem. A significant error rate and substantial vulnerabilities indubitably exist in all major programs and are likely to remain for the foreseeable future.

Interaction of Software

The interaction of code with other code amplifies complexity, just as the interaction of legislation with other legislation has the same effect. Software is extensible²⁹ – it is written to interact (“interface”) with other software, so even if an operating system were without apparent vulnerability, vulnerabilities are likely to be introduced as applications are used to perform tasks (for example, as an Excel spreadsheet or an Adobe PDF reader is used in conjunction with a Microsoft operating system). Some of these vulnerabilities will only be apparent when programs operate together. The result is combinatorial explosion. “It is not sufficient merely to prove a program correct; you have to test it too.

Verizon calculates that in two out of three cases of data breach, the loss took “months or more to discover.” Similarly, in 7 out of 10 cases victims only learned about their loss from third parties.

Moreover, to be really certain that a program is correct, you have to test it for all possible input values, but this is seldom feasible. With concurrent programs, it’s even worse: You have to test for all internal states, which is, for all practical purposes, impossible.”³⁰

Low Visibility of Attacks and Attackers

All the preceding characteristics (centralization, communicative capacity, size, complexity and interaction) help attackers to hide their work. Verizon calculates that in two out of three cases of data breach, the loss took “months or more to discover.” Similarly, in 7 out of 10 cases victims only learned about their loss from third parties.³¹ And this database of course does not include cases that are not recognized at all.³²

Temporal Linkages

Software systems are chained to the past and exposed to the future. It is typically impractical to abandon past systems (“legacy systems”) and associated data. To be useful, most new systems have to integrate old ones, and this tends to expose them to the vulnerability of the old. Even if safe and secure at present, they are vulnerable to unpredictable interactions in the future. “[S]ome attacks that are not possible in today’s state of the world may become possible in the future and invalidate the design environment in which a standard was crafted.”³³

Successful strategies must proceed from the premise that cyberspace is continuously contested territory in which we can control memory and operating capabilities some of the time but cannot be assured of complete control all of the time or even of any control at any particular time.

The consequences of all the factors described above are apparent. A recent IBM report describes trends and magnitudes of vulnerability discovery:

Since 1997, the IBM X-Force has been documenting public disclosures of security vulnerabilities. Back then, there were a handful of vulnerabilities to document each week. Now, over fifteen years later, we document an average of over 150 vulnerabilities per week. ... Our database now contains 70,000 unique vulnerabilities and continues to climb at a steady pace averaging 7,700 vulnerabilities per year over the past five years. ... There were 3,436 vulnerabilities that had public exploits available in 2012, which comprised the total number of public exploits for the year. This is 42% of the total number of vulnerabilities ...³⁴

Vulnerabilities Outside the Software

The problems described above may be called structural vulnerabilities. In addition to these challenges, operational vulnerabilities³⁵ for cyber systems will arise from subversion by authorized but malevolent, corrupted or

unwittingly manipulated users,³⁶ by fraudulently accessed credentials or by reconfiguration of the system's hardware while it was being designed, manufactured or shipped ("supply chain" intrusion).³⁷ Moreover, people with access are likely to include not only qualified inside employees, but also vendors performing remote service and those (frequently contractors) charged with maintenance and operation of the system.³⁸ Because a complex system also will use software drawn from third-party vendors, these contractors (as well as their software) may intentionally or inadvertently assist subversion.³⁹ Credentialing, oversight and screening systems for these vendors and contractors are likely themselves to be dependent on hardware and software that could be subverted. And all these problems are compounded by the globalization of supply chains and subcontractor, service and production relationships. For example, more than three-quarters of the field programmable gate arrays in the F-35 are made in China and Taiwan.⁴⁰

In sum, strategies for cybersecurity can reasonably aim to reduce and manage risks by improving security capabilities and practices. But cybersecurity risks cannot be completely or assuredly eradicated. Successful strategies must proceed from the premise that cyberspace is continuously contested territory in which we can control memory and operating capabilities some of the time⁴¹ but cannot be assured of complete control all of the time or even of any control at any particular time. Policymakers must make a judgment about when to intervene and when to allow market forces to determine exposure to this risk. They must also judge how much they are willing to sacrifice efficiency and effectiveness in cyber systems to enhance security.

IV. CYBERDEFENSE AND ITS LIMITS

In general, inherent risk can be reduced or increased by human choices about how a technology is constructed, employed and safeguarded. For example, high-speed automobile traffic is inherently unsafe, but for its benefits Americans annually accept thousands of deaths. Licensing, policing, road design and other efforts kept traffic fatalities to about 50,000 a year toward the end of the 20th century, and a renewed priority then cut this toll to 25,000 annually by the use of seat belts, changes to vehicle materials and structures, campaigns against drunken driving, etc.

Unfortunately, little attention was paid to security when the architecture of digital computing systems and networks was designed.⁴² As a result, efforts to enhance cybersecurity are commonly imposed as additions or modifications to existing systems. Deficiencies are compounded by only incrementally introducing protective measures, with limited data to assess their value and with incentives for each proponent of a security innovation to exaggerate its effectiveness. There is little systemic overview and little incentive for those who do not have legal liability to initiate protective measures.

An overview of protective actions may begin by thinking about the work that an electronic attacker must do.⁴³ Increasing the attacker's workload is likely to decrease the number of successful attacks. The attacker must in many instances identify a target (though, particularly if the goal is economic gain, targets can be attacked randomly). The attacker must then gain electronic access. He or she then needs to acquire credentials, escalate his or her privileges or identify a loophole in the recipient system's software that enables the invader to gain a foothold for executing instructions. He or she needs to employ a set of instructions that can exploit this loophole to gain control over the

data and/or the executable code of the recipient computer. For targets that can defend themselves, the successful attacker must accomplish these tasks with low enough visibility and/or high enough speed to avoid triggering a pre-emptive reaction.

A common means of gaining electronic access is by misleading users into downloading subversive instructions. This is commonly labeled "social engineering."⁴⁴ The loopholes are called "vulnerabilities,"⁴⁵ and the subversive instructions are called "exploits."

Early commercial and government defensive efforts focused on denying unauthorized access. This method has matured and still has its uses,⁴⁶ but it is rooted in an anachronistic intuition to defend information resources as one defends physical spaces. Controls attempt to limit access by using physical systems (cards and keys), passwords, firewalls and the like. These defenses impede unsophisticated attackers or those who seek only soft targets. But they are technically vulnerable⁴⁷ to well-resourced persistent attacks and remain porous because legitimate users (who pass freely through these defenses) so frequently can be made unsuspecting sources of infection.⁴⁸

Defensive effort responded by focusing on the next link in the chain: countering exploits. Anti-viral software presumed a degree of access but sought to recognize attacks, identify their signatures and thwart their code by selective filtering or disabling responses to their instructions.⁴⁹ A closely related effort focused on recognizing vulnerabilities that were being exploited and distributing "patches" to repair them.

But these defenses operate only in a reactive cycle: Attackers discover vulnerabilities; they exploit them; when attacks proliferate they are recognized by vendors that write and distribute code that eliminates the vulnerability or thwarts the exploit; typically this protection against

malware is circulated more than three weeks after initial attacks.⁵⁰

These approaches yield some benefits: The system achieves equilibrium of a sort in the battle against many attackers. But major deficiencies are evident. Attacks are effective until the anti-virus software or patches are issued; in many cases patches are not issued;⁵¹ even when issued, protective software is frequently not installed⁵² or belatedly installed;⁵³ when installed, anti-viral software imports its own vulnerabilities;⁵⁴ cyber attacks and defenses are interactive⁵⁵ – new anti-virus software can quickly be countered by attacker modifications;⁵⁶ and new viruses require new counters, starting the whole cycle all over again (to the profit of both malware creators and anti-virus defenders).⁵⁷

The deficiencies in the existing methods of cyberdefense have been increasingly exposed as state-sponsored and state-run attacks have become more frequent and use more sophisticated and extensive resources.

The deficiencies in the existing methods of cyberdefense have been increasingly exposed as state-sponsored and state-run attacks have become more frequent and use more sophisticated and extensive resources. These attackers can more rapidly exploit the interval before response. And when responses are published, these attackers frequently and readily shift their signatures to work around the new defense. Most significantly, states have the resources and patience to systematically

attempt to discover and stockpile previously unrecognized vulnerabilities.⁵⁸ (These are so-called zero-day vulnerabilities.)⁵⁹

In response to these industrial-strength efforts, defenders have moved to try to uncover and correct vulnerabilities *before* attackers discover them. Closely linked to this are initiatives to modify cyber architectures so they more effectively resist exploits. A number of vendors, government agencies and contractors use supercomputer and other resources to attack their own software and discover its vulnerabilities, so they can then create patches before others recognize offensive capabilities.⁶⁰ In addition, markets have developed where vulnerabilities are sold to either defenders or attackers at varying prices according to their perceived power. Software vendors, government contractors and criminal groups have all been reported as buyers.⁶¹ Some buyers use their acquisition to patch a system and destroy the vulnerability. Others acquire it to exploit it. As one astute government observer puts it, “Bugs become good or bad depending on whether they are exposed to light [for defenders to patch] or darkness [for attackers to exploit].”⁶²

In an even more proactive effort, the Defense Advanced Research Products Agency (DARPA), Microsoft and others have developed architectural modifications so that whole classes of vulnerabilities are eliminated or attackers’ tools confounded by random variables and other techniques.⁶³ Address space layout randomization (ASLR), use of random canaries to protect against buffer overflow, data execution prevention (DEP) and sandboxing systems that limit privileges for intruders are representative.⁶⁴ The first two are exemplary of “moving target defenses” – they make the cyber system less predictable and therefore less exploitable by attackers. ASLR does this by randomly assigning and repeatedly changing internal addresses for key data and functions so attack software cannot readily

locate them.⁶⁵ Exploits taking advantage of a long-established class of vulnerabilities from buffer stack overflow were eliminated *as a class* by introducing a random number variable (a “canary”) that attackers cannot readily discern.⁶⁶ DEP, embedded in software, hardware or both, prevents code being run from some memory regions.⁶⁷ Sandboxing adds a new problem for attackers: Their penetrations do not grant them broad access to an operating system.⁶⁸

Many parts of these cybersecurity efforts and markets are so obscure that it is impossible to describe the offense-defense balance with great confidence. Appendix I describes why the data are so limited and distorted. The broad assessment of sophisticated observers is that increased and proactive cybersecurity efforts by vendors and governments are yielding rewards.⁶⁹ Moreover, as bugs are weeded out and later versions benefit from more robust architectures, vulnerabilities in established operating systems and applications are becoming harder to identify and exploit. These gains are exponentially powerful in combination.⁷⁰ One researcher describes this effect:

[E]ach additional defense requires [an attacker] to find another vulnerability. In the case of an application running DEP, ASLR and inside of a sandbox I will need to find around four vulnerabilities: the initial vulnerability, a vulnerability that lets me read/predict the stack canary, a vulnerability to read how ASLR has randomized the addresses, and a vulnerability to break out of the sandbox. ... One “exploit” will actually require four different vulnerabilities to be exploited. Modern exploitation requires exploit chaining.⁷¹

However, while improvements have reduced the number of vulnerabilities and mitigated exploits, they have not eliminated vulnerabilities or totally prevented exploits.⁷² Cyber systems remain open to attack from those with the resources and the will to discover vulnerabilities and to exploit them. As in

In an even more proactive effort, the Defense Advanced Research Products Agency (DARPA), Microsoft and others have developed architectural modifications so that whole classes of vulnerabilities are eliminated or attackers’ tools confounded by random variables and other techniques.

all security domains, there is a constant interaction between defense and offense, with improvements in one provoking a response in the other.⁷³

While these contests are occurring in mature cyber systems, new systems and new applications of information technology open new vulnerabilities. This is partly because they expand the targets for attack, partly because novel applications are likely to have unforeseen security weaknesses and partly because the ambitious miniaturization associated with many new applications (for example, devices worn or inserted in the body) makes hardware support for security more costly than in larger equipment. Unfortunately, it is also because the developers of new applications are succumbing to the same pressures and temptations that afflicted their predecessors. Immediate priorities to get to market quickly, economically and effectively cause security investments to be disfavored.⁷⁴ Though a bottom line can only be drawn intuitively and imprecisely, it is reasonably clear: We are increasing our vulnerabilities faster than we are closing them.

A last group of defenses assumes penetration and vulnerability. Under the rubrics of “detective controls” and “active defense,”⁷⁵ these protective

efforts aim to thwart exploits and data exfiltration by automated scrutiny of ongoing operations, identification of suspect instructions and data flows, and isolation or nullification of malevolent attempts at manipulation.⁷⁶ It will be seen, however, that these valuable initiatives have limitations similar to the previously described defenses. Though active defense benefits from operating closer to real time, it depends on signature recognition or recognizably improper data movements. Insecurity remains from zero-day attacks, sleeper exploits that do not involve data exfiltration, shrewdly disguised exfiltrations, etc.

Defensive efforts in major mature systems have grown more sophisticated and effective. However, competition is continuous between attackers and defenders. Moreover, as new information technologies develop we are not making concomitant investments in their protection. As a result, cyber insecurities are generally growing, and are likely to continue to grow, faster than security measures. Especially protected high-value IT systems can be better-secured. But even in these systems, improvements simply reduce vulnerability; they do not eliminate it.⁷⁷

V. THREE OTHER CAUSES OF CYBER INSECURITY

Three other factors powerfully contribute to cyber insecurity.

1. CYBERSPACE IS ADVERSARIAL, CONTESTED AND CROWDED TERRITORY. OUR ADVERSARIES (CRIMINALS, MALEVOLENT GROUPS, NUMEROUS OPPOSING STATES) CO-EVOLVE WITH US.

Co-evolution is a familiar experience for national security decision-makers. The United States encountered it in its contest with the Soviet Union as the Soviets built a tank or tactical aircraft and the United States countered with a more modern version or defense system and the Soviets countered again. U.S. forces in Iraq and Afghanistan adapted as terrorists evolved improvised explosive devices, terrorists responded to the adaptations and we adapted again. But the number, diversity, low visibility, extent and speed of interactivity of actors in cyberspace are unprecedented. When discovered, attacks lead quickly to imitation and defenses are constantly probed, both randomly and against selected targets. In this hothouse environment the pace of competitive evolution is unprecedented. Even successful defensive efforts are soon tested and often subverted or circumvented. A recent National Academy of Sciences study rightly concluded: “[C]ybersecurity is a never-ending battle. A permanently decisive solution to the problem will not be found in the foreseeable future.”⁷⁸

2. LIMITS OF U.S. GOVERNMENTAL POWER.

American national security policymakers came of age during decades in which American security presumed the probability, if only we invested properly, of American technological dominance. This advantage came as a consequence of the United States’ emerging from World War II with little battle damage, a large industrial establishment closely tied to military priorities, and half of the world’s gross domestic product.

It was sustained for three-quarters of a century through repeated, though often erratically stopped and started, investments of money, energy and talent. With only occasional exceptions,⁷⁹ U.S. technical dominance of the Soviet Union was broadly achieved and fueled a Cold War victory. American technical superiority has never been questioned with respect to terrorists.

Cyberspace is adversarial, contested and crowded territory. Our adversaries (criminals, malevolent groups, numerous opposing states) co-evolve with us.

Though the United States dominated information technologies in the later 20th century and retains substantial cyber advantages today, this ingrained premise cannot indefinitely apply to IT. Cyber skills and resources have massively proliferated outside the United States. IT is much more fluid, egalitarian, distributed and dynamic than the technologies encountered when the United States, for example, succeeded during the last half-century in controlling the high seas and airspace over battlefields. However strong U.S. cyber capabilities may be, they cannot credibly be premised on enduring dominance and they cannot presume that the United States will not be the victim of technological surprise.

Moreover, the U.S. government cannot control threatening cyber behaviors using traditional tools intended to shore up national security within the geographic boundaries of America. Many cyber attacks originate on American soil, but most do not. The United States has little ability to create cyber boundaries at its borders.

Finally, the defensive power our nation now has rests significantly in private hands. Both defensive capability and innovation are controlled less by the priorities and pace of the Department of Defense (DOD) than by private sector markets, marketing cycles and commercial priorities.

3. SPEED OF CHANGE

No technologies have ever spread so fast as cellphones and digital information systems. The first transistors (the hardware that enables computer memory and processing capabilities) were sold 60 years ago. Global production is now estimated at a rate of 8 trillion transistors per second.⁸⁰ During the last four decades, driven by Moore's law, innovative companies and intense consumer demand, the world has moved from mainframes to distributed desktops to personal computers to mobile systems, to the cloud and to the "Internet of Things." The Internet, a complex emergent communication system that has never been "a static thing," grew from some 16 million users in 1995, to 880 million in 2005, to 2.7 billion in the spring of 2013.⁸¹ Its communicative capability has penetrated financial, military, industrial and social systems, increasing their cyber dependency at an unprecedented rate.

It is not surprising that market, regulatory, bureaucratic and legacy technical systems that might over time adapt to provide more security have been too slow and too reactive to be effective. Even markets – much the fastest of these systems – have a cycle time that cannot keep pace. To offer a distant security analogy, when gunpowder was introduced in Europe it set in motion changes that outmoded personal armor, traditional fortresses and battlefield and naval doctrine. Responses evolved to include new approaches to military mass and mobility, new ideals of leadership (to replace outdated concepts of individual heroism and chivalry), new systems for training officers and men (to include, for example, an understanding of ballistics), new architectures for defensive systems, new industrial

establishments to produce guns and munitions and new relationships between governments and those governed (to facilitate and respond to conscription into mass militaries). Digital information systems and their remarkable child, the Internet, are no less powerful, novel or proliferated than gunpowder. They demand comparable changes. But while gunpowder technology took approximately two centuries to develop and spread, cyber technologies have proliferated in approximately two decades – exponentially faster.

If problems of cybersecurity are seen simply as failures to replace old security regimes with new ones, any progress is likely to be quickly outdated by changes in the technologies one is trying to secure. Plausible paradigms for the past cannot be reshaped quickly enough to cope with a new present. New security approaches must have a plasticity that permits them to cope not only with what is, but also with what unpredictably will come to be.

In sum, the United States' defensive plans must recognize some unfamiliar limits. First, competition and adversaries' offensive co-evolution create currents that impede U.S. progress toward cybersecurity. Second, though premises of American technological superiority are deeply ingrained in the nation's security establishment, the United States cannot presume American dominance of cyber technology. In this domain, America should not expect that it would be able to avoid technological surprise. Third, U.S. government power will be constrained by the facts that the private sector is the engine of innovation inside this country and that extensive capabilities are located outside our borders. Moreover, the novelty, speed and unpredictability of information technologies make them particularly difficult to regulate, orchestrate through long-term research programs and defend against. The limits of American governmental power will be felt especially severely in this widely proliferated and extraordinarily diverse technology.

VI. CREATING A MINIMAL SHARED CYBERSECURITY STANDARD AND SOME RECOMMENDATIONS

The Standard

Because IT dependency and concomitant insecurities have come so quickly, the United States lacks a shared understanding of acceptable and unacceptable risk and of the proper roles of the federal government and the private sector. Decision-makers and citizens typically accept a normal level of familiar risks – for example, the risks to shipping in the Persian Gulf or the murder rate in a U.S. city. Without much reflection, those norms often become implicit standards. The nation invests to maintain or improve (usually only marginally) a status quo about which we have a common understanding. Typically, U.S. local and national leaders are animated to change their approach or greatly increase investments only when confronted by dramatic changes such as a crime wave, a terrorist attack or a beneficial event such as the breakup of the Soviet Union. No such standard is now established for cybersecurity. The absence of a standard cripples efforts at consensus and therefore disrupts strategies, undermines legislative proposals, makes budget allocations difficult to size and defend, etc.

Though there is a good case for doing much more, it ought to be possible to articulate at least a minimal standard of protection that would garner wide support and provide an anvil against which to hammer out agreed programs for action. Approaching the issue from a national security perspective, I suggest this: The United States cannot allow the insecurity of our cyber systems to reach a point where weaknesses in those systems would likely render the United States unwilling to make a decision or unable to act on a decision fundamental to our national security.⁸²

The suggested standard implies, for example, that if we thought an opponent could use cyber tools to

render the U.S. nuclear arsenal impotent, or to turn the country's missiles back upon the United States, then we would be unable to act to protect our interests. In this case, we would judge ourselves to be intolerably insecure in cyberspace.⁸³

It should be observed that while commentators often address catastrophic or existential risk, this standard focuses our attention on the likelihood that what other nations or terrorist groups would seek from a sophisticated cyber attack is to deter us, not to destroy us. The cyber threat to our national security is less likely to be from vandalism than from efforts to inhibit our freedom of action. The national security case for protecting critical cyber systems is that this is essential to maintaining our will and capacity to act to defend vital American interests.

This proposition is minimal. Much more ambitious standards could be articulated from a national security perspective and from economic, crime prevention and other perspectives. For example, the United States could pursue the ideal of permitting little or no penetration or corruption of military cyber systems, or of assuring that military equipment was not subject to cyber sabotage. In the domestic arena America could establish goals for the minimization of identity theft, loss of intellectual property or disruption of service.

The minimal national security standard advanced here is chosen as the focus of the observations and recommendations that follow for three reasons. First, this is indisputably a responsibility of the federal government. If this standard is not being met, the resulting weaknesses cannot simply be left to state and local or private decision-making. Second, it has the best chance of securing assent. If the United States cannot agree to respond to any other weaknesses of cybersecurity, it would seem the nation could at least agree to try to address this one. Third, because it is so minimal, this is a most illuminating case. Whatever needs to be done to

meet this standard should be considered for more ambitious standards. Whatever we accomplish can provide a model for these greater efforts.⁸⁴ More ominously and significantly, deficiencies that undermine the abilities to meet even this standard point to weaknesses likely to demand attention in other cases.

The United States cannot allow the insecurity of our cyber systems to reach a point where weaknesses in those systems would likely render the United States unwilling to make a decision or unable to act on a decision fundamental to our national security.

What does the United States need to do to meet this standard? The federal government is doing many useful things to strengthen cyber capabilities. The following recommendations do not attempt to describe or evaluate these. The recommendations are not an effort at a comprehensive program. Rather, they advance propositions about where the logic of the first parts of this paper suggests U.S. government strategies should be strengthened or altered and how that can be accomplished.

1. For critical U.S. government systems, presume cyber vulnerability and design organizations, operations and acquisitions to compensate for this vulnerability. Do this by a four-part strategy of abnegation, use of out-of-band architectures, diversification and graceful degradation. Pursue

the first path by stripping the “nice to have” away from the essential, limiting cyber capabilities in order to minimize cyber vulnerabilities. For the second, create non-cyber interventions in cyber systems. For the third, encourage different cyber dependencies in different systems so single vulnerabilities are less likely to result in widespread failure or compromise. And for the fourth, invest in discovery and recovery capabilities. To implement these approaches, train key personnel in both operations and security so as to facilitate self-conscious and well-informed tradeoffs between the security gains and the operational and economic costs from pursuing these strategies.

The early parts of this paper show why cyber vulnerabilities cannot be completely eradicated or assuredly detected. How should the United States government respond to this modern weakness – the poisonous side of 21st-century dependencies on IT? A first answer is to limit our diet. The United States cannot effectively function in the modern age without IT capabilities, but in critical systems steps should be taken that forsake some efficiencies, speed or capabilities in order to achieve greater security.

IT systems typically offer more functional capabilities than are required for their missions. These excess capabilities come bundled in common architectures or seem desirable to custom designers because they may be valuable for unlikely but possible needs. A strategy of abnegation is founded on the presumption that critical systems should be supported by cyber capabilities that are no more extensive than required to perform their core mission.⁸⁵ Enabling a system to perform unnecessary functions unnecessarily increases cyber vulnerabilities. (In technical terms, additional functions increase “attack surfaces.”) Unnecessary functions also expand abusive opportunities for whoever can capture control of the system. Users also should be sharply

differentiated in the capabilities accessible to them. The development of the cloud, with its ability to provide “thin client” capabilities that authorized users need, but only what they need, should assist this priority. It is neither necessary nor wise that in critical security situations most users have most of the range of capabilities available to them.⁸⁶

Second, because cyber systems cannot assuredly be reliable, a strong presumption should be created that critical systems integrate non-cyber safeguards in their access and operation. Non-cyber components, also described as “out-of-band” measures,⁸⁷ can include, for example, placing humans in decision loops, employing analog devices as a check on digital equipment and providing for non-cyber alternatives if cyber systems are subverted. The digital information systems that drive information technologies import only digital vulnerabilities. We can protect ourselves by forcing attackers to cope with system attributes that are outside the reach of computer code.

Critical systems should be designed to degrade as predictably and gracefully as possible. An important way of doing this is through diversity.⁸⁸ For example, if within the U.S. our missile arsenal the Navy and the Air Force employ different cyber systems, a vulnerability in Navy missile command and control is not as likely similarly to affect Air Force missile command and control. In the language of computer experts, investments need to be made to avoid common mode failure. This principle applies especially to assuring that safety systems, security systems and operational systems are strongly differentiated so operational failures and compromised systems are checked rather than replicated by safety and security systems.

Principles of diversity and out-of-band systems should be familiar from the construction of the nation’s nuclear arsenal, built as it is to offer a triad of alternatives and varied mechanisms of

command and control. But the recommended presumptions are counter to present trends in cyber acquisitions, incentives and architectures.

Finally, more attention needs to be given to the loss or subversion of a critical cyber system. A goal of graceful degradation is facilitated by diversity, but detection programs, restoration programs and workarounds need to be well-considered and prepared for in advance.⁸⁹ Often the critical variable will be not the security of a system but rather the remaining efficacy of a network of systems.

The strategies of abnegation, use of out-of-band measures, diversification and preparation for graceful degradation are now recognized and partially implemented, but their systematic pursuit and prioritization would change present practice. This change would not be easy to implement. Narrowly designed systems are more expensive than commonly marketed, broadly constructed ones. Integrating out-of-band measures is easier said than done and can cause operational degradation through delay or increased likelihood of error. Diversity of systems increases costs and complications for defenders⁹⁰ as well as attackers.⁹¹ Graceful degradation is an expensive and complicated goal.

Because cost and administrative benefits exert such a strong hydraulic force in everyday operations, policymakers need to define and designate critical systems and create presumptions in favor of these protective strategies. They and key subordinate operators also need to understand and choose among the tradeoffs. In critical systems, choices for greater cybersecurity or greater cyber insecurity are too consequential to be left to technologists or to be resolved by project budget calculations. Conversely, they are too important to be left to those whose training and responsibilities predispose them to prioritize present operations at the expense of longer-term and probably

ill-understood security risks. Decision-makers need a general understanding of *both* operations and cyber technology. Operators and those responsible for cybersecurity need a particularized understanding of both for the activities they control. This paper aims to provide a basis for the generalized understanding. U.S. military and civilian systems need to proliferate cross-training to provide the particularized understandings required.

2. Identify critical nongovernmental systems in light of the minimal standard, publicly explain why some of these systems pose national security risks and use incentives, standards, attack information and regulatory authorities to improve cybersecurity in these systems, including through the strategies described in Recommendation 1. In doing this, consider not just present risk but also areas of increasing risk and broaden discussion to include endemic risk – the long-term effects of loss of intellectual property and competitive position from cyber penetration.

The distribution of cybersecurity responsibilities among public and private-sector actors has not been established. There is ample room for debate, but as the government moves to get its own house in order, it should be evident that the minimal principle applies to private-sector activities (for example, the nation's power grid or the Internet itself) that are critical to the ability to maintain national security. American policymakers should articulate the minimal standard, emphasize that it includes private-sector systems, explain how these are at risk and then use the range of persuasion, incentives (including U.S. government acquisition preferences) and regulatory tools to improve these systems' resistance to catastrophic attack.⁹²

This instinct underlay a requirement under a presidential executive order issued in early 2013 that the Department of Homeland Security

(DHS) assess industries and companies "where a cybersecurity incident could reasonably result in catastrophic regional or national effects on public health or safety, economic security, or national security."⁹³ The resulting classified report admirably advances understanding but leaves the United States short of where it needs to be.

The shortfall occurs in three dimensions. First, the executive order aims at illuminating only acute risk. Enough has been observed and documented to make it evident that a further important hypothesis needs to be advanced and evaluated. It is that the characteristics of the new cyber age and the mercantilist priorities of other nations, particularly (but not exclusively) China,⁹⁴ combine to generate an effect that neither alone would achieve: large-scale, low-visibility, highly consequential campaigns to abscond with business secrets that can change the balance of economic power between companies, industries and ultimately nations.

This hypothesis may be called one of endemic risk and is not so salient or likely to be agreed upon as the acute risk recognized in the minimal standard. This is because it occurs in an area of blindness for the U.S. government. Unlike China (and a number of other states), whose government gives priority primarily to economic competition and secondarily to security competition, the U.S. government gives priority to security and regards economic competition as dominantly a private-sector concern. Concomitantly, America clearly differentiates government and private-sector actions and actors, while China and other states do not see such a sharp bifurcation.

These perspectives leave the United States ill – equipped to blend economic and security concerns (though there are rhetorical and symbolic steps in this direction).⁹⁵ The blinkered focus on traditional security causes cyber attacks to be viewed predominantly in terms of the loss of

classified or militarily sensitive information. Conceptually and sometimes organizationally, cybersecurity is then treated as a subset of counterintelligence. When cyber attacks are considered in the business context, losses are typically assessed individually and economically. The cumulative national consequences of economic loss are rarely considered.⁹⁶

Endemic loss demands its own evaluation. A satisfactory effort will require consideration not only of the magnitude of the data exfiltrated but also how it is used and with what consequence. It may be that loss is less consequential than presumed because data may frequently need to be matched with operational skills (“tacit knowledge”) to be used in practice. On the other side of the coin, data thefts may be paired with non-cyber actions in a larger campaign that is not comprehended if each instance is considered only as an individual transaction. An assessment of these variables is very different from what was required by the executive order but is a necessary complement to it.

Second, the United States does not have plans for acting to reduce either the acute risks recognized in response to the executive order or the endemic risk not addressed by the order. For the acute risks, a reasonable hypothesis is that critical private systems need to implement the strategies applicable to governmental systems: creation of lean systems, integration of non-cyber controls, maintenance of variance in cyber systems and preparation for graceful degradation. Broad action on the endemic risks needs the predicate of assessing and publicizing the magnitude of these risks as recommended below. But planning should be initiated now for a robust strategy of counterattack from an array of existing governmental powers if that risk is shown to be substantial.⁹⁷

Because industries greatly vary in their incentives and disincentives, degrees of concentration,

resiliency, cyber budgets and cyber sophistication, action plans need to vary industry by industry.⁹⁸ They also need to be accepted, indeed championed, by relevant oversight agencies, and this oversight needs to be supported by Congress. This requires declassification of important parts of the DHS study and strong White House leadership to articulate and act on its findings. And, as just suggested, it requires much more White House attention to, and declassified explanation of, the endemic risk.

To a student of security studies, it is an evident deficiency and oddity of cybersecurity that “the threat” is neither comprehensively nor publicly well described. Only piecemeal descriptions of attacks are available, many of these are published by private actors, integration of public and private information is poor and classification presents a large barrier to public explanation.⁹⁹ The United States cannot develop good strategies without a broadly shared understanding of its problems. Ignorance, badly integrated information and classified materials do not provoke action in a democracy or profit-driven marketplace.¹⁰⁰

The two most notable efforts at providing a base for public discussion – a 2010 National Intelligence Estimate and a detailed 2013 Mandiant case study of a group of Chinese cyber attackers – suggest both the rarity and the value of such efforts. Richly detailed Kremlinology during the Cold War was followed by considerable efforts to comprehend and characterize terrorist groups in the first decade of this century. But the Chinese cyber problem was regularly pushed aside, and public descriptions of other cyber activities have either been left to journalists or largely ignored.¹⁰¹

Third, policymakers need to recognize that though there is great value in the present snapshot of private-sector critical vulnerability, the greatest risks and opportunities inhere in the rapid rate of change in IT integration into industries. The

discussion in Section V above indicates that the speedy growth of cyber dependencies means that tomorrow's risks will often be different from today's. On the other side of the coin, the imperative need and opportunity is not only for reformation of present systems but also for the design and implementation of new systems as industries evolve. Some industries that are not now critical risks will become so as they become more cyber dependent.¹⁰²

3. Tolerate private-sector decisions about cyber risk in systems that are not critical or are adequately resilient despite their cyber risk . A perfect world would have no cyber insecurity anywhere, but public-sector pursuit of that goal is unrealistic, distracting and drains resources.

Security analysts are by nature insecure. There is a great risk that in seeing cyber vulnerabilities from software, hardware and humans everywhere, policymakers concerned with security will attempt to regulate everything. This could stifle the nation's greatest source of innovation and lead to a reduction in security from heavy-handed, ill-conceived or ineffective regulation and from costs (including costs in national power) that outweigh benefits. It has been usefully observed that, for example, Google Glass, automobiles connected to the Internet, health care devices¹⁰³ and health care records of all kinds are subject to hacking, with potentially destructive results. However, unless these results can scale to levels that are societally catastrophic, they are unlikely to have national security effects and responses can be left to consumer and producer priorities or, in cases of life, death or grave injury, to ordinary regulatory procedures. Most cyber security problems are not national security problems. It would be counterproductive to treat them as though they were.

4. Recognize that regulatory regimes rapidly become outdated, so wherever possible use

incentives and education to improve security. Where regulation is required, tie it as much as possible to ends (responsibility for achieving a desired state of security) rather than means (particular required process improvements).

Regulation has a strong tendency to be static. The cyber world is dynamic. Consequently, regulation should be an instrument of last resort and when employed should specify ends (for example, resilience sufficient to avoid catastrophic failure after successful cyber attack), leaving maximum freedom for choices of means. Incentive strategies are necessarily complex and varied: Different industries and companies within industries are incentivized by different potential offerings from government. Some will be attracted by shared information about attack vectors and defenses. Others may be animated by joint research and development projects. Still others will be animated by tax incentives or grants. This, like other recommendations offered here, underscores the need for varied strategies administered by the government agencies with the greatest understanding of the particular industries they are attempting to affect.

5. Seek to develop norms and restraints limiting the use of cyber weapons to achieve effects outside the cyber domain. As a first step, initiate efforts to buttress a fragile norm that appears to exist between Russia, China and the United States: Apparently, we have not used cyber as a means of physical attack against one another. Articulate a norm of renouncing cyber attacks on civilian infrastructure and discuss this goal with China and Russia even if other kinds of cyber conflict with these countries continue or intensify. As a second step, evaluate whether uncertainties inherent in the cyber foundations for nuclear command, control and warning may erode confidence to a point where the present regime of mutually assured destruction risks degradation into a regime of MUD. If that

alternative is mutually seen to be more insecure, pursue agreements to refrain from cyber intrusions into nuclear command, control and warning systems.

Discussion about deterrence is an important part of trying to bring greater security to the cyber domain. Deterrence in this realm is much desired (because technical defenses alone will not assure protection)¹⁰⁴ and plausible through cyber counterattack (attackers are no more secure than defenders) but difficult to tailor (because of difficulties in attack visibility, characterization and attribution) and challenging to constrain against uncontrolled escalation.¹⁰⁵ Debate now focuses on how these goals are not currently achieved, whether and how cyber attacks might be better deterred by using cyber or other means¹⁰⁶ and the extent to which analogies drawn from nuclear deterrence and control might or might not apply to cyber weaponry.¹⁰⁷

Discussion of international restraints and confidence building is also occurring but is less robust.¹⁰⁸ This recommendation points to neglected opportunities for identifying and strengthening nascent norms around which there are common interests.¹⁰⁹ The recommended journey is neither easy nor certain, but it can yield important rewards and is an important complement to technical protections and deterrent efforts.

Three observations suggest reasons for investing in this effort. The starting point is to recognize that some cyber stability now exists, but it is fragile and inchoate. Present restraints on the use of cyber weapons are an emergent characteristic of complex technological, economic and political interactions rather than a result of agreement between potential adversaries. A priority should be to set stability on a foundation of deeper and broader understanding between nations. This effort can draw inspiration from the fact that some nuclear strategic stability was established before it was well understood.

Policymakers and theorists quickly recognized its importance, but it took experience, great tension and risk,¹¹⁰ maturation and discussion for the United States and the Soviet Union to move together from an implicit and precarious situation to a more stable state.

Second, while commendable effects at building cyber norms have focused on efforts inside cyberspace, better possibilities for initial consensus may arise from agreements limiting cyber actions with consequences outside the cyber domain. China, Russia and the United States have common vulnerabilities that make it in their mutual interest to try to limit cyber-physical conflict.¹¹¹

Third, if these three specially strong cyber powers can identify unacceptable behavior against one another, it may dampen third-party engagement in this behavior by making misattribution less likely and by prompting a broader norm against this misconduct. During the last two-thirds of a century, national security issues have typically focused on single identified enemies. The analysis in Section III points in a different direction. Though the intent and capabilities of particular opponents are still relevant, cyber risks are pervasive for both the United States and its opponents because they are inherent in the technology.¹¹² Recognizing this can point to possible common ground for agreement.

To develop and then evaluate opportunities for agreement on cyber norms, two priorities should be pursued.

Attempt to enhance cyber strategic stability with Russia and China by discussing agreed constraints on the use of cyber weapons against civilian physical targets. Even as the United States and China quarrel about cyber espionage against each other, there is an unspoken arena of strategic stability: Apparently neither country has used cyber mechanisms to conduct physical attacks against the other.

This restraint is vulnerable. It is unarticulated and occurs in a context of apparent contrary behaviors involving other countries. Most notably, the United States and Iran have reportedly engaged in cyber attacks against each other's civilian infrastructure.¹¹³ The Chinese, Russian and American restraint with respect to one another is as if, during the Cold War, the United States and the Soviet Union attempted to maintain nuclear stability with each other while reportedly engaged in tactical nuclear exchanges with third parties. Schizophrenia is not a recipe for stability.¹¹⁴

Moreover, while nuclear weapons are held by few states and highly controlled within those states, cyber weapons are not confined to states or well-controlled by them, particularly within and outside the Chinese and Russian governments. Compounding this problem, others who might engage in destructive cyber attacks could willfully or inadvertently create the false impression that one of these three nations conducted these attacks.¹¹⁵

The United States should initiate discussions¹¹⁶ with China and Russia about the delineation of mutually agreed red lines.¹¹⁷ Any resulting understandings should be buttressed with confidence-building measures and escalatory speed bumps.¹¹⁸ One area of agreement could be against the use of cyber weapons to damage civilian systems.¹¹⁹ Agreement could then be used as a foundation for developing other understandings that could limit cyber misbehavior in these relationships.¹²⁰ An agreement with either country could pave the way for agreement with the other and facilitate other cooperative efforts to limit destructive cyber attacks involving other nations and subnational groups.¹²¹

The movement from MAD to MUD. If cyber insecurity is seen to be prevalent, particularly against attacks by major states, then it is likely to affect confidence in nuclear arsenals and related

command, control and warning systems. Emphasis on security-enhancing initiatives such as those urged in the first recommendation may mitigate these effects, but questions will still be raised about whether an assured retaliatory capability is really assured for the United States and for other major powers.¹²²

Deterrent theorists and policy-makers should begin to consider now how this might affect stability within the established framework of MAD. Detailed studies about vulnerabilities will necessarily be classified, but an open discussion can usefully consider how cyber uncertainties on both sides of major state conflict may diminish nuclear deterrence. For example, second-strike nuclear capabilities may be questioned if they rest on uncertain foundations of cyber command and control systems or are dependent on warning systems that can be spoofed. If so, a world of MAD may be replaced by one of MUD.

The effects of this change are not unidirectional. Mutually unassured destruction may have deterrent effects if neither side will be sure of its ability to execute a first strike as envisioned. But on balance MUD is likely to be less stable and less attractive to Russia, China and the United States. If that is the case, agreements to forswear intrusions into one another's nuclear command, control and warning systems may be mutually desired and achievable.¹²³

Progress toward these agreements will be complicated by other kinds of cyber conflict and by ambiguities about the boundaries of protected systems and definitions of proscribed and permitted conduct.¹²⁴ Many will regard difficulties, even impossibilities, of verifiability and attribution as fatal problems. Other forms of arms control and confidence-building depend on the visibility, verifiability, state control of the relevant weapons and attribution, but all these characteristics will be impaired in cyberspace.¹²⁵ There will be questions

To improve performance, U.S. technology investments have to be better complemented with investments to understand those who are attacking the United States – and to understand ourselves.

about nations' self-interests, with some urging that a perceived cyber attack advantage should not be surrendered even if it might be transitory, unreliable or ultimately destabilizing.¹²⁶ Perhaps most fundamentally, even if all parties see an accord as desirable, the well-known phenomenon of "prisoner's dilemma" – the inability to trust the others enough to find a common improvement – may keep the United States from reaching a better state. But the opportunities are significant and our diplomatic, technical and national security skills should be employed in trying to realize them. The alternative is a steady degradation in strategic stability as cyber risk expands through the civilian infrastructure and we move to a world of unstable MUD.

6. Because cybersecurity is not just a technical problem, invest in socioeconomic research to understand private industry and personal incentives, inhibitions and options for improving cybersecurity. Invest similarly to better understand the behaviors, incentives and inhibitions of government, criminal, ideological group and individual cyber attackers. Systematically evaluate foreign government responses to cyber insecurity and assess whether lessons learned abroad could be applied here.

Though cyberspace is a battlefield and cyber technologies are weapons and means of defense, as with all security issues the critical actors are people and the critical determinants of their behavior are political, economic, psychological and sociological variables. Accordingly, to understand and shape cyber defenses, the United States needs to comprehend and attempt to influence the cultures, perceptions, incentives and disincentives of defenders and attackers. Without sound diagnoses the United States cannot offer sound prescriptions and is prone to monotonic approaches, prescribing the same remedies for different situations. To improve performance, U.S. technology investments have to be better complemented with investments to understand those who are attacking the United States – and to understand ourselves.

Because we imperfectly comprehend the IT activities and incentives of most American businesses and individuals, we are disconcerted by the persistent failures of many to invest in cybersecurity in proportion to security professionals' perceptions of risk. But the challenge is not simply one of documenting or dramatizing the risk. For example, inquiry is likely to show that most successful exploits are against vulnerabilities that are already patched, but patching in most American companies and government agencies occurs only after significant delays. These delays are not just a consequence of indifference or funding failures, but often are caused by procedures designed to cope with concerns about interoperability, security and downtime. (Patching is dis-incentivized if, for instance, a vendor is penalized for downtime. Patching is also risky if done piecemeal and costly if systems are taken offline.) Better understanding of these priorities and procedures may yield improvements that are not visible if a purely technical approach is taken to security.¹²⁷

Studies of this kind would also illuminate how one size does not fit all in prescriptions

for cybersecurity. For example, the cyber circumstances and sensitivities of brokerage firms and credit card companies are very different from those of companies that generate and distribute wholesale power.¹²⁸ The former are necessarily very open to and engaged with the public; their assets and risks are dominantly in the cyber realm; they accept some cyber fraud as a cost of doing business but place great emphasis on immediate response when attacks cumulate beyond accepted margins; to maintain this posture, their software is rapidly and frequently modernized, heavily financed and well understood. The latter companies are less integrated with the public; their assets and perceived risks are primarily physical (for example, generators and their risk of breakdown); they do not have much history of cyber attack; they commonly operate with antiquated and poorly documented software; key software components are at best upgraded once a year during an annual shutdown; slow cycles of capital investment and heavy regulation of capital recoupment impede software investment; and the diversity of their physical systems and cyber connections creates difficulty in implementing defenses but also provides them with a measure of protection.

These circumstances cause very different interactions with government. Financial firms commonly know as much as or more than government officials about their daily risks. Constantly stressed, these firms have evolved to live in an environment of persistent cyber attack. But their adaptation is to everyday life and they may understate their exposure to catastrophic black swan risks, especially risks that affect the Internet or society as a whole. They want very precise information about attackers and attack signatures, are unlikely to be attracted by government capital and fear any regulatory intervention. Power companies have not been similarly stress-hardened to cyber attack. They need a richer understanding of the cyber

vulnerabilities of their own systems, capital to address these issues and assistance persuading state regulators to accept these investments in their rate base. In both industries (and probably all industries), larger firms are more knowledgeable and better defended than smaller ones, and government programs that may be appropriate for one size or type of enterprise may be ineffective or counterproductive for another. To understand these and an array of other variables, the government should invest in industry studies that integrate technological and business perspectives.¹²⁹

Similarly, the United States needs to study its opponents and the illicit eco-systems in which they operate. These include individual hackers, loosely and tightly organized groups asserting that they are motivated by the public good,¹³⁰ criminals acting as individuals or as members of cartels, business entities and criminals acting with state support or encouragement¹³¹ and nation-states. Resources, skills, relationships and markets vary significantly within each of these categories, but capabilities tend to be greater and better masked the further one moves along the list.¹³²

Technical analyses of attack paths and signatures abound.¹³³ But composite insights and understanding of business models are less available.¹³⁴ For example, lower-end attackers do not normally discover new vulnerabilities or exploits. Instead, they use established attack tools. These are purchased from open markets¹³⁵ and dark markets¹³⁶ or derived from patch announcements from vendors, accounts of state-initiated exploits and publications trumpeting vulnerabilities that “white hats” have discovered.¹³⁷ Furthermore, low-end attackers commonly pursue targets en masse and rarely pursue a single target in isolation or for extended periods.¹³⁸

By contrast, high-end attackers invest substantially in R&D: They investigate and stockpile

vulnerabilities. When they deploy exploits they do so against an individual entity or a class of targets and pursue these targets for extended periods, earning them the sobriquet of “advanced persistent threat.”¹³⁹ An early assessment described the modus operandi:

[t]he attackers selected the data for exfiltration with great care. Though they had the opportunity, they did not simply “take what they could get” and leave, rather, they chose specific files, often ignoring related information in adjacent directory locations, activity which suggests these attackers were disciplined and operating from a specific list of collection requirements, a characteristic usually only found in highly professional operations.

During the incident ... the attackers did not open and review file contents – though they had the required file permissions – but instead navigated immediately to the files or folders they wanted and began the steps necessary to exfiltrate them, suggesting that they had reviewed the directory contents offline and that they had already gained access to this firm’s network to conduct detailed reconnaissance, including the possible exfiltration of file directory listings.

These types of operational techniques are not characteristic of amateur hackers operating in widely dispersed geographic areas.¹⁴⁰

Metaphorically, the United States may think of the lower-end actors as like open-sea fishermen who harvest their prey with nets, value volume and adopt techniques that minimize their costs. High-end actors are like fly fishermen motivated by factors beyond narrow economic calculation as they shape hooks and craft lures in the patient pursuit of individual targets. In the language of security, low-end actors can be countered by vigorous self-defense, police work and strategies (including punishments) that raise their costs.

High-end actors demand the attention of the national security establishment because they can do substantial damage, are relatively insensitive to their own costs and often have security-related espionage goals.¹⁴¹ The United States has not, however, adequately invested in such things as understanding how high-end attackers set goals and use information; how their economic and security motivations intertwine and are prioritized and deconflicted; and the extent to which such efforts are centrally directed, controlled and visible to senior leaders.

Similarly, the United States should invest systematically in evaluating how other nations are attempting to cope with cyber insecurity. An astute observer has argued that the United States can profit from the Australian model of limiting government purchases to “white-listed” software that has been tested and found acceptable.¹⁴² The Finns have established procedures for quarantining infected consumers, allowing them access only to sites that support remediation. The British have developed collaborative government-industry exchanges with some sectors that have no such relationships in the United States. Each nation operates in its own cultural, economic and political context, so lessons are not easily assimilated. But foreign experiences can be richly suggestive and the United States is now only harvesting them serendipitously.

Heavier investment in sociological, political and economic research will not appeal to many technologists and will not appear on technological agendas. Work of this kind will rarely meet academic “scientific” standards and sometimes (particularly in studying attackers) it will require unconventional research strategies, involving controversial ethical and legal decisions.¹⁴³ For these reasons, it will not be well-supported elsewhere. It should, however, be a policymaker priority and requires government support. The ninth recommendation suggests a vehicle for this support.

7. Fund a data collection consortium that will illuminate the character and magnitude of cyber attacks against the U.S. private sector, using the model of voluntary reporting of near-miss incidents in aviation. Use this enterprise as well to help develop common terminology and metrics about cybersecurity.

Channels have developed in recent years for sharing information between public and private actors about cyber attacks. Principal among these are the Enduring Security Framework, classified and unclassified Defense Industrial Base programs run by DOD and the National Cybersecurity and Communications Integration Center run by DHS.¹⁴⁴ Less-visible efforts involve some Information Sharing and Analysis Centers (ISACs)¹⁴⁵ and more informal exchanges such as those between the Nuclear Regulatory Commission and nuclear power companies, the Department of Energy and power companies, and the FBI and Secret Service with companies that are experiencing attacks.¹⁴⁶

Though valuable, these relationships tend to be ad hoc, information is not transferred from industry to industry, common metrics are not created as a foundation for analysis¹⁴⁷ and the separate efforts do not contribute to a public comprehension of what is happening. A commentator summarized his frustration after the recent cyber theft of credit cards from Target:

I have to wonder how many times this scene played out in 2013: an individual forensics firm analyzes a sophisticated retail breach involving point-of-sale malware – collecting mountains of interesting and useful data about the threats, threat actors and their methods – but at the end of the day has no mechanism by which to share that information with others in the retail and security community. ... [I am] incredulous at how the industry as a whole still sucks at sharing important information.¹⁴⁸

A model for how government-industry relationships could be improved is provided by a little-publicized successful effort in the aviation world. While regulatory requirements for aviation accident reporting are firmly established through the National Transportation Safety Board, there are no requirements for reporting the vastly more numerous and often no less informative near misses. Efforts to establish such requirements inevitably generate resistance: Airlines would not welcome more regulation and fear the reputational and perhaps legal consequences of data visibility; moreover, near accidents are intrinsically more ambiguous than accidents. An alternative path was forged in 2007 when MITRE, a government contractor, established an Aviation Safety Information Analysis and Sharing (ASIAS) system receiving near-miss data and providing anonymized safety, benchmarking and proposed improvement reports to a small number of initially participating airlines and the Federal Aviation Administration (FAA). A professional journal summarizes subsequent developments:

Today, membership has grown to 44 airlines representing 96 percent of commercial airspace operations and 131 safety data sources, according to the FAA. The MITRE Corp. analyzes and safeguards proprietary airline data; integrates it with MITRE's own aviation safety databases covering weather, radar tracks, airspace and traffic and other public data; conducts studies; and builds analysis capabilities. Airline data is shared over MITRE secure servers and includes pilot safety reports and FDR data.¹⁴⁹

Another journal records a US Airways participant's perspective about what is achieved:

For him, the key advantage of ASIAS has been the company's ability to tap ASIAS databases "to look at aggregated data or different airports or different types of data and to also compare and

use the Web portal dashboards” to analyze issues such as unstabilized approaches.

He explains, “I might be thinking we’re doing really well, but I can compare US Airways against the aggregate. ... I can see that maybe I have a problem at one airport, but am I the only one that has that problem?”¹⁵⁰

It is hard to predict whether a contractor-based sharing of cyber attack data would succeed, but it has attributes that make it worth trying. A third-party intermediary seems to mitigate distrust of government uses of the data; growth by accretion enables industry leaders to set examples for laggards; and the voluntary non-punitive nature of the effort seems to lower resistance. Because the entire structure is contractual, it requires no legislation.

8. Invest in research and development in conjunction with private industry and other nations to make cyber architectures more robust.

Federal agencies now substantially invest in research and development directly relating to cyber defense.¹⁵¹ They engage a vast web of subordinate bureaucratic offices, national laboratories, contractors, federally financed research and development organizations and universities. Alongside these actors, and sometimes supported by or in coordination with them, the private sector is investing considerable resources in cybersecurity through major software enterprises such as Microsoft and Google, hardware developers such as Intel, Internet service providers such as Verizon, anti-malware firms, and scores of small startup enterprises and academia. These firms do substantial work outside of the United States as well as within it.

The two key questions are whether this research should be refocused and/or more centrally directed. The recommended answers are respectively “yes” and “no.”

It is not surprising that it is hard to comprehend, much less evaluate, the universe of cyber research. Expenditures are difficult to compile and impossible to accurately allocate between operational, infrastructure, defensive and offensive investments;¹⁵² large parts of it are proprietary or classified; and it evolves, so even if a fuzzy snapshot were obtainable, it would have a short shelf life and capture different elements at different times.¹⁵³

Periodically, there are calls to better coordinate or centralize at least the federal parts of these expenditures. A recent congressional commission urged centralization,¹⁵⁴ and the most recently promulgated White House cyber strategy¹⁵⁵ states as its fourth “initiative”:

No single individual or organization is aware of all of the cyber-related R&D activities being funded by the Government. This initiative is developing strategies and structures for coordinating all cyber R&D sponsored or conducted by the U.S. government, both classified and unclassified, and to redirect that R&D where needed. This Initiative is critical to eliminate redundancies in federally funded cybersecurity research, and to identify research gaps, prioritize R&D efforts, and ensure the taxpayers are getting full value for their money as we shape our strategic investments.¹⁵⁶

The attractions of this rhetoric are understandable, but the congressional commission and the quoted White House initiative point in the wrong direction. To the extent cybersecurity is an engineering problem applying well understood principles to a well-defined goal, central coordination and efforts to “eliminate redundancies” are rewarding priorities. But to the extent that the field is not a mature science and competing concepts still require basic research, it is better to have less coordination and more competition. The latter is the more accurate diagnosis and prescription for cybersecurity today.

The question is how to drive the focus of competing research efforts toward the most fundamental problems of cybersecurity rather than simply to incremental, short-term improvement. We require strategies to develop architectures that will be intrinsically more secure than the present software and hardware that power the Internet and related information technologies.¹⁵⁷ Private-sector research and development is relevant to that effort but unlikely to optimally invest because large companies have heavy time discounts and vested interests in maintaining the existing systems. Small companies can be more radical, but radical change usually requires long-term investment beyond their resources. Moreover, researchers in most enterprises have limited or no classified access and therefore have only an incomplete view of cutting-edge developments in attack and defense.

The federal government can help to improve priorities by following these precepts:

- Map federal *and private* cybersecurity investments to the extent possible so as to help all participants better understand where investment is occurring and lacking;
- Do not attempt to strongly coordinate this investment, but instead recognize that the nascent nature of cybersecurity warrants trying and even retrying different paths to similar ends – redundancy is tolerable, even warranted;
- In lieu of strong coordination, identify critical opportunities for systemically enhancing security and incentivize effort in both the public and private sectors on these problems. This list will evolve.¹⁵⁸

The ninth goal of the Comprehensive National Cybersecurity Initiative (CNCI) shows a greater appreciation of these principles than the fourth initiative. It says:

One goal of the CNCI is to develop technologies that provide increases in cybersecurity by orders

of magnitude above current systems and which can be deployed within 5 to 10 years. This initiative seeks to develop strategies and programs to enhance the component of the government R&D portfolio that pursues high-risk/high-payoff solutions to critical cybersecurity problems. The Federal Government has begun to outline Grand Challenges for the research community to help solve these difficult problems that require ‘out of the box’ thinking. In dealing with the private sector, the government is identifying and communicating common needs that should drive mutual investment in key research areas.¹⁵⁹

This commendable statement of course leaves much unsaid, including: Who will identify these challenges? Who will administer the allocation of funds? How much funding? For which challenges? And how will research initiatives be taken from proof of concept through development, a transition commonly described as a requirement to voyage through “the valley of death”?

A stronger White House staff and a stronger federal cyber career workforce will increase the likelihood that these questions are productively answered. Our final recommendation addresses these issues.

9. Avoid delusions that centralized control through “a czar” or single agency are achievable or desirable. Instead, enhance federal cyber capacity by strengthening White House oversight, relocating the cyber coordinator outside the National Security Council and building more support capability for the White House. Make the U.S. government more cyber strong by creating a new federally funded research and development center (FFRDC) to recruit and retain cyber experts who are less likely to be attracted to work in federal civilian agency bureaucracies. That FFRDC can also be a focal point for communicating with and engaging private-sector experts.

In the face of political, structural and operational difficulties, the federal government has gotten better at cyber defense. Improvement has included the growth of expertise, focus and private sector outreach in the FBI, the intelligence community, the Department of Homeland Security, the Department of Energy, the Department of Defense and other agencies. This has been abetted by growing recognition of the cost and persistence of cyber attacks, the allocation of budget resources to strengthen cyber defenses and the creation of a cyber coordinator position in the White House.

It is not surprising that weaknesses remain. The problems are hard, they have (as described above) come upon the nation quickly, and bureaucracies resist rapid change. Executive orders, legislation and other promulgations articulate desiderata, but the Cabinet departments' personnel, processes and priorities have a great deal of inertia. Changing bureaucratic behavior and capabilities is not like remodeling your kitchen. It is more like building your muscles, losing 50 pounds or teaching yourself a foreign language. Recommendations in this realm accomplish more when they recognize rigidities, are fashioned modestly and proceed incrementally.

However, initiatives in this regard are crucial: To the extent cyber challenges evolve rapidly, unpredictably, and with only partial visibility, U.S. government resilience depends significantly on the skills and the orchestration of the federal workforce. Accordingly, this paper concludes by advancing modest but achievable recommendations for addressing two defensive priorities: building a better workforce¹⁶⁰ in civilian agencies concerned with cyber defense and strengthening White House abilities to coordinate the diverse and competitive federal agencies.

The federal workforce is strikingly unbalanced. One aspect of the imbalance has been hotly debated: The National Security Agency (NSA),

The federal bureaucratic organism is strongly muscular on one side (its military and intelligence functions) and malformed and malnourished on the other (its civilian side). Accordingly, it should be no surprise that it walks with a limp and often stumbles.

an intelligence organization, is intertwined with Cyber Command, a military organization. Because NSA has a longer history and stronger culture than Cyber Command, the intelligence mission has dominated the priorities, promotion opportunities and the budgets of these organizations. To counter this, a recent review group and a number of other assessments have recommended severing the two organizations.¹⁶¹ A presidential decision has resolved this issue, determining, at least for the moment, that the advantages of this combination outweigh the disadvantages.

A second aspect of this arrangement is more important. The muscularity of NSA and Cyber Command provides a large, talented and well-established military and civilian workforce focused on intelligence and military missions. The Information Assurance Directorate of NSA, charged with defense, has less power and prestige than the offensive side of NSA, but it is nonetheless well-funded and well-staffed and offers meaningful career opportunities (including movement between the defensive and offensive sides). The glaring deficiency is that there is no equivalent resource for the domestic agencies of the federal government.

To continue the earlier metaphor, the federal bureaucratic organism is strongly muscular on one side (its military and intelligence functions) and malformed and malnourished on the other (its civilian side). Accordingly, it should be no surprise that it walks with a limp and often stumbles.¹⁶²

Strong central leadership could correct at least some of this imbalance. However, the organism has only a rudimentary central nervous system. The only entity that can coordinate the whole is the White House. Efforts have been concentrated there within the National Security Council. But resources to meet this management challenge are badly incommensurate with what is needed. The cyber office consists of one regular political appointee¹⁶³ and a handful of professionals detailed for two years from across the national security agencies. The intelligence, foreign policy, domestic policy, budgetary, protective, research and development, representational and other activities demanded of White House coordinators can hardly be conducted by such a small and transitory group.

It is tempting to attempt to counter this at one stroke by creating a cyber czar with a dedicated workforce to address all cyber security problems (or at least all defensive cyber security problems). That path, however, overlooks several secondary difficulties and one primary problem. Secondly, the transaction costs of creating a new entity are immense: Everything is delayed while a newly proposed enterprise seeks authorization, budgets, administrative capability, etc. Moreover, when created, the new regime begins to become segmented, restricted and calcified just like the old ones: Existing bureaucracies and congressional constituencies assert old lines of authority and resist innovation, for example, and the hiring process becomes burdened with procedures and preferences just as at present.

Beyond this, the primary argument against consolidation is that the challenge is not one

of building a centralized cyber defense. It is to build a variety of defenses within the range of public and private environments that depend on cyber technology. Central guidance, research and development can help this process, but varied solutions have to be applied in varied ways by existing bureaucracies and the industries to which they relate.

Forgoing radical change (at least for now) does not, however, foreclose requirements and opportunities for improvement. A pervasive difficulty is that the federal work environments and hiring and promotion practices are not well-matched to recruitment, development and retention of cyber skilled talent – especially of what the Defense Science Board called “top-tier talent who can be certified to perform at the elite or extreme cyber conflict levels.”¹⁶⁴ In a highly competitive labor market,¹⁶⁵ the federal civil service impedes free flow in and out of the private sector and between positions within the federal personnel system; cybersecurity opportunities are commonly treated as a subset of the information technology field, though the two fields need to recruit and reward different talents;¹⁶⁶ hiring has strong presumptions in favor of some experiences (for example, veterans’ service) and credentials (for example, college graduation); promotion is routinely tied to time in grade; pay is not well-calibrated to capability; and culture and expectations about working hours, attire, benefits and more are quite different from those of the commercial cyber marketplace, not to mention environments in which hackers work.¹⁶⁷ There is more at stake than psychological benefits: Cybersecurity is a field that disproportionately depends on on-the-job training;¹⁶⁸ professionals learn a great deal from one another and benefit particularly from environments where they are congenially collocated.

NSA combats these difficulties with some special hiring authorities,¹⁶⁹ the glamour of its mission and the quality of some its special equipment

and by having created a critical mass of first-rank cyber experts. Civilian departments have no such strengths.¹⁷⁰ Throughout government¹⁷¹ there is a bifurcation of cyber expertise: older cyber-cognizant (but often not expert) leaders preside over young, tech-savvy people with few people in between; among the younger generation there is a lot of burn-out, little retention and much defection to civilian life.¹⁷² Attrition among the best and brightest detailed to the White House cyber office is illustrative: One-third of the “directors” leave government immediately after a two-year White House tour.¹⁷³

A federally funded research and development center could improve this situation by assembling a critical mass of cyber talent (on the order of several hundred skilled professionals), establishing flexible personnel policies (including flow in and out of the private sector)¹⁷⁴ and creating an environment where training, work patterns and promotion were more attractive than within the present federal system. (I note that I am and have been affiliated with some FFRDCs.¹⁷⁵ Readers may want to consider these affiliations when evaluating this recommendation.) Employees of this FFRDC could work for varying terms on assignment throughout the federal government and then return to the FFRDC or rotate to other agencies as needed and as helpful for their professional development. An FFRDC could also be a focal point for offering training to regular federal employees¹⁷⁶ and for identifying and engaging private-sector expertise to assist the federal government routinely and on an emergency basis. These qualities could be enhanced by establishing an office in Silicon Valley as well as Washington, by managing the FFRDC with a board drawing significantly from the private sector and by appointing a CEO with strong cyber talents and abilities to attract those with cyber skills.

FFRDCs now contribute valuable help on cyber issues.¹⁷⁷ The entity proposed here could develop from a present FFRDC or from other existing

enterprises. However, a new and more focused enterprise with a visionary, cyber-trained leader would be more likely to create a new culture and a critical mass of cyber skills. Labeled as a “Cyber Corps,” the enterprise could recruit better, establish a stronger sense of professional identity¹⁷⁸ and path for professional education,¹⁷⁹ and improve retention by creating more visible and sustainable career paths.¹⁸⁰

The proposed FFRDC could also enhance the power of the White House cyber office by dedicating a tiger team to supporting that enterprise. The intelligence, foreign policy, domestic policy, budgetary, protective, research and development, representational and other activities demanded of White House coordinators is now managed with hardly even the basics of administrative support or technically trained research assistants.

Beyond this, it is constraining to locate the cyber office within the National Security Council rather than independently in the White House. This arrangement impedes the White House cyber group’s outreach to society at large,¹⁸¹ biases it toward the national security agencies and positions its leader at a level insufficient to effectively lead Cabinet principles. Making the office an independent entity within the White House would strengthen it.

Administrative changes will not create dramatic improvements, but they will facilitate them. These changes should be initiated.

VII. CONCLUSION

Information technologies are representative of – and intertwined with – several new technologies in this new millennium. Biotechnology, robotics, big-data analysis and additive manufacturing are now evident other members of this group. These technologies share five characteristics: While they can do immense good, they also can be commandeered to do much bad; they are rapidly proliferating to state and non-state actors; weapons derived from this new knowledge, or that exploit new dependencies, can achieve large-scale effects with small investments; these weapons are relatively easily disguised in production and use (therefore undermining prevention, detection, attribution and deterrence); and these possibilities have arisen more quickly (and can be expected to continue to evolve more quickly) than national security systems have adapted.

Facing this technological flood, it feels as though the United States is working ever faster to plug holes in its dikes. Since the odds are against this response, it is tempting to try to drain the ocean. But the problems posed by these technologies cannot be addressed by adopting Luddite positions. Pondering a mid-20th century antecedent of this class – nuclear weapons – the physicist Freeman Dyson observed that once we develop them, along with their blessings we are cursed by having them forever.¹⁸² Nor can the United States deal with these challenges simply by associating them with a particular adversary. The weapons they provide can be used by anyone for any purpose.

The question is not how the United States keeps others from accessing these technologies, nor how it walls them off, nor how it lives without them. It is how the United States copes with them, assuming that others master them as well.

The answers are not simple or wholly satisfying. The United States will make its peace with the new

technologies by understanding them and finding ways to limit their potentially pernicious and especially their potentially disastrous effects. This paper has tried to enhance that understanding with regard to cyber technology. It points to some straightforward priorities that can make the United States more secure. Beyond that, it suggests a mode of thought about how to live with insecurity.

ENDNOTES

1. George Tenet, "Information Security Risks, Opportunities, and the Bottom Line" (Sam Nunn Nations Bank Policy Forum, Atlanta, April 6, 1998), https://www.cia.gov/news-information/speeches-testimony/1998/dci_speech_040698.html.
2. The terms "cyberspace" and "cybersecurity" are used in different ways by different authors. This text adopts the approach advanced by the National Research Council: "This report defines cyberspace broadly as the artifacts based on or dependent on computer and communications technology; the information that these artifacts use, store, handle, or process; and how these various elements are connected. Security in cyberspace (i.e., cybersecurity) is about technologies, processes, and policies that help to prevent and/or reduce the negative impact of events in cyberspace that can happen as the result of deliberate actions against information technology by a hostile or malevolent actor." National Research Council, *At the Nexus of Cybersecurity and Public Policy: Some Basic Concepts and Issues* (Washington: The National Academies Press, prepublication copy, 2014), 1, <http://nebula.wsimg.com/584842e1269950244b00b6fdeb299527?AccessKeyId=E2B9E747A9A6F0184889&disposition=0&alloworigin=1>.
3. See for example, John Arquilla and David Ronfeldt, eds., *In Athena's Camp: Preparing for Conflict in the Information Age* (Santa Monica, CA: RAND, 1997). More recently the Defense Science Board warned: "There is no exaggerating our dependence on DoD's information networks for command and control of our forces, the intelligence and logistics on which they depend, and the weapons technologies we develop and field. In the 21st century, modern armed forces simply cannot conduct high-tempo, effective operations without resilient, reliable information and communication networks and assured access to cyberspace." Department of Defense, *Quadrennial Defense Review Report* (February 2010), 37, <http://www.defense.gov/qdr/qdr%20as%20of%2029jan10%201600.pdf>.
4. This issue was recently highlighted by the indictment of Chinese military officers on charges of economic espionage. *United States of America v. Wang Dong, Sun Kailiang, Wen Xinyu, Huang Zhenyu, Gu Chunhui*, No. 14-118 (U.S. District Court, Western District of Pennsylvania, May 1, 2014), <http://www.justice.gov/iso/opa/resources/5122014519132358461949.pdf>. More generally, it has been observed that "[w]hile the threat to intellectual property is often less visible than the threat to critical infrastructure, it may be the most pervasive cyber threat today. . . . As military strength ultimately depends on economic vitality, sustained intellectual property losses erode both U.S. military effectiveness and national competitiveness in the global economy." Department of Defense, *Department of Defense Strategy for Operating in Cyberspace* (July 2011), 4, <http://www.defense.gov/news/d20110714cyber.pdf>. The cyber exfiltration of intellectual property is discussed at length in Commission on the Defense of Intellectual Property (2013), http://www.ipcommission.org/report/IP_Commission_Report_052213.pdf which Chapter 2 of this report (pp. 23ff) discusses the difficulties measuring the value of American intellectual property losses and provides estimates in the range of 300 billion dollars annually.
5. "[W]e must understand that the attacks we should be most concerned with are not designed to disable their digital targets, but to manipulate them in an unintended way to achieve a desired physical outcome. Many professionals have limited their thinking to dealing with the loss of individual elements or capabilities of their control systems and have failed to fully embrace the implications of calculated misuse." Michael Assante, President and CEO of National Board of Information Security Examiners of the United States Inc., testimony to the Committee on Homeland Security and Governmental Affairs, U.S. Senate, November 17, 2010, 3, <http://www.hsgac.senate.gov/download/2010-11-17-assante-testimony>.
6. Robert J. Butler, "Securing the Grid: Opportunities and Risks in Operational Technology" (Center for a New American Security, March 2014), <http://www.cnas.org/securing-the-grid#.U5YKMJSwl1c>.
7. Michael B. Kelley, "NSA: Snowden Stole 1.7 MILLION Classified Documents And Still Has Access To Most Of Them," *businessinsider.com*, December 13, 2013, <http://www.businessinsider.com/how-many-docs-did-snowden-take-2013-12#1K4fuM>.
8. Robert Parry, "Looking Backward on Bradley Manning," *globalresearch.ca*, August 23, 2013, <http://www.globalresearch.ca/looking-backward-on-bradley-manning/5346540>.
9. The growth in magnitude of stored data is suggested: "In 1986, you could fill the world's total storage using the world's total bandwidth in two days. Today, it would take 112 days of the world's total bandwidth to fill the world's total storage, and the measured curve between 1986 and today is all but perfectly exponential." Daniel E. Geer Jr., "People in the Loop: Are They a Failsafe or a Liability?" (Suits and Spooks Anti-Conference, Arlington, Virginia, February 8, 2012), <http://geer.tinho.net/geer.suitsandspooks.8ii12.txt>. The compression of stored data also facilitates espionage. Commenting on an engineer found to have 250,000 stolen documents in his house, the Office of the National Counterintelligence Executive observed that these would have required four, four-drawer filing cabinets but could be captured on one CD: "On average, one page of typed text holds 2 kilobytes (KB) of data; thus, 250,000 pages x 2 KB/page = 500,000 KB, or 488 megabytes (MB). A data CD with a capacity of 700 MB retails for \$0.75, and a flash drive with a capacity of 4 gigabytes costs about \$13.00." Office of the National Counterintelligence Executive, *Foreign Spies Stealing US Economic Secrets in Cyberspace: Report to Congress on Foreign Economic Collection and Industrial Espionage, 2009-2011* (October 2011), 2, http://www.ncix.gov/publications/reports/fecie_all/Foreign_Economic_Collection_2011.pdf.
10. On average, breaches in 2011 stole data on 1.1 million identities; in 2012 that number was 604,826. The median breach stole data on 8,350 identities. Symantec, "Internet Security Threat Report 2013," vol. 18 (Symantec, April 2013), 33, https://www.symantec.com/content/en/us/enterprise/other_resources/b-istr_appendices_v18_2012_221284438-en-us.pdf. Regarding Target, "hackers stole about 40 million debit and credit card numbers from cards swiped at its stores between Nov. 27 and Dec. 15. The thieves also took personal information — including email addresses, phone numbers, names and home addresses — for another 70 million people." The Associated Press, "Target data breach info begins popping up, but source still hard to track," *wjla.com*, January 22, 2014, <http://www.wjla.com/articles/2014/01/target-data-breach-info-begins-popping-up-but-source-still-hard-to-track-99508.html#ixzz2rF44jm36>.

11. When University of California, San Diego, Professor Stefan Savage and his colleagues replicated bot spamming networks, they observed: "75% of bot-originated spam was being immediately dropped on the floor, most of the remainder was filtered by spam filters, and only a very small fraction of users actually clicked on the links contained in such messages and an even smaller fraction ever decided to place an order. Yet in spite of this it was clear that the raw volume of this activity could produce significant revenue." Rik Farrow, "Interview with Stefan Savage: On the Spam Payment Trail," *login*, 36 no. 4 (August 2011), 12-13, <http://cseweb.ucsd.edu/~savage/papers/LoginInterview11.pdf>.

12. The story is told with remarkable detail by James Cortada. Cortada uses the term ICT – information communication technology – as well as IT. James W. Cortada, *The Digital Flood: The Diffusion of Information Technology Across the U.S., Europe and Asia* (New York: Oxford University Press, 2012).

13. Tenet, "Information Security Risks, Opportunities, and the Bottom Line."

14. Dave Aucsmith, Senior Director of Microsoft's Institute for Advanced Technology in Governments, "Rethinking Cyber Defense . . . Lessons Learned" (Microsoft Federal Forum, Washington, March 4, 2014), as reported by Dan Verton, "What Microsoft knows about cybersecurity might surprise you," *fedcoop.com*, March 6, 2014, <http://fedcoop.com/what-microsoft-knows-about-cybersecurity-and-why-it-might-surprise-you/>. See agenda summary at https://presentations.inxpo.com/Shows/microsoft/GMO/Global_MSC/2014/2014_3_4_Virtual_Federal_Forum/registration/FedForumAgenda2014VirtualVersion.pdf.

15. National Research Council, *At the Nexus of Cybersecurity and Public Policy: Some Basic Concepts and Issues*, 13: "The digital representation of information has a number of important security consequences . . . The fact that a given sequence of bits could just as easily be a program as data means that a computer that receives information assuming it to be data could in fact be receiving a program, and that program could be hostile. Mechanisms in a computer are supposed to keep data and program separate, but these mechanisms are not foolproof and can sometimes be tricked into allowing the computer to interpret data as instructions. It is for this reason that downloading data files to a computer can sometimes be harmful to the computer's operation – embedded in those data files can be programs that can penetrate the computer's security, and opening the files may enable such programs to run."

16. Symantec calculates that in 2012 one in every 414 emails was a phishing attack. Symantec, "Internet Security Threat Report 2013," 75.

17. On misdirection, see for example, "DNSChanger Malware," http://www.fbi.gov/news/stories/2011/november/malware_110911/DNS-changer-malware.pdf. Symantec reports that one in every 532 websites it scanned was infected with malware. Symantec, "Internet Security Threat Report 2013," 29, http://www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v18_2012_21291018.en-us.pdf. In the same report, Symantec writes:

"How does a hacker add his code to a legitimate website? Toolkits are available that make it easy. For example, in May 2012, the LizaMoon toolkit used a SQL injection technique to affect at least a million websites. Other approaches include:

- Exploiting a known vulnerability in the website hosting or content management software
- Using phishing, spyware, or social engineering to get the webmaster's password
- Hacking through the Web server backend infrastructure, such as control panels or databases
- Paying to host an advertisement that contains the infection

This last technique, known as malvertising, means that legitimate websites can be impacted without even being compromised. This form of attack appears to be very common." Symantec, "Internet Security Threat Report 2013," 27. "The Elderwood Project" describes a widespread campaign of targeting particular websites for infection. Gavin O'Gorman and Geoff McDonald, "The Elderwood Project" (Symantec, 2012), http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/the-elderwood-project.pdf.

18. In a famous 1984 speech accepting an award as a software writer, Ken Thompson asserted and convincingly illustrated his conclusion that no software writer could ever fully trust another's code. "You can't trust code that you did not totally create yourself. (Especially code from companies that employ people like me.) No amount of source-level verification or scrutiny will protect you from using untrusted code. . . . A well installed microcode bug will be almost impossible to detect." Ken Thompson, "Reflections on Trusting Trust," *Communications of the ACM*, 27 No. 8 (August 1984), 761-763, <http://cm.bell-labs.com/who/ken/trust.html>.

19. For a timeline of the evolution of the number of transistors in an assortment of microprocessors, graphics processing units, and similar technologies, see "Willus.com's Brief History of the Desktop PC back to 1970," Willus.com, 2014, <http://willus.com/archive/timeline.shtml>.

20. Jonathan Corbet, Greg Koah-Hartman, and Amanda McPherson, "Linux Kernel Development: How Fast it is Going, Who is Doing It, What They are Doing, and Who is Sponsoring It," (The Linux Foundation, March 2012), <http://go.linuxfoundation.org/who-writes-linux-2012>.

21. Steve McConnell observes that for a system with a million lines of code, the requirements specification "would typically be about 4,000-5,000 pages long, and the design documentation can easily be two or three times as extensive as the requirements. It's unlikely that an individual would be able to understand the complete design for a project of this size – or even read it." Steve McConnell, *Code Complete, Second Edition* (Redmond, WA: Microsoft Press, 2004), 12, [http://testa.robterra.free.fr/My%20Books/Computer%20programming/Java/Code%20Complete%20%20Second%20Edition%20By%20Steve%20Mcconnell%20\(Microsoft%20Press%202004\).pdf](http://testa.robterra.free.fr/My%20Books/Computer%20programming/Java/Code%20Complete%20%20Second%20Edition%20By%20Steve%20Mcconnell%20(Microsoft%20Press%202004).pdf).

22. Steve Lohr and John Markoff, "Windows Is So Slow, But Why?," *The New York Times*, March 27, 2006, http://www.nytimes.com/2006/03/27/technology/27soft.html?adxnnl=1&pagewanted=all&adxnnlx=1382805118-0jnNRGXEPip3xoW+BDp8Q&_r=0. The number of lines is not static as code is frequently amended, in part to counter vulnerabilities. One commentator observes:

"In 2003, something on the order of five thousand new security vulnerabilities were reported to CERT, and that number per year has only grown since then. . . . The number of vulnerabilities introduced by all this additional code added to MS Windows systems [in 2003] by [a] very conservative estimate . . . is one per 1000. This means that MS Windows was adding new bugs at a rate of about 4,550 per year. That means that MS Windows alone gained almost as many vulnerabilities as were actually *discovered*, for all software reported to CERT, in the year 2003. Given that MS Windows is actually a fairly small part of CERT's total database of bugs, the implications are dismaying."

Chad Perrin, "The danger of complexity: More code, more bugs," techrepublic.com, February 1, 2010, <http://www.techrepublic.com/blog/it-security/the-danger-of-complexity-more-code-more-bugs/#>. A more encouraging situation is described by Andy Ozment and Stuart E. Schechter in "Milk or Wine: Does Software Security Improve with Age?" Their view is that some security enhancement is gained over time as vulnerabilities are discovered and patched. Software is accordingly said to be "more like wine than milk – it improves with age." Andy Ozment and Stuart E. Schechter, "Milk or Wine: Does Software Security Improve with Age?" (paper presented at the 15th USENIX Security Symposium, Vancouver, British Columbia, Canada, July 31–August 4, 2006), https://www.usenix.org/legacy/events/sec06/tech/full_papers/ozment/ozment_html/index.html.

23. "Code wants to be wrong. We will never have 100% error-free software." Omer Ben-Shalom et al., "Rethinking Information Security to Improve Business Agility" (Intel, January 2011), <http://www.intel.com/content/www/us/en/enterprise-security/intel-it-enterprise-security-rethinking-information-security-to-improve-business-agility-paper.html>.

"[I]nstead of a nice, clean, organized communication hierarchy with everything neatly calling only its 'immediately surrounding' levels, almost everything can communicate with almost everything else on all sorts of disjoint levels." Hoglund and McGraw, *Exploiting Software: How to Break Code*, 43. The authors also observe: "[R]isk cannot be measured out of context. A classic example of an environmental effect is demonstrated by taking a program that has been successfully run with no security problem for years on a proprietary network and putting it on the internet. The risks change immediately and radically." Hoglund and McGraw, *Exploiting Software: How to Break Code*, 42. McConnell, *Code Complete, Second Edition*, 518–520, summarizes sources of error.

24. For example, endemic memory vulnerabilities in the most prevalent languages, C+ and C++, were countered by moving to C#, Java and other languages. Mozilla writes that its language Rust has been in development for several years and is rapidly approaching stability. . . . [It] is intended to fill many of the same niches that C++ has over the past decades . . . it is *safe by default*, preventing entire classes of memory management errors that lead to crashes and security vulnerabilities." Mozilla, "Mozilla and Samsung Collaborate on Next Generation Web Browser Engine," The Mozilla Blog on Mozilla.org, April 3, 2013, <https://blog.mozilla.org/blog/2013/04/03/mozilla-and-samsung-collaborate-on-next-generation-web-browser-engine/>. See, generally, <http://www.rust-lang.org/> and, for other efforts, <http://safecode.cs.illinois.edu/> and <http://cyclone.thelanguage.org/wiki/Why%20Cyclone/>.

25. It is a subject of theoretical debate whether one could be created.

26. Eric Rosenbach and Robert Belk, "U.S. Cybersecurity: The Current Threat and Future Challenges," in *Securing Cyberspace: A New Domain for National Security*, eds. Nicholas Burns and Jonathon Price (Queenstown, MD: Aspen Institute, 2012), 50, estimated three to six errors per thousand lines of code. Steve McConnell concluded that the "industry average" is about one to 25 errors per thousand lines of code, with Microsoft's Applications Division said to have brought this down to .5 errors for released product, according to a 1992 report. McConnell, *Code Complete, Second Edition*, 21. The situation is dynamic, and neither the 1992 assertion nor McConnell's 2004 estimate may obtain a decade later. For purposes of estimation, one error per thousand lines of code as a rule of thumb still seems useful to obtain a general understanding of the problem.

27. McConnell, *Code Complete, Second Edition*, 653.

28. Not only are the estimates crude, but also variables like especially extensive checking can reduce the rate. "Harlan Mills pioneered 'cleanroom development,' a technique that has been able to achieve rates as low as 3 defects per 1000 lines of code during in-house testing and 0.1 defect per 1000 lines of code in released product (Cobb and Mills 1990). A few projects – for example, the space-shuttle software – have achieved a level of 0 defects in 500,000 lines of code using a system of formal development methods, peer reviews, and statistical testing (Fishman 1996)." McConnell, *Code Complete, Second Edition*, 521.

29. Greg Hoglund and Gary McGraw, *Exploiting Software: How to Break Code* (Addison-Wesley, 2004).

30. Joshua Bloch, "Extra, Extra – Read All About It: Nearly All Binary Searches and Mergesorts are Broken," <http://googleresearch.blogspot.com/2006/06/extra-extra-read-all-about-it-nearly.html>.

31. Verizon, "2013 Data Breach Investigations Report" (Verizon, 2013), http://www.verizonenterprise.com/resources/reports/rp_data-breach-investigations-report-2013_en_xg.pdf. The precise numbers are that 66 percent took months or more to discover and 69 percent were discovered by outsiders.

32. Unfortunately, this visibility is not symmetrical. While attackers are covert, to counter attacks on their operating systems, defenders must broadcast their knowledge to immense numbers of users. Oracle reports that Java runs on 3 billion computers. Brian Krebs, "Critical Java Update Plugs 51 Security Holes," *Krebsonsecurity.com*, October 16, 2013, <http://krebsonsecurity.com/2013/10/java-update-plugs-51-security-holes/>.

33. Dan Geer, "Vulnerable Compliance," *login*, 35 no. 6 (December 2010), 28, <https://www.usenix.org/system/files/login/articles/geer.pdf>.

34. "IBM X-Force 2012 Trend and Risk Report" (IBM, March 2013), 50 and 54, <http://public.dhe.ibm.com/common/ssi/ecm/en/wgl03027usen/WGL03027USEN.PDF>.

35. The terms "vulnerabilities" and "exploits" are widely used in the profession. A useful distinction between "inherent vulnerabilities" and "operational vulnerabilities" was made in a January 2013 Defense Science Board report. Defense Science Board, *Resilient Military Systems and the*

Advanced Cyber Threat (January 2013), <http://www.acq.osd.mil/dsb/reports/ResilientMilitarySystems.CyberThreat.pdf>. I have relabeled “inherent vulnerabilities” as “structural vulnerabilities,” because software and hardware can be modified to reduce weaknesses. In addition, kinetic attacks can destroy cyber equipment, but that concern is outside the subject of this paper.

36. As described in Appendix I, generalizations from all existing databases must be made with caution, but Verizon’s “2013 Data Breach Investigations Report,” one of the most precise and useful sources, suggestively found insider misconduct in one in every seven of the major 2012 data breaches it identified. See particularly Pages 5, 20, 23 and 29.

Insider risks are extensively discussed in George J. Silowash, et al., “Common Sense Guide to Mitigating Insider Threats, 4th Edition,” Paper 677 (Software Engineering Institute, 2012), <http://repository.cmu.edu/sei/677>. The Carnegie Mellon Software Engineering Institute has a program actively studying this risk. For a view from the standpoint of the attacker, see Kevin Mitnick and William Simon, *The Art of Deception* (Kineticstomp, Indianapolis, (2002), 82 http://fr.thehackademy.net/madchat/esprit/textes/The_Art_of_Deception.pdf: “Every computer system in the world has at least one human that uses it. So if the attacker is able to manipulate people who use the systems, the obscurity of the system is irrelevant.”

37. Song Yun et al., “Towards Hardware Trojan: Problem Analysis and Trojan Simulation” (paper presented at the second International Conference on Information Engineering and Computer Science, Wuhan, China, December 25-26, 2010), summarizes the hardware risk: “[T]he problem of hardware Trojan has . . . become an increasing concern. . . . A hardware Trojan . . . is a malicious modification of the circuitry of an integrated circuit. Moreover, as more and more digital systems are using 3rd party Intellectual Property (IP) blocks, the hardware Trojan problem has been more complex since the IP blocks are untrusted.” For broad education on this issue, see Mohammad Tehranipoor and Cliff Wang, eds., *Introduction to Hardware Security and Trust* (Bücher: SpringerLink, 2011).

38. The Stuxnet virus that crippled the Iranian nuclear program illustrates an aspect of this set of vulnerabilities. Public accounts assert that the Iranian facilities’ computers were “air-gapped” (that is, isolated) and the malware did not enter the system via the Internet but instead by contamination through a vendor or other user knowingly or unknowingly connecting with, for example, his own contaminated computer.

39. Mandiant reports: “In 2012, companies spent \$134 billion on outsourcing business processes such as finance, accounting, HR, and procurement. Combine that with the estimated \$252 billion organizations spent on IT outsourcing in 2012 and it adds up to a lot of organizations allowing outside vendors unfettered access to large portions of their networks. . . . During our investigations in 2012, we found an increase in the number of outsourced and managed service providers who were compromised and used as a primary access point for attackers to gain entry to their victims’ networks. . . . In many instances, the attackers initially gained access to the service provider solely as a means to find a way into their real target — the client of the service provider. In those cases, we have seen the attackers compromise the first victim — the outsourced service provider — gather the intelligence they need to facilitate their compromise of the second victim, and then lay dormant at the first victim for months or even years, only accessing backdoors at those companies if they need to regain access to the second victim.”

Mandiant, “M-Trends 2013: Attack the Security Gap,” (Mandiant, 2013), 4, http://www.greycastlesecurity.com/resources/documents/2013_M_Trends.pdf. And see the case study of “an energy company” on pages 14-16.

40. Dan Kaufman, “An analytical framework for cyber security,” unclassified briefing (DARPA Public Release Center, November 2011), <http://www.darpa.mil/WorkArea/DownloadAsset.aspx?id=2147484449>.

41. A presumption of breach has been advocated by others. See Assante, testimony to the Committee on Homeland Security and Governmental Affairs, where he uses the term “contested territory.” In 2012, Jeffrey Carr wrote:

“I was recently invited to participate in a closed Congressionally-mandated meeting of a dozen or more intelligence and technology experts to discuss what the research and development priorities of the U.S. Intelligence Community should be for the next 10 years. While a lot of ideas were tossed about and shot down, one of a handful that rose to the surface was the need to re-think our security paradigm from the long-standing one of trying to keep bad guys out of our networks to assuming that they’re already inside. This is known in government circles as “Assumption of Breach”. Debora Plunkett of the NSA’s Information Assurance Directorate has said as much back in December 2010. Price Waterhouse Coopers has been an advocate of that strategy as well.”

Jeffrey Carr, “Assumption of Breach: The New Security Paradigm,” Digital Dao blog on jeffreycarr.blogspot.com, July 12, 2012, <http://jeffreycarr.blogspot.com/2012/07/assumption-of-breach-new-security.html>. Department of Defense, *Department of Defense Strategy for Operating in Cyberspace*, 6-7, states a desire for “a complete spectrum of cyberspace scenarios [in] exercises and training” and speaks about “a presumption of breach” but as a general operation principle says only that “intrusions may not always be stopped at the network boundary.” It does not establish the recommended presumption.

42. “The Internet protocol was designed virtually assuming that all users were trustworthy.” President’s Council of Advisors on Science and Technology, *Report to the President: Immediate Opportunities for Strengthening the Nation’s Cybersecurity* (November 2013), 5, http://www.whitehouse.gov/sites/default/files/microsites/ostp/PCAST/pcast_cybersecurity_nov-2013.pdf. See also, Michael Warner, “Cybersecurity: A Pre-history,” *Intelligence and National Security*, 27, no. 5 (October 2012), 781-799, <http://www.tandfonline.com/doi/abs/10.1080/02684527.2012.708530>. As Warner observes, some prescient early developers raised these issues, but they were not significantly addressed in the face of other priorities.

43. An attacker can outsource a number of these tasks by buying ready-made tool kits. See, for example, Chris Grier et al., “Manufacturing Compromise: The Emergence of Exploit-as-a-Service” (paper presented at the 19th ACM Conference on Computer and Communications Security, Raleigh, North Carolina, Oct. 16-18, 2012), describing the market for exploits; and Juan Caballero et al., “Measuring Pay-per-Install: The Commoditization of Malware Distribution” (paper presented at the 20th USENIX Security Symposium, San Francisco, California, August 8-12, 2011), <http://www.icir.org/vern/papers/ppi-usesec11.pdf>, describing sale of access (“pay per install”) to infected computers.

44. Social engineering can also exploit other targets, for example, by inducing help desks to provide passwords. Mitnick and Simon's *The Art of Deception* provides a consummate social engineer's education in the craft.

45. "For the context used in the software security industry and in [Microsoft], a vulnerability is a security exposure that results from a product weakness that the product developer did not intend to introduce and should fix once it is discovered." Microsoft Security Response Center, "Software Vulnerability Exploitation Trends" (Microsoft Corporation, 2013), <http://www.microsoft.com/en-us/download/confirmation.aspx?id=39680>.

Vulnerabilities can be of many different kinds. Many stem from loopholes in programming languages, some from logical errors in the construction of system architectures and others from unanticipated effects arising from the interaction of systems.

46. Most robustly, sophisticated corporations invest in distinguishing vital from merely valuable confidential information, create enclaves for especially prized information, keep the number of authorized users in low double digits and demand at least two-factor authentication (for example, passwords and a physical token) for authorized access. However, this process is costly and never-ending, most information needs to be shared more widely within the company and with subcontractors, and even tightly controlled enclaves have operational vulnerabilities.

47. For example, Ars Technica gave three individuals a list of 16,000 cryptographically hashed passcodes, and they had deciphered the majority of them in the span of one day. Dan Goodin, "Anatomy of a hack: How crackers ransack passwords like 'qeadzcrwsfxv1331,'" *Ars Technica*, May 27, 2013, <http://arstechnica.com/security/2013/05/how-crackers-make-minced-meat-out-of-your-passwords/>.

48. Defense Science Board, *Resilient Military Systems and the Advanced Cyber Threat*, 51, summarized its observations:

"While well-intentioned and strongly supported, [Department of Defense] initiatives have not had the desired impact on the overall [information assurance] posture of the Department. Defensive measures implemented at the boundaries between the NIPRNet and the Internet proved to be only marginally effective in blocking successful intrusions or reducing the overall attack surface of DoD networks and systems. Mobile platforms (smart phones, tablets, etc.) exacerbate this already challenging problem. Red teams, conducting operations during military exercises or at the request of Military Department and Agency officials, continue to have a nearly perfect success rate breaking into the systems."

49. "Hacker Intelligence Initiative, Monthly Trend Report #14" (Imperva, December 2012), http://www.imperva.com/docs/HII_Assessing_the_Effectiveness_of_Antivirus_Solutions.pdf, provides a short overview. Imperva (at Page 3) traces the history to 1988 and references Gartner's calculation that in 2011, consumers and businesses spent \$7.4 billion on anti-viral products.

50. See "Hacker Intelligence Initiative, Monthly Trend Report #14," 8-9. Microsoft routinely issues corrections on the second Tuesday of every month, "Patch Tuesday." Wolfgang Kandek, "January 2014 Patch Tuesday Preview - Update," Qualys Blog on community.qualys.com, January 9, 2014, <https://community.qualys.com/blogs/laws-of-vulnerabilities/2014/01/09/january-2014-patch-tuesday-preview>.

Oracle issues patches every quarter. Qualys comments on the latter: "[T]hese quarterly releases typically address over 100 vulnerabilities in their large software line. For example, 127 were addressed in October of 2013. Analyzing the applicability of these flaws to one's software infrastructure and addressing them are a major concern for any organization that uses Oracle products." Qualys, "Eschelbeck's New 'Laws of Vulnerabilities' Research Demonstrates Growing Threat to Internal Networks," [qualys.com](http://www.qualys.com), July 28, 2004, <http://www.qualys.com/company/newsroom/news-releases/usa/2004-07-28/>.

51. Software companies abandon old operating systems as they introduce new ones, but the old systems continue to be used. "Apple computers running 10.5 or less get no updates (comprising about half the installed base). Any Microsoft computer running XP gets no updates (comprising about half the installed base)." Daniel E. Geer Jr., "On Abandonment" (IEEE Computer and Reliability Societies, July-August 2013), 88, <http://geer.tinho.net/ieee/ieee.sp.geer.1307.pdf>.

52. Millions will be unprotected because they operate counterfeit software (and therefore typically do not receive patches), do not subscribe to virus protection, subscribe but do not download or delay downloading. Some of those who do download will use interacting programs and applications that may not appropriately integrate the patches received for a primary program. Anti-virus modifications will be timely for many but too late to protect other users. On Page 32, the "IBM X-Force 2012 Trend and Risk Report" provides a nicely illustrative timeline of a Java exploit discovered and disclosed to Oracle, which issued a patch on February 14, 2012. The discoverer published his work in late February.

"A month later, after the details of the vulnerability were released, a working exploit was integrated into the Blackhole exploit kit, and within a few days, into the Phoenix exploit kit. Then, in early May, an exploit for this same vulnerability was seen in the RedKit41 exploit kit. Considering that a patch was available from Oracle on February 14th, it indicates that attackers believe that organizational and individual patch uptake is infrequent enough to be successful with exploits for recently patched vulnerabilities."

Also see IBM's similar discussion of other exploits just after the quoted text. Qualys estimated that two months after a patch was released to counter the Conficker worm, 30 percent of computers running the Windows operating system remained unpatched. "Conficker Working Group: Lessons Learned" (The Rendon Group, June 17, 2010), 4, http://www.confickerworkinggroup.org/wiki/uploads/Conficker_Working_Group_Lessons_Learned_17_June_2010_final.pdf.

Recently Microsoft has moved to providing patches even for counterfeit systems. Many vendors do not.

53. In 2004, Gerhard Eschelbeck, CTO of Qualys, presented a detailed analysis of patch repairs of vulnerabilities. Among other things, he reported that "the half-life identifies the length of time it takes users to patch half of their systems, reducing their window of exposure. The half-life of critical vulnerabilities for external systems is 21 days and for internal systems [i.e., corporate intranets] is 62 days. This number doubles with lowering degrees of severity." Qualys, "Eschelbeck's New 'Laws of Vulnerabilities' Research Demonstrates Growing Threat to Internal Networks." Commenting on the

recent theft of credit card information from Target, an expert observed yet greater delays in that context:

"Basically the retailer would receive notification of a patch, and go off to an outside server and pull it down. Very similar to you getting notified of a patch being available on a Windows machine. Except in these situations, when you're dealing with critical retail software and installations, patching is very formal process, because uptime and reliability is very important. So they ... would install the patch on their internal lab servers, run and test it to verify that everything would run smoothly. Then they would then certify the patch ... and ... distribute it to stores 1-5 tonight, and then go 6-10 then next day, or whatever their mechanism was to distribute them. But it would typically take them 2-3 months to finish the distribution to all the stores."

Comment by Tom Arnold in Brian Krebs, "These Guys Battled BlackPOS at a Retailer," *Krebsonsecurity.com*, February 4, 2014, <http://krebsonsecurity.com/2014/02/these-guys-battled-blackpos-at-a-retailer/#more-24517>. The discussion below returns to this issue in Recommendation 6.

54. "Anti-virus software is one of the most complicated applications. It has to deal with hundreds of file types and formats," including executables, documents, media files, etc. Using a public database, Feng Xue reported 165 recognized anti-virus vulnerabilities over a four-year period. (See particularly examples he presents at Pages 9ff.) He concluded: "[I]t is clear that anti-virus software can be targeted just like[sic] other components or services of computer systems." Feng Xue, "Attacking Antivirus" (Nevis Networks Inc., 2008), <http://www.blackhat.com/presentations/bh-europe-08/Feng-Xue/Whitepaper/bh-eu-08-xue-WP.pdf>. More recently, one group "fuzzed" (randomly attacked) samples of anti-virus software and reported: "Almost all of the tested tools seemed to be easy to crash using our relatively simple automated techniques. Some of the observed failures had information security implications, and should be considered as vulnerabilities. This is alarming considering the tested products were advertised as security products." The group further noted that: "Anti-virus tools run at high privileges, increasing the impact of potential compromise. Anti-virus tools are commonly installed organisation-wide on all able computers, including (or especially) on computers in critical and high-profile roles. Usage of Anti-virus tools is commonly mandated by organisational policy, contract and other administrative and/or legal requirements. US HIPAA legislation is commonly interpreted to mandate use of anti-virus software." University of Oulu, "PROTOS Genome Test Suite c10-archive," Abstract, https://www.ee.oulu.fi/research/ouspg/PROTOS_Test-Suite_c10-archive#Results_from_the_Test_Runs. In his classic, *The Mythical Man-Month*, Frederick Brooks offers a more general observation: "All repairs [to soft-ware] tend to destroy the structure, to increase the entropy and disorder of the system." Frederick Brooks, *The Mythical Man-Month* (Addison-Wesley, 1995), 122.

55. The Defense Science Board recognized this dynamic in another context. When it recommended monitoring of networks it was compelled also to acknowledge that:

"It is not unusual for a sophisticated adversary, who has infiltrated a network, to monitor in real time as the network owners try to kick them out. Frequently, the adversary then implements a counter to the network owner's defensive actions and can be back in the network in a matter of minutes or hours." Defense Science Board, *Resilient Military Systems and the Advanced Cyber Threat*, 61.

An illuminating case study for the U.S.-China Economic and Security Review Commission reports that before a major exfiltration, the target company detected activity attributed to these types of sophisticated attackers, but it was generally characterized by low volumes of traffic through compromised hosts and appeared primarily focused on maintaining access and presence. The company's information security staff detected and countered these compromises in the months before the final data exfiltration; however, the attackers appear to have simply created new entry vectors or reverted back to other pre-established means of accessing the company's network. Bryan Krekel et al., "Capability of the People's Republic of China to Conduct Cyber Warfare and Computer Network Exploitation" (Northrop Grumman Corporation Information Systems Sector, October 9, 2009), 61, <http://www2.gwu.edu/~nsarchiv/NSAEBB/NSAEBB424/docs/Cyber-030.pdf>.

56. See, for example, Chet Hosmer, "Polymorphic & Metamorphic Malware" (Wetstone, 2008), http://www.blackhat.com/presentations/bh-usa-08/Hosmer/BH_US_08_Hosmer_Polymorphic_Malware.pdf.

National Research Council, *At the Nexus of Cybersecurity and Public Policy: Some Basic Concepts and Issues*, 37: "The Einstein program of the Department of Homeland Security (DHS) is an example of a signature-based approach to improving cybersecurity. ... Einstein monitors Internet traffic going in and out of government networks and inspects a variety of traffic data (i.e., the header information in each packet but not the content of a packet itself) and compares that data to known patterns of such data that have previously been associated with malware. If the match is sufficiently close, further action can be taken (e.g., a notification of detection made or traffic dropped). This signature-based technique for detection has two primary weaknesses. First, it is easy to morph the code without affecting what the program can do so that there are an unlimited number of functionally equivalent versions with different signatures. Second, the technique cannot identify a program as malware if the program has never been seen before."

57. Dan Kaufman, "An analytical framework for cyber security."

58. Symantec concludes that "Attackers use as many zero-day vulnerabilities as they need, not as many as they have." Symantec, "Internet Security Threat Report 2013," 5. Symantec may be particularly influenced in this conclusion by its recognition of a group, apparently Chinese, that used seven zero-day vulnerabilities in a series of attacks on selected targets, supplementing prior efforts even before the earlier vulnerabilities were widely patched. O'Gorman and McDonald, "The Elderwood Project." And see Assante, testimony to the Committee on Homeland Security and Governmental Affairs, 4.

In another highly publicized series of intrusions, *The New York Times* reported on attacks it experienced and had assessed by Mandiant. "Over the course of three months, attackers installed 45 pieces of custom malware. The Times — which uses antivirus products made by Symantec — found only one instance in which Symantec identified an attacker's software as malicious and quarantined it, according to Mandiant." Nicole Perlroth, "Hackers in China Attacked The Times for Last 4 Months," *The New York Times*, January 30, 2013, <http://www.nytimes.com/2013/01/31/technology/chinese-hackers-infiltrate-new-york-times-computers.html>.

59. Symantec observed 14 new zero-day vulnerabilities in 2012. As a second-order effect, when some of these vulnerabilities are announced and patched, criminal groups use what they learn from the announcement to attack

systems that have not protected themselves. Symantec, "Internet Security Threat Report 2013," 5.

60. A memorandum from Bill Gates in 2002 raised security priorities at Microsoft. He said:

"Every week there are reports of newly discovered security problems in all kinds of software, from individual applications and services to Windows, Linux, Unix and other platforms. We have done a great job of having teams work around the clock to deliver security fixes for any problems that arise. Our responsiveness has been unmatched – but as an industry leader we can and must do better. Our new design approaches need to dramatically reduce the number of such issues that come up in the software that Microsoft, its partners and its customers create. . . . So now, when we face a choice between adding features and resolving security issues, we need to choose security. Our products should emphasize security right out of the box . . . "

Bill Gates, "Memorandum" (2002), as reproduced in "Bill Gates: Trustworthy Computing," *Wired Magazine* (January 17, 2002), <http://archive.wired.com/techbiz/media/news/2002/01/49826>.

Microsoft's proactive security efforts led eventually to the use of supercomputers to "fuzz" (randomly attack) its software to identify bugs. The most sophisticated version of this software is described by Patrice Godefroid.

"Since 2008, SAGE has been running 24/7 on approximately 100-plus machines/cores automatically fuzzing hundreds of applications in Microsoft security testing labs. This is more than 300 machine-years . . ." Patrice Godefroid et al., "SAGE: Whitebox Fuzzing for Security Testing," *Communications of the ACM*, 55 no. 3 (March 2012), 44, http://research.microsoft.com/en-us/um/people/pg/public_pfiles/cacm2012.pdf. See also the description of the work of the IBM X-Force research and development team in "IBM X-Force 2012 Trend and Risk Report," 3.

61. A long list of open-market buyers can be found at Bugcrowd, "The Bug Bounty List" [bugcrowd.com](https://bugcrowd.com/list-of-bug-bounty-programs), 2013, <https://bugcrowd.com/list-of-bug-bounty-programs>; and Andy Greenberg, "Shopping for Zero-Days: A Price List For Hackers' Secret Software Exploits," *Forbes.com*, March 23, 2012, <http://www.forbes.com/sites/andygreenberg/2012/03/23/shopping-for-zero-days-an-price-list-for-hackers-secret-software-exploits/> provides a convenient summary of prices. Microsoft's program is advertised at Microsoft, "Security TechCenter: Microsoft Bounty Programs," Microsoft.com, June 26, 2013, <http://technet.microsoft.com/en-us/security/dn425036>: "Microsoft will pay up to \$100,000 USD for truly novel exploitation techniques against protections built into the latest version of our operating system. Learning about new exploitation techniques earlier helps Microsoft improve security by leaps, instead of capturing one vulnerability at a time as a traditional bug bounty alone would." Google writes: "To honor all the cutting-edge external contributions that help us keep our users safe, we maintain a Vulnerability Reward Program for Google-owned web properties, running continuously since November 2010" and provides a table of prospective payments at Google, "Vulnerability Reward Program," [google.com](http://www.google.com/about/appsecurity/reward-program/), <http://www.google.com/about/appsecurity/reward-program/>. On government buying, see Joseph Menn, "Special Report: U.S. cyberwar strategy stokes fear of blowback," *Reuters.com*, May 10, 2013, <http://www.reuters.com/article/2013/05/10/us-usa-cyberweapons-specialreport-idUSBRE9490EL20130510>:

"The Department of Defense and U.S. intelligence agencies, especially the NSA, are spending so heavily for information on holes in commercial computer systems, and on exploits taking advantage of them, that they are turning the world of security research on its head, according to longtime researchers and former top government officials.

Many talented hackers who once alerted companies such as Microsoft Corp to security flaws in their products are now selling the information and the exploits to the highest bidder, sometimes through brokers who never meet the final buyers. Defense contractors and agencies spend at least tens of millions of dollars a year just on exploits . . . "

62. Interview with Mike Walker, program manager, DARPA Information Innovation Office, July 23, 2013.

63. Some of this work has roots in insights from academia, working either independently or with commercial entities. DARPA and other U.S. government offices have provided support and/or worked in conjunction with commercial and academic efforts.

64. These classes of tools have a variety of components. ASLR, for example, may include bottom-up randomization, top-down randomization, heap randomization, etc.

65. Hovav Shacham, Matthew Page, et al., "On the Effectiveness of Address-Space Randomization," (paper presented at CSS'04, Washington, District of Columbia, October 25-29, 2004), <http://benpfaff.org/papers/asrandom.pdf>.

66. Crispin Cowan, Perry Wagle, Carlton Pu, et al., "Buffer overflows: attacks and defenses for the vulnerability of the decade," *Proceedings of the Foundations of Intrusion Tolerant Systems (OASIS'03)* (2003).

67. "Buffer overflow attacks, in which an attacker forces a program or component to store malicious code in an area of memory not intended for it, are some of the most common exploits seen today. DEP is a Windows feature that enables the system to mark one or more pages of memory as non-executable. Marking memory regions as non-executable means that code cannot be run from that region of memory, which makes it harder for exploits involving buffer overruns to succeed. DEP was introduced in Windows XP SP2 and has been included in all subsequent releases of Windows desktop and server operating systems." Microsoft Security Intelligence Report, "Protecting Your Software," Microsoft.com, http://www.microsoft.com/security/sir/strategy/default.aspx#!section_3_3. See also Microsoft Support, "A detailed description of the Data Execution Prevention (DEP) feature in Windows XP Service Pack 2, Windows XP Tablet PC Edition 2005, and Windows Server 2003," [microsoft.com](http://support.microsoft.com/kb/875352), <http://support.microsoft.com/kb/875352>.

68. Adobe substantially reduced its vulnerability by moving to sandboxing in its recent versions. Adobe Systems Inc., "Protected Mode," [adobe.com](http://www.adobe.com/devnet-docs/acrobatetk/tools/AppSec/protectedmode.html), February 14, 2014, <http://www.adobe.com/devnet-docs/acrobatetk/tools/AppSec/protectedmode.html>.

69. For example, "[I]n 2012, we saw the overall percentage of vulnerabilities disclosed by [large enterprise software vendors] decrease to 26% [of all disclosed vulnerabilities]. This is the first decrease in this category over the past five years and is something we will watch closely in 2013 to see if this is a one-time occurrence or whether it represents a downward trend as enterprise

software vendors continue to implement secure development practices within their software development lifecycle, allowing vendors to identify and remedy vulnerabilities before the code makes its way into a new software release." "IBM X-Force 2012 Trend and Risk Report," 57. Microsoft Security Response Center, "Software Vulnerability Exploitation Trends," 16, reported that in the last quarter of 2012 its Malicious Software Removal Tool, "which executes on more than 600 million Windows systems around the world each month," found an order of magnitude more infections on systems running its Windows XP systems than on those running Windows 8.

70. "The combination of ASLR and DEP creates a fairly formidable barrier for attackers to overcome in order to achieve reliable code execution when exploiting vulnerabilities." Microsoft Security Intelligence Report, "Protecting Your Software."

71. Email from Mike Frantzen, Kudu Dynamics, August 21, 2013.

72. The broad assessment of the participants interviewed for this paper is that increased and proactive efforts by vendors and governments are yielding rewards. Moreover, mainstream systems are maturing as bugs are weeded out and later versions benefit from more robust architectures. As a result, vulnerabilities in established operating systems are becoming harder to identify and to exploit.

73. A Microsoft research report is illustrative: "The increasing prevalence of DEP and ASLR has forced attackers to identify techniques that can be used to exploit vulnerabilities even when these features are enabled. These techniques have led to an increasing number of exploits that attempt to bypass ASLR by relying on images that have not opted into ASLR or by leveraging a vulnerability to disclose information about the layout of an application's address space. Similarly, the use of return-oriented programming (ROP) has become common in exploits that seek to bypass DEP. The increasing use of ROP was a focal problem for the winning solutions that were submitted to the Microsoft BlueHat Prize competition in 2012—some of which were integrated into the Microsoft Enhanced Mitigation Experience Toolkit (EMET) 3.5 technical preview." Microsoft Security Response Center, "Software Vulnerability Exploitation Trends," 11.

74. Bruce Schneier, "The Internet of Things Is Wildly Insecure — And Often Unpatchable," *wired.com*, January 6, 2014, <http://www.wired.com/Opinion/2014/01/Theres-No-Good-Way-To-Patch-The-Internet-Of-Things-And-Thats-A-Huge-Problem/>.

75. See Department of Defense, *Department of Defense Strategy for Operating in Cyberspace*, 7: "Active cyber defense is DoD's synchronized, real-time capability to discover, detect, analyze, and mitigate threats and vulnerabilities. ... It operates at network speed by using sensors, software, and intelligence to detect and stop malicious activity before it can affect DoD networks and systems. As intrusions may not always be stopped at the network boundary, DoD will continue to operate and improve upon its advanced sensors to detect, discover, map, and mitigate malicious activity on DoD networks."

76. The National Institute of Standards and Technology summarizes this set of activities as "intrusion prevention systems," saying they "are primarily focused on identifying possible incidents, logging information about them, attempting to stop them, and reporting them to security administrators." Karen Scarfone

and Peter Mell, "Guide to Intrusion Detection and Prevention Systems (IDPS): Recommendations of the National Institute of Standards and Technology," National Institute of Standards and Technology Special Publication 800-94 (February 2007), <http://csrc.nist.gov/publications/nistpubs/800-94/SP800-94.pdf>. The federal government's present priorities are described in J. Michael Daniel, "Cybersecurity: Executive branch departments and agencies will achieve 95% implementation of the Administration's priority cybersecurity capabilities by the end of FY 2014. These capabilities include strong authentication, Trusted Internet Connections (TIC), and Continuous Monitoring," *performance.gov*, <http://goals.performance.gov/content/cybersecurity>. The key goal is summarized: "Transform the historically static security control assessment and authorization process into an integral part of a dynamic enterprise-wide risk management process. This change allows departments and agencies to maintain an ongoing near-real-time awareness and assessment of information security risk and rapidly respond to support organizational risk management decisions."

77. The co-chairs of a recent Defense Science Board report summarized the board's conclusion: "[T]he United States cannot be confident that our critical Information Technology (IT) systems will work under attack from a sophisticated and well-resourced opponent utilizing cyber capabilities in combination with all of their military and intelligence capabilities (a 'full spectrum' adversary)." Defense Science Board, *Resilient Military Systems and the Advanced Cyber Threat*, 1.

78. National Research Council, *At the Nexus of Cybersecurity and Public Policy: Some Basic Concepts and Issues*, 2.

79. Particularly in space (where Sputnik was the most memorable surprise) and in Soviet operation of a massive biological weapons program only belatedly discovered

80. This number was calculated for this manuscript and provided to the author by the Intel Corporation Research Group and reported in email correspondence with the author, May 28 and June 3, 2014.

81. Miniwatts Marketing Group, "History and Growth of the Internet from 1995 till Today," *Internetworldstats.com*, March 24, 2014, <http://www.internetworldstats.com/emarketing.htm>.

82. Former Air Force Chief of Staff Larry Welch arrived at a similar standard: "The first priority is to identify those segments of cyberspace where freedom of action is essential to mission accomplishment." Larry D. Welch, "Cyberspace — The Fifth Operational Domain," *IDA Research Notes* (Summer 2011), <https://www.ida.org/~media/Corporate/Images/Publications/2011researchnotessummer.ashx>.

83. The standard also has implications for civilian systems on which military capability depends. This point is discussed in this paper's second recommendation

84. For example, if the strategies suggested in Recommendation 1 are pursued in federal systems, they will establish a precedent and demand for IT configurations of this kind, and that will incentivize similar changes in commercial markets for critical systems. The effort will also illuminate the costs and benefits of moving to this way of doing business.

85. Others have reached the same conclusion though with other terminology or in other contexts. For example, National Research Council, *At the Nexus of Cybersecurity and Public Policy: Some Basic Concepts and Issues*, 36: "The most basic way to improve cybersecurity is to reduce the use of information technology (IT) in critical contexts," though the only example offered is of disconnecting from the Internet. In the private sector, large companies commonly establish software systems empowering employees to do some things and not others. (Their challenge is to truly minimize capabilities without granting too many waivers.) Joe Markowitz has put the point pithily in the first of his "three laws of information assurance," saying: "Capability equals vulnerability; each new capability introduces a new vulnerability; AND 'excess' capability means unnecessary vulnerability." Joe Markowitz, "Cyber Insecurity: Blame the Victim," undated PowerPoint briefing.

86. There is, to offer a caricatured illustration, no reason why we should equip those responsible for nuclear missiles with the ability to play games on their computers. Less obviously, many federal users should not have communicative capabilities, copying capabilities, built-in cameras, built-in abilities to use disks and thumb drives and many other attributes irrelevant to their jobs. However, computer standardization and the savings from buying at scale have left many of these capabilities in place (in some instances straightforwardly, in others requiring a screwdriver and modest skills to restore blocked capabilities).

87. Defense Science Board, *Resilient Military Systems and the Advanced Cyber Threat*, 56, recommends innovators "provide out-of-band command and control for most sensitive systems."

88. Dan Geer, "Heartbleed as Metaphor," lawfareblog.com, April 21, 2014, <http://www.lawfareblog.com/2014/04/heartbleed-as-metaphor/>. The principal Department of Defense cyberstrategy document recognizes this, reciting: "Multiple networks can add diversity, resiliency, and mission assurance to cyberspace operations." Department of Defense, *Department of Defense Strategy for Operating in Cyberspace*, 6. But the flow of acquisition decisions is in the opposite direction.

89. As concisely put by P.W. Singer and Allan Friedman: "The idea is to build systems where the parallel for offense and defense isn't from warfare, but biology. When it comes to bacteria and viruses in our bodies, human cells are actually outnumbered by as much as 10 to 1. But the body has built up an amazing capacity of both resistance and resilience, fighting off what is most dangerous and, as Vint Cerf, the computer scientist who is literally one of the 'fathers of the Internet,' puts it, figuring out how to 'fight through the intrusion.' No computer network will mimic the human body perfectly, but DARPA and other groups are working on 'intelligent' computer security networks that learn and adapt to resist cyberattacks." P.W. Singer and Allan Friedman, "Cult of the Cyber Offensive," ForeignPolicy.com, January 15, 2014, http://www.foreignpolicy.com/articles/2014/01/15/cult_of_the_cyber_offensive_first_strike_advantage.

90. Common standards and standardized equipment generate efficiencies. The resulting cost savings are reinforced by some gains in security effectiveness: System administrators can more readily regulate, monitor, assess and patch standardized systems. Training is easier and more effective.

91. Diversification makes attackers' tasks more difficult and can mitigate their achievement if successful. The first factor is more relevant to security risks from low-end attackers and the second from high-end attackers.

92. The resistance to regulation is substantial and understandable. Asked what role they think the federal government should play "protecting private sector networks against nation state or other attacks," about a third of corporate senior IT professionals favor some regulatory and related action, a third want "advice coordination and research support only," and a third want monitoring and intervention "only when clear need" or they want the government to "leave private sector alone." Dan Geer and Mukul Pareek, "The Index of Cyber Security: Annual Report," cybersecurityindex.org, August 2013, http://geer.tinho.net/ICS_Annual_Report_2013.pdf. But after much struggle, Section 113 of the Dodd-Frank Act authorized "enhanced prudential standards" for companies whose "material financial distress . . . could pose a threat to U.S. financial stability." Office of the Federal Register, Financial Stability Oversight Council, *Authority to Require Supervision and Regulation of Certain Nonbank Financial Companies*, Billing Code 4810-25-P (2013), 1, <http://www.treasury.gov/initiatives/fsoc/documents/nonbank%20designations%20-%20final%20rule%20and%20guidance.pdf>. The crisis of 2008 taught America that if the nation is to remain financially sound, some firms are too important to fail. Does the United States need to await a comparable crisis before recognizing that the same thing is true with respect to national security?

93. White House, Executive Order 13636, Section 9 (February 12, 2013), <http://www.whitehouse.gov/the-press-office/2013/02/12/executive-order-improving-critical-infrastructure-cybersecurity>.

94. The Commission on the Theft of American Intellectual Property, "The IP Commission Report," 3, estimates that China is "between 50% and 80% of the problem," saying that "[t]he major studies range in their estimates of China's share of international IP theft; many are roughly 70%, but in specific industries we see a broader range." See also Pages 15ff.

95. See, for example, Hillary Rodham Clinton, "Speech on the Future of US Foreign Policy" (Washington, November 30, 2012), http://www.foreignpolicy.com/articles/2012/11/30/hillary_clintons_remarks_at_foreign_policy_transformational_trends_forum. See also the recent indictment of Chinese military officers without expectation that they will stand trial. See *United States of America v. Wang Dong, Sun Kailiang, Wen Xinyu, Huang Zhenyu, Gu Chunhui*.

96. Exceptions include The Commission on the Theft of American Intellectual Property, "The IP Commission Report," and biennial reports to Congress required by the Intelligence Authorization Act for Fiscal Year 1995, Section 809(b), of which the most recent example is the Office of the National Counterintelligence Executive (2011).

97. Relevant present powers include federal prosecutorial powers under the Economic Espionage Act and tariff authorities to bar products based on illegally obtained trade secrets. A dozen possibly deterrent responses are summarized by the U.S.-China Economic and Security Review Commission, *2013 Report to Congress* (November 2013), 253-257, http://origin.www.uscc.gov/sites/default/files/annual_reports/Complete%202013%20Annual%20Report.PDF.

On the Economic Espionage Act, see "Criminal Law – Economic Espionage – Ninth Circuit Upholds First Trial Conviction Under § 1831 of the Economic Espionage Act of 1996," *Harvard Law Review*, 125 no. 8 (June 20, 2012), 2177: "Despite a startling escalation in the rate and scale of economic espionage, surprisingly few cases have been prosecuted under the Act." Among relevant tariff authorities, see 19 U.S.C. § 1337, "Unfair Practices in Import Trade," <http://www.gpo.gov/fdsys/granule/USCODE-2011-title19/USCODE-2011-title19-chap4-subtitleI-partII-sec1337/content-detail.html>.

98. This is a limitation of the National Institute of Standards and Technology, *Framework for Improving Critical Infrastructure Cybersecurity, Version 1.0* (February 12, 2014), <http://www.nist.gov/cyberframework/upload/cybersecurity-framework-021214.pdf>. While asserting that it "is not a one-size-fits-all approach," the framework correctly describes itself as written for all "organizations – regardless of size, degree of cyber risk or cybersecurity sophistication." The variance between industries (discussed under Recommendation 6) demands more variance in our approaches. Hopefully, Version 2.0 of the Framework will offer more detail and point the way to more varied engagement by federal agencies concerned with diverse industries.

99. See the discussion in Appendix I.

100. There are costs to public exposure. It may induce excessive fear, inform opponents about America's understanding and methods and distract the nation's energies in internal debate. All these things occurred in the United States' confrontation with the Soviet Union and in confronting terrorism, but the United States managed to limit these costs while providing significant disclosure. The need for balance underscores the importance of making these decisions outside the intelligence community, which naturally is predisposed to classification.

101. Similarly, while there are a handful of commendable academic efforts and no doubt there is some classified work, we have limited understanding of the ecosystem of cyberattackers to include the relationship between criminals. See Ross Anderson et al., "Measuring the Cost of Cybercrime" (paper presented at the Workshop on the Economics of Information Security, Berlin, June 25–26, 2012), http://weis2012.econinfosec.org/papers/Anderson_WEIS2012.pdf; Caballero et al., "Measuring Pay-per-Install: The Commoditization of Malware Distribution"; Geer, "People in the Loop: Are They a Failsafe or a Liability?"; and Dan Guido, "The Exploit Intelligence Project" (ISEC Partners, July 25, 2011), http://www.trailofbits.com/resources/exploit_intelligence_project_2_slides.pdf.

102. This risk is not dealt with by the executive order's requirement that an annual review be done of existing dependencies. We need to make the effort to forecast nascent dependencies and reduce them before they mature and become less tractable. There are also issues about dependencies on systems abroad whose failure would have national security consequences for the United States. That topic is beyond the scope of the executive order and of this paper but will require attention in the years ahead.

103. "Security researcher and Type 1 diabetic Jerome Radcliffe demonstrated a potentially lethal vulnerability in wireless medical devices by hacking his own insulin pump and glucose meter – showing a live audience that a malicious hacker could potentially use security gaps in the devices to end someone's life." Joshua Philipp, "Poor Cybersecurity in Health Devices a Life-threatening Problem," *The Epoch Times*, August 15, 2011, <http://www.theepochtimes.com/>

[n2/technology/poor-cybersecurity-in-health-devices-a-life-threatening-problem-60411.html](http://www.theepochtimes.com/n2/technology/poor-cybersecurity-in-health-devices-a-life-threatening-problem-60411.html).

104. The Defense Science Board and others have reached the same conclusion. See Defense Science Board, *Resilient Military Systems and the Advanced Cyber Threat*, 6: "Since it will be impossible to fully defend our systems against [high-level state] threats, deterrence must be an element of an overall risk reduction strategy." Recommendation 1 will help reduce attack efficacy and certainty – that is, it will contribute a measure of "deterrence by denial" – but this will not alone be fully protective.

105. Stephen J. Cimbala, "Nuclear Crisis Management and 'Cyberwar' Phishing for Trouble?," *Strategic Studies Quarterly*, 5 no. 1 (Spring 2011), 123–125, usefully emphasizes that control is generally difficult to sustain in crisis and may be particularly difficult to assert over attacks and responses at cyberspeeds. In addition, the challenge may be less one of establishing responses than of conveying the prospect of those responses without that prospect being destabilizing.

106. The use of other means is sometimes called "cross-domain deterrence." On nuclear deterrence see particularly Defense Science Board, *Resilient Military Systems and the Advanced Cyber Threat*, 40ff. More modest deterrent initiatives are discussed in Franklin Kramer and Melanie Teplinsky, "Cybersecurity and Tailored Deterrence," Issue Brief (Atlantic Council, December 2013), http://www.atlanticcouncil.org/images/publications/Cybersecurity_and_Tailored_Deterrence.pdf. The *Strategic Studies Quarterly* devoted its Spring 2011 issue to cyberdeterrence. *Strategic Studies Quarterly*, 5 no. 1 (Spring 2011), <http://www.au.af.mil/au/ssq/2011/spring/spring11.pdf>.

107. See generally, Emily Goldman and John Arquilla, eds., "Cyber Analogies," NPS-DA-14-001 (Naval Postgraduate School, February 28, 2014), <http://calhoun.nps.edu/public/handle/10945/40037>.

108. A good summary and point of view are presented in Joseph Nye, "The Regime Complex for Managing Global Cyber Activities" (Global Commission on Internet Governance, May 2014), especially Pages 9ff, http://issuu.com/cigi/docs/gcig_paper_no1/1?e=4209517/7932585. Recently, the undersecretary of state for arms control and international security has commissioned a study on strategies for building international norms for cyber stability, confidence-building, etc. U.S. Department of State, "Terms of Reference: ISAB Study of a Framework for International Cyber Stability," state.gov, July 17, 2013, <http://www.state.gov/t/avc/isab/213008.htm>.

109. David C. Gompert and Phillip C. Saunders, *The Paradox of Power: Sino-American Strategic Restraint in an Age of Vulnerability* (Washington: National Defense University Press, 2011), point in a similar direction. See particularly Pages 123ff arguing that China has common interests with the United States in establishing restraints on cyber attacks, that these interests will grow over time and that the United States should make a concerted effort to discuss these norms with China.

110. See, for example, Michael Dobbs' account of the pivotal day in the Cuban missile crisis, October 27, 1962. Dobbs quotes Arthur Schlesinger Jr.'s judgment that it was "the most dangerous day in human history." Dobbs summarizes:

"The day began with Fidel Castro dictating a telegram urging Khrushchev to use his nuclear weapons against their common enemy; it ended with

the Kennedy brothers secretly offering to give up U.S. missiles in Turkey in exchange for a Soviet climb down in Cuba. In between these two events, Soviet nuclear warheads were transported closer to Cuban missile sites, a U-2 spy plane was shot down over eastern Cuba, another U-2 strayed over the Soviet Union, a Soviet nuclear-armed submarine was forced to the surface by U.S. Navy depth charges, the Cubans began firing on low-flying U.S. reconnaissance aircraft, the Joint Chiefs of Staff finalized plans for an all-out invasion of Cuba, and the Soviets brought tactical nuclear weapons to within fifteen miles of the U.S. naval base at Guantanamo Bay. Any one of these incidents could have led to a nuclear exchange between the two superpowers." Michael Dobbs, *One Minute to Midnight* (New York: Alfred A. Knopf, 2008), xiii-xiv.

111. China may recognize this. In June 2013 it subscribed to a U.N. report's conclusion that international law applies to cyberspace. See the discussion in U.S.-China Economic and Security Review Commission, *2013 Report to Congress*.

112. Iran provides a case in point. Jeff Bardin, "Iranian Cyber Capabilities: Treadstone 71" (presentation at GovSec West conference, Dallas, Texas, October 2012), <http://www.youtube.com/watch?v=RhK7mrgID3Q>, presents a plausible and relatively detailed assessment of Iranian efforts. The presentation describes how protest movement uses of social media around the 2009 election sensitized the authorities to the need to control telecommunications. This led to Iranian Revolutionary National Guard IT investments in an offensive program. Bardin's view is that through the Islamic Revolutionary Guard Corps and other entities Iran is now among the top five nations in offensive cyber capabilities, spending some \$100 million per year.

113. Bruce Schneier, "The Story Behind The Stuxnet Virus," *Forbes* (Oct. 7, 2010), <http://www.forbes.com/2010/10/06/iran-nuclear-computer-technology-security-stuxnet-worm.html>.

114. As well explored by Elbridge Colby, "Defining Strategic Stability: Reconciling Stability and Deterrence," in *Strategic Stability: Contending Interpretations*, eds. Elbridge A. Colby and Michael S. Gerson, (Carlisle, PA: Strategic Studies Institute and U.S. Army War College Press, February 2013), 47, <http://www.strategicstudiesinstitute.army.mil/pdffiles/PUB1144.pdf>, there has long been an ambivalence in U.S. policy between the aim of avoiding nuclear use and the nation's using the threat of nuclear weapons as a deterrent against conventional attacks. However, the United States is not just threatening destructive cyberattacks; it is apparently engaged, both offensively and defensively, in their use.

The United States is also considering delegation of cyber tools to field-level units. See Brian Weeden, "Cyber offence and defence as mutually exclusive national policy priorities," *Disarmament Forum*, four (2011), 19, <http://www.unidir.org/files/publications/pdfs/confronting-cyberconflict-en-317.pdf>.

115. Eric Sterner observes: "Rather than symmetrical bipolar relationships, cyberspace is governed by a potentially infinite number of asymmetrical, multilateral, and bilateral relationships that are constantly in flux. ... [A]mbiguity will be prevalent before, during, and after engagements.

Perhaps the greatest problem encountered when applying strategic deterrence models to cyberspace is the difficulty of identifying the challenger and appropriate retaliatory targets. This was not a problem in traditional models of deterrence, whether nuclear or conventional. Theoretically, an

attacker's identity would always be known; only nation-states possessed the capability of launching significant military attacks. Actors in cyberspace, however, are 'created' in cyberspace. They may or may not correspond to the creator's identity in the real world." Eric Sterner, "Retaliatory Deterrence in Cyberspace," *Strategic Studies Quarterly*, 5 no. 1 (Spring 2011), 66, <http://www.au.af.mil/au/ssq/2011/spring/spring11.pdf>.

116. Cyber working groups, established recently in bilateral discussions with Russia and China, provide natural forums for such discussions. It seems more likely to be fruitful to proceed separately with each country rather than in a three-party discussion.

117. An important but secondary issue is not addressed here: Does an agreed red line against destructive attacks draw a line at exploitation of any vulnerability, at installation of software that can be programmed for destructive behavior or at the initiation of destructive behavior (not just the capability for destructive behavior)? The first and second positions are technically so close that they may not represent a meaningful distinction. It may be that only the third position is viable. This paper assumes the third position, but as noted in the last paragraph of this recommendation, the issue requires much more extended discussion.

118. This is consistent with, but an intensification of, present policy. That policy has been rather haltingly stated as "DoD and the Department of State are actively engaged with Allies, partners, and other states to build transparency and confidence with traditional adversaries." Department of Defense, *Cyberspace Policy Report to Congress* (November, 2011), 5, http://www.defense.gov/home/features/2011/0411_cyberstrategy/docs/NDA%20Section%20934%20Report_For%20webpage.pdf. Some small measures have been initiated with Russia to establish communications channels designed to build mutual confidence on cyber matters. See White House, "Fact Sheet: U.S.-Russian Cooperation on Information and Communications Technology Security," [whitehouse.gov](http://www.whitehouse.gov/the-press-office/2013/06/17/fact-sheet-us-russian-cooperation-information-and-communications-technol), June 17, 2013, <http://www.whitehouse.gov/the-press-office/2013/06/17/fact-sheet-us-russian-cooperation-information-and-communications-technol>. Development of confidence-building efforts might be encouraged through study by groups such as the JASONs and the National Academy of Sciences.

119. This discussion focuses on peacetime norms. David Gompert and Phillip Saunders argue for an agreement "that at no time should networks critical to civilian and economic well-being be subject to attacks, except in retaliation." They also discuss considerations for and against cyber constraints in the course of military conflict, observing that "[T]he danger of escalation from military to general cyber war provides one of the most powerful incentives for mutual restraint." Gompert and Saunders, *The Paradox of Power: Sino-American Strategic Restraint in an Age of Vulnerability*, 143, 145.

120. For example, one important intermediate issue will be the use of cyber attacks to destructively manipulate data and operations inside the victim's network.

121. There is active discussion in this area. See, for example, Patrick Lin, Fritz Allhoff and Neil C. Rowe, "Computing Ethics: War 2.0: Cyberweapons and Ethics," *Communications of the ACM*, 55 no. 3 (March 2012), 24, <http://www3.nd.edu/~cpence/ewt/Lin2012.pdf>; and John Mallery, "High-impact Functional Norms For Cyber Threat Reduction" (presentation at the panel on "Cyber Conflict: Models, Deterrence, Norms and Threat Reduction," at the

seventh International Forum of the Partnership of State Authorities, Civil Society and the Business Community in Ensuring Information Security and Combating Terrorism, Garmisch-Partenkirchen, Munich, Germany, April 24, 2013). Also see Michael N. Schmitt, ed., *Tallinn Manual on the International Law Applicable to Cyber Warfare* (Cambridge University Press, 2013).

122. These questions may be amplified by other factors, including improvements in anti-satellite weaponry and ballistic missile defenses.

123. This negotiation will be complicated by ambiguity and debate about whether intrusion for intelligence gathering can be distinguished from intrusion for potential sabotage. An agreement may be realizable only by forswearing both.

124. Some would diminish these difficulties, but raise others, by advocating the more sweeping approach of defining any state use (or state-sanctioned use) of cyber to attack as a "use of force" as prohibited in U.N. Charter Article 2(4) ("All members shall refrain in their international relations from the threat or use of force against the territorial integrity or political independence of any state, or in any other manner inconsistent with the purposes of the United Nations.")

125. The difficulties of controlling cyber weapons stem not only from their proliferation but also from their unpredictability. The complexity and interactivity described in Section III make software effects less predictable than the explosive effects of nuclear and conventional weapons. Cyber deterrence, response and confidence building all need to give unintended consequences greater weight in their planning than in equivalent thinking about other weapons. This both strengthens the motivation for confidence building and makes it more difficult to achieve.

126. Jim Miller comments that "If this were only a 3-way agreement, our allies and partners would ask why we have left them out of this protective agreement, and would likely be concerned more broadly about our commitment to their security. On the other hand, if we extend this understanding to all countries, we could be tying our hands regarding the use of cyber against other actors. One possibility is that we could expand any understanding but state that we have an exclusion — that this applies only to countries that are in accord with their international obligations under the UN charter including the NPT (much as we do for our negative security assurance regarding nuclear use). This would put Iran and North Korea in particular outside of the agreement." Correspondence with the author, May 23, 2014.

127. Some other examples would include ill-considered but prevalent factors in how IT security budgets are determined, resistance and optimal response to IT security centralization, the psychological and economic causes that prompt excessive retention of insecure legacy systems and police agency priorities that often privilege arrests and undervalue achievements that "merely" thwart disruption and theft.

128. Companies engaged in transportation, oil and gas exploration, high-volume online retail sales, etc., will fall on a spectrum between these examples, administering their systems in ways that have some overlap and some differences from these examples.

129. These paragraphs benefited from the insight of Stefan Savage of the University of California, San Diego, and Brandon Wales, director of the DHS

Homeland Infrastructure Threat and Risk Analysis Center. The center would be a logical place for contracting for the recommended studies.

130. For example, Anonymous and the Internet Watch Foundation in Britain, which uses cybertools to thwart pornography online.

131. States can directly fund or contract with commercial enterprises, finance terrorist or other low-visibility groups, supply or receive technical information or simply provide encouragement to "patriotic hackers." For example, during the 2008 Russo-Georgian War, Russian cyber activities against Georgia apparently were not conducted by military personnel but were coincident with military operations.

132. Defense Science Board, *Resilient Military Systems and the Advanced Cyber Threat*, sketches a similar hierarchy of capabilities.

133. See, for example, North American Network Operators' Group, "About NANOG," nanog.org, July 8, 2014, <https://www.nanog.org/about/home>; and Project Honey Pot, projecthoneypot.org, July 8, 2014, http://www.projecthoneypot.org/about_us.php.

134. The National Science Foundation and, abroad, the British government and the European Union have made some investments in this research. The science foundation has supported a notable collaborative effort between groups at the University of California, Berkeley, and the University of California, San Diego. The character and evolution of that work is well-described in Farrow, "Interview with Stefan Savage: On the Spam Payment Trail." There, Professor Stefan Savage traces the history of the group from the time "[w]e came to see that our community had a . . . poor understanding of the value chain for economically motivated attackers and thus didn't understand that our various technical interventions actually played minor roles, at best, in mitigating their actions. . . . [W]e started thinking much more holistically about our security work, trying in particular to understand what the underlying economic models were and how we might access those through measurement." Farrow, "Interview with Stefan Savage: On the Spam Payment Trail," 7-8. And see Savage's summary of his perspective on Page 19.

135. For a general description, see Greenberg, "Shopping for Zero-Days: A Price List For Hackers' Secret Software Exploits."

136. For a readily accessible example of what can be acquired online, see Kali Linux, an open source project that conducts digital forensics and penetration testing, <http://www.kali.org>.

137. In an exemplary analysis in 2011, Dan Guido reported that in each of the preceding two years the great majority of criminal attacks used only about a dozen exploits and virtually all of them were previously known. Guido, "The Exploit Intelligence Project."

138. Clay Wilson, "Botnets, Cybercrime, and Cyberterrorism: Vulnerabilities and Policy Issues for Congress," *CRS Report RL32174* (Washington, DC: Library of Congress, Congressional Research Service, January 29, 2008).

139. O'Gorman and McDonald, "The Elderwood Project," and Mandiant, "APT1: Exposing One of China's Cyber Espionage Units," http://intelreport.mandiant.com/Mandiant_APT1_Report.pdf, offer case study of Chinese attackers. Mandiant labeled the group APT1. The phrase is well explained at Damballa,

"Advanced Persistent Threats: A Brief Description," damballa.com, <https://www.damballa.com/knowledge/advanced-persistent-threats.php>.

140. Krekel et al., "Capability of the People's Republic of China to Conduct Cyber Warfare and Computer Network Exploitation," 59.

141. Notwithstanding Mandiant's private-sector achievement in unraveling a Chinese operation, high-end attackers also may normally best be detected by a state intelligence apparatus that can assimilate a large number of cases.

142. James A. Lewis, "Raising the Bar for Cyber Security" (Center for Strategic and International Studies, February 12, 2013), http://csis.org/files/publication/130212_Lewis_RaisingBarCybersecurity.pdf.

143. Difficulties for the Berkeley-San Diego group are described in Farrow, "Interview with Stefan Savage: On the Spam Payment Trail," 11-13 and 17.

144. U.S. Department of Homeland Security, "About the National Cybersecurity and Communications Integration Center," <http://www.dhs.gov/about-national-cybersecurity-communications-integration-center>.

145. These are loosely affiliated in a National Council of ISACs and enumerated in its website at <http://www.isaccouncil.org/aboutus.html>.

146. Models also exist abroad, as with quiet cooperation between British telecommunications companies and the British government.

147. These issues around metrics are discussed further in Appendix I.

148. Krebs, "These Guys Battled BlackPOS at a Retailer."

149. See Wayne Rosenkrans, "Sharing the Wealth," *AeroSafety World* (April 2013), 40ff, and Mitre "Fact Sheet: Aviation Safety Information Analysis and Sharing," <http://www.mitre.org/sites/default/files/publications/ASIAS1.pdf>. ASIAS is complemented by semiannual industry meetings, industry safety discussions and government review boards. It has moved recently to share data in international forums.

150. Rosenkrans, "Sharing the Wealth," 45.

151. Legislation authorizing these investments can be traced to 15 U.S.C. § 5501 "High-Performance Computing Act of 1991," <http://www.gpo.gov/fdsys/pkg/STATUTE-105/pdf/STATUTE-105-Pg1594.pdf>; 15 U.S.C. § 5513, "Next Generation Internet Research Act of 1998," <http://www.gpo.gov/fdsys/pkg/PLAW-105publ305/html/PLAW-105publ305.htm>; and 20 U.S.C. § 9801, "America COMPETES Act of 2007," <http://www.gpo.gov/fdsys/pkg/PLAW-110publ69/pdf/PLAW-110publ69.pdf>. Agencies centrally involved include the National Science Foundation; Cyber Command and its doppelgänger, the National Security Agency; the Departments of the Army, Navy and Air Force; the Defense Advanced Research Projects Agency; the CIA's Science and Technology Directorate; Intelligence Advanced Research Projects Agency; the Department of Energy (including the national laboratories that it oversees); the Department of Commerce (particularly through its subordinate agency, the National Institute of Standards and Technology); the National Institutes of Health; and the Department of Homeland Security.

152. See the House Appropriations Committee's Surveys and Investigations Report. In a 2013 report, the President's Council of Advisors on Science and

Technology comments: "The manner in which investments are aggregated and reported in the Federal Government in some cases tends to obscure the distinction between (a) R&D in [net-working and information] disciplines, and (b) the acquisition of [networking and information technology] infrastructure used to conduct R&D in other areas. . . . Misclassifying infrastructure expenditures as R&D expenditures may lead policymakers to believe that the government is investing far more in the latter than is actually the case." President's Council of Advisors on Science and Technology, *Designing a Digital Future: Federally Funded Research and Development in Networking and Information Technology* (January 2013), 18, <http://www.whitehouse.gov/sites/default/files/microsites/ostp/pcast-nitrd2013.pdf>.

153. Whatever overview exists comes from heroic and necessarily inadequate efforts from the Networking and Information Technology Research and Development Program (NITRD), a "subcommittee" of the National Science and Technology Council, which is itself nested under the White House Office of Science and Technology Policy. NITRD is staffed by three people who collect and organize relevant budget information but do not have directive authority. Their most recent summary shows a total request to Congress of almost \$4 billion for information technology and network research sponsored by federal agencies. "Subcommittee on Networking and Information Technology Research and Development," see generally <http://www.nitrd.gov/>. The President's Council of Advisors on Science and Technology observes: "Recall that there is no NITRD budget, per se. There is only a summary of NITRD-related portions of agency budgets (so-called cross-cuts), as categorized by the [agencies]." President's Council of Advisors on Science and Technology, *Designing a Digital Future: Federally Funded Research and Development in Networking and Information Technology*, 21.

154. National Commission for the Review of the Research and Development Programs, *Report of the National Commission for the Review of the Research and Development Programs of the United States Intelligence Community*, <http://www.fas.org/irp/eprint/ncrdic-cyber.pdf>.

155. The White House, "The Comprehensive National Cybersecurity Initiative," [whitehouse.gov](http://www.whitehouse.gov), 2009, <http://www.whitehouse.gov/cybersecurity/comprehensive-national-cybersecurity-initiative>.

156. The White House, "The Comprehensive National Cybersecurity Initiative."

157. Initiative No. 9 of the Comprehensive National Cybersecurity Initiative recognizes this, positing that one goal of the national effort "is to develop technologies that provide increases in cybersecurity by orders of magnitude above current systems and which can be deployed within 5 to 10 years. This initiative seeks to develop strategies and programs to enhance the component of the government R&D portfolio that pursues high-risk/high-payoff solutions to critical cybersecurity problems. The Federal Government has begun to outline Grand Challenges for the research community to help solve these difficult problems that require 'out of the box' thinking. In dealing with the private sector, the government is identifying and communicating common needs that should drive mutual investment in key research areas."

158. As four present examples, these might include opportunities for bringing safe languages to scale, further restricting abilities for attackers who circumvented sandboxes and other safeguards to acquire privileges to compromise code, enhancing accuracy and security of attribution and using

cryptography and formal methods (for example, certifying compilation) to inhibit supply chain contamination.

159. The White House, "The Comprehensive National Cybersecurity Initiative."

160. The size, character and demand for this workforce are difficult to assess and are changing. A National Academy of Sciences report attempts to describe the national picture with collateral comments on the federal cyberworkforce. See National Research Council, *Professionalizing the Nation's Cybersecurity Workforce? Criteria for Decision-Making* (Washington, DC: The National Academies Press, 2013), 7, http://www.nap.edu/catalog.php?record_id=18446.

161. The President's Review Group on Intelligence and Communications Technologies, *Liberty and Security in a Changing World* (December 12, 2013), http://www.whitehouse.gov/sites/default/files/docs/2013-12-12_rg_final_report.pdf.

162. The morale and recruitment problems inflicted by Edward Snowden's disclosures do not improve this situation. They only further weaken the federal government generally.

163. This person's title is "Special Assistant to the President and Cyber Coordinator."

164. Defense Science Board, *Resilient Military Systems and the Advanced Cyber Threat*, 9.

165. Geer and Pareek, "The Index of Cyber Security: Annual Report," 18, a survey of senior IT professionals in the private sector, found that half subscribe to the proposition that "Qualified [information security professionals] are difficult or impossible to find regardless of their compensation."

166. I am indebted to Lara Schmidt for observing that the two fields are sharply differentiated in the private sector. With guidance from the federal Chief Information Officers Council, specialized cybersecurity positions have been defined within the IT field. Illustrations are provided by a National Initiative for Cyber Education presentation, slides 9 and 10, at <http://www.isaca-washdc.org/presentations/2013/201304-session6.pdf>.

167. Similar difficulties burden the uniformed military.

168. Daniel Manson and Ronald Pike, "The Case for Depth in Cybersecurity Education," *acm Inroads*, 5 no. 1 (March 2014), 47, <http://niccs.us-cert.gov/sites/default/files/documents/files/ACM%20Inroads%20The%20Case%20for%20Depth%20in%20Cybersecurity%20Education.pdf>.

169. 10 U.S.C. § 1601, "Civilian intelligence personnel: general authority to establish excepted positions, appoint personnel, and fix rates of pay," and the provisions immediately following that.

170. "Cyber In-Security: Strengthening the Federal Cybersecurity Workforce" (Partnership for Public Service and Booz Allen Hamilton, July 2009), 4, 8 and 11, http://www.boozallen.com/media/file/CyberIn-Security_2009.pdf, observed the difficulties civilian agencies have in recruiting and retaining cyberskilled employees in the face of competing offers from the NSA.

171. This is true of the NSA as well.

172. Sometimes they return as contractors paid at substantially higher levels, but contractor profit and overhead make them much more expensive and contractors cannot assume certain important duties.

173. Six of the last 20 left after two years, according to the author's informal sample.

174. FFRDC employees are not part of the federal civil service system. They work on contract for the federal government and sometimes move into regular federal appointments for terms of two or four years under the Intergovernmental Personnel Act.

175. I am vice chairman of the RAND Corporation, which operates three FFRDCs sponsored by the Department of Defense, and several years ago worked for the Center for Naval Analyses (CNA), which operates an FFRDC sponsored by the U.S. Navy. I retain the title of senior adviser at CNA.

176. Assante, testimony to the Committee on Homeland Security and Governmental Affairs, 8: "I have never understood why we have not embraced better training and development methods for our frontline [information technology] security and operations staff. We train pilots using advanced simulators to deal with difficult conditions and mechanical failures. Why do we not use simulators to allow security and operational staff to experience low frequency but high consequence attacks against the systems they defend and operate? Why do we not use performance-based examinations to qualify our most important resources?"

177. Among them, MITRE and the Institute for Defense Analyses have notably provided this service. The National Institute of Standards and Technology proposes to create an additional such FFRDC. The National Laboratories also provide valuable support.

178. Called for by "Securing Cyberspace for the 44th Presidency" (Center for Strategic and International Studies, December 2008), http://csis.org/files/media/csis/pubs/081208_securingcyberspace_44.pdf; and "Cyber In-Security: Strengthening the Federal Cybersecurity Workforce."

179. On this subject generally, see National Research Council, *Professionalizing the Nation's Cybersecurity Workforce? Criteria for Decision-Making*, 14.

180. In a 2011 speech, Frank Kramer recommended a "think tank skunk works" in an effort to achieve "analysis, operations and technology development." Franklin Kramer, "Cyber Conflict: Challenging the Future" (Black Hat Conference, Washington, January 18, 2011), transcript at <http://www.atlanticcouncil.org/news/transcripts/franklin-kramer-us-should-aim-for-cyber-resilience>. Others have made similar recommendations. The focus here is on enriching the workforce available for cyber defense for the federal government.

181. The National Security Council staff is subject to outreach and security procedures that are stricter than for the White House generally.

182. Freeman Dyson, *Disturbing the Universe* (New York: Basic Books, 1981).



J U L Y 2 0 1 4

Surviving on a Diet of Poisoned Fruit
Reducing the National Security Risks of America's Cyber Dependencies





Appendix

A NOTE ON DEFICIENCIES OF DATA AND METRICS

53

APPENDIX: A NOTE ON DEFICIENCIES OF DATA AND METRICS

It is difficult to drive when you cannot see the road. Data about cyber vulnerabilities, attacks and consequences are sparse, irregularly collected,¹ drawn from changing populations, inconsistently labeled and aggregated,² erratically shared and published,³ and most commonly analyzed by parties with an economic interest in conclusions that are drawn.⁴ International data sharing is rudimentary among even close allies and distrustful across political and cultural barriers. Victims often imperfectly comprehend what they have lost – or that they have been attacked.⁵ When aware, any inclination companies may have to share their experience is inhibited by fear that doing so will give away competitive information, hurt their reputations, frighten their customers, lower their stock price, induce more attacks,⁶ invite lawsuits and/or trigger regulatory scrutiny.⁷ Government officials keep much of their knowledge confidential or classified lest their sources and methods be discerned, international tensions be inflamed,⁸ prosecutions be compromised or they inadvertently violate privacy restrictions or legal boundaries between what is permissible abroad and at home. Most perpetrators have incentives to cover their tracks.⁹ Further confusing things, some perpetrators (hackers, for instance) hype their capabilities while other actors understate their sources and insight.¹⁰

As a result of these difficulties, most cybersecurity information is indeterminate, inconsistent, over-interpreted or all three. A careful evaluation by Microsoft researchers concludes: “Our assessment of the quality of cyber-crime surveys is harsh: they are so compromised and biased that no faith whatever can be placed in their findings. We are not alone in this judgment. Most research teams who have looked at the survey data on cyber-crime have reached similarly negative conclusions.”¹¹ Butressing this view, they note

that “Ryan and Jefferson, who perform a meta-study of fourteen cyber-crime surveys, write ‘In the information security arena, there is no reliable data upon which to base decisions. Unfortunately, there is unreliable data that is masquerading as reliable data.’”¹² More recently, two leading authorities on vulnerability data pithily summarized their view in a subtitle: “Why Vulnerability Statistics Suck.”¹³ Among other things, their essay observed:

As maintainers of two well-known vulnerability information repositories, we’re sick of hearing about research that is quickly determined to be sloppy after it’s been released and gained public attention. In almost every case, the research casts aside any logical approach to generating the statistics. They frequently do not release their methodology, and they rarely disclaim the serious pitfalls in their conclusions.¹⁴

Perhaps the strongest manifestation of disaffection with reported data came from Carnegie Mellon University’s Software Engineering Institute when, in 2008, it ended a decade long practice of collecting and publishing vulnerability information.¹⁵

The absence of a baseline has strong adverse consequences for individual, corporate and national decisionmaking. Because past¹⁶ and present reality cannot credibly be ascertained and future circumstances cannot be predicted, insurance cannot reasonably be offered¹⁷ and the return on investment for security expenditures cannot be well-established.¹⁸ Without a baseline, government, nongovernmental organizations and private-sector assessments are unconstrained in reflecting their political, ideological and commercial agendas rather than logical inferences. Credible, shared metrics¹⁹ for cyber success, failure, protective efforts

and vulnerability are only beginning to be developed.²⁰ Even a common vocabulary has proved elusive.

In the absence of agreed data and metrics, some think cybersecurity problems are “hyped”; others think that they are woefully understated. Without a common framework the field fragments. Silos are reinforced by tool-centered thinking:²¹ Those invested in a particular problem see virtues of that approach in all circumstances. As Bill Crowell nicely puts it, “Cybersecurity is a thousand points of light, together yielding no illumination.”²²

ENDNOTES

1. Two exceptions are data collected by states when cyber breaches reveal consumer information (see <http://www.ncsl.org/issues-research/telecom/security-breach-notification-laws.aspx>) and Department of Defense compilations of breaches of industrial contractors.
2. “*Abstraction bias* is a term that we crafted to explain the process that VDBs [vulnerability databases] use to assign identifiers to vulnerabilities. Depending on the purpose and stated goal of the VDB, the same 10 vulnerabilities may be given a single identifier by one database, and 10 identifiers by a different one. This level of abstraction is an absolutely critical factor when analyzing the data to generate vulnerability statistics. This is also the most prevalent source of problems for analysis, as researchers rarely understand the concept of abstraction, why it varies and how to overcome it as an obstacle in generating meaningful statistics. Researchers will use whichever abstraction is most appropriate or convenient for them; after all, there are many different consumers for a researcher advisory, not just VDBs. Abstraction bias is also frequently seen in vendors, and occasionally researchers in the way they disclose one vulnerability multiple times, as it affects different software that bundles additional vendor’s software in it.” Steve Christey and Brian Martin, “Buying Into the Bias: Why Vulnerability Statistics Suck,” *media.blackhat.com*, July 11, 2013, <https://media.blackhat.com/us-13/US-13-Martin-Buying-Into-The-Bias-Why-Vulnerability-Statistics-Suck-WP.pdf>.
3. Exceptions of some value are referenced throughout this paper. See, for example, Verizon’s annual “Data Breach Investigations Reports” and IBM’s “X-Force Trend and Risk Reports.”
4. Ross Anderson’s work is a notable exception and accordingly referred to at several points in this paper. Noting that rapid changes may be occurring in cybercrime, Anderson and his co-authors say: “[W]e believe that our work is a principled start to being able to measure the cost of cybercrime. We propose to continue updating our estimates, and to produce new versions of this paper every few years.” Ross Anderson et al., “Measuring the Cost of Cybercrime” (paper presented at the Workshop on the Economics of Information Security, Berlin, June 25–26, 2012), 25, http://weis2012.econinfosec.org/papers/Anderson_WEIS2012.pdf.
5. Discussing breaches, methods, motives and attackers, Verizon’s “2013 Data Breach Investigations Report” comments that “All of the above still takes forever and a day to discover, and that discovery is rarely made by the victim.” Verizon, “2013 Data Breach Investigations Report” (Verizon, 2013), 6, http://www.verizonenterprise.com/resources/reports/rp_data-breach-investigations-report-2013_en_xg.pdf.
6. After distributed-denial-of-service attacks attributed to Iran “knocked at least five . . . banks’ Web sites off-line . . . PNC bank C.E.O. James Rohr stated that ‘we had the longest attack of all the banks’ and warned that ‘cyber-attacks are a very real, living thing, and if we think we are safe that way, we’re just kidding ourselves.’ Shortly afterward, the attacks on PNC escalated, causing further problems. Neither Rohr nor any other high-level executive of any victim bank has since made any such conspicuous and pointed statement. ‘The lesson from Rohr’s statement was, don’t talk,’ says one former national-security official.” Michael Joseph Gross, “Silent War,” *Vanity Fair* (July 2013), <http://www.vanityfair.com/culture/2013/07/>

new-cyberwar-victims-american-business. Kevin Mitnick recounts: "Pacific Bell eventually found out about the access we had gained. Yet we were never arrested and charged because, I later learned, company management was afraid of what would happen if others found out what I had been able to do and started trying to duplicate my efforts." Kevin Mitnick, *Ghost in the Wires: My Adventures as the World's Most Wanted Hacker* (New York: Little, Brown and Company, 2011), 56.

7. A reporter who has been vigorous in tracking attacks writes: "Most companies have preferred not to talk about or even acknowledge violations of their computer systems, for fear of panicking shareholders and exposing themselves to lawsuits — or for fear of offending the Chinese and jeopardizing their share of that country's exploding markets." Michael Joseph Gross, "Enter the Cyber-dragon," *Vanity Fair* (September 2011), <http://www.vanityfair.com/culture/features/2011/09/chinese-hacking-201109>.

8. A recent 45-minute Australian news report features experts and officials describing cyberthefts and penetrations said to be Chinese but not attributed as such for fear of repercussions and disclosing sources and methods. Andrew Fowler and Peter Cronau, "Hacked!," *Four Corners*, May 27, 2013, <http://www.abc.net.au/4corners/stories/2013/05/27/3766576.htm>.

9. "[T]errorist crimes are hyper-salient because the perpetrators go out of their way to be as annoying as possible, while most online crooks go out of their way to be invisible." Anderson et al., "Measuring the Cost of Cybercrime," 26.

10. *Vanity Fair* reports "a surreal new creation of bureaucracy: government-directed 'hacktivism,' in which an intelligence agency secretly provides information to a group of private-sector hackers so that truths too sensitive for the government to tell will nevertheless come out." Gross, "Enter the Cyber-dragon."

11. Dinei Florencio and Cormac Herley, "Sex, Lies and Cyber-crime Surveys," MSR-TR-2011-75 (Microsoft, June 2011), 6, <http://research.microsoft.com/apps/pubs/default.aspx?id=149886>. Peter Maass recounts how various unfounded dollar estimates for cyber loss have made their way into statements by officials, including President Barack Obama and the head of Cyber Command. Peter Maass and Megha Rajagopalan, "Does Cybercrime Really Cost \$1 Trillion?," *ProPublica.org*, Aug. 1, 2012, <http://www.propublica.org/article/does-cybercrime-really-cost-1-trillion>.

The Verizon and IBM data summaries do a better-than-average job of noting biases in their data. For example, Verizon writes: "The number of . . . contributors [to the data pool] has doubled each year since 2009 and tripled for this current edition. We want to stress the importance of this point because changes in overall trends (like threat actions observed year over year) may be caused by changes within the sample set or in the threat environment (or both). With that out of the way, we can get on to the data." Verizon, "2013 Data Breach Investigations Report," 23. See also "IBM X-Force 2012 Trend and Risk Report" (IBM, March 2013), 19.

12. *Ibid.*

13. Christey and Martin, "Buying Into the Bias: Why Vulnerability Statistics Suck."

14. *Ibid.*, 1.

15. The data through the third quarter of 2008 comes from CERT (Computer Emergency Response Team), a division of the Software Engineering Institute at Carnegie Mellon University. CERT maintains information databases of information regarding software vulnerabilities and malicious code. See "The CERT Division," Software Engineering Institute at Carnegie Mellon University, <http://www.cert.org/>.

16. Jason Healey describes a collection of essays he edited as "the only major attempt in twenty-five years to codify [cybersecurity] history." He continues:

"In other areas of national security, new military personnel, diplomats, and policymakers are taught to avoid old mistakes through a formal study of history. . . . Just as we teach young cadets and military officers the implications of Gettysburg . . . so too must we pass along the lessons of [cyber attacks]. Yet the opposite has been the case. Cyber history has been forgotten, ignored as irrelevant, or intentionally falsified, even as a crush of new personnel flood into the field." Jason Healey, ed., *A Fierce Domain: Conflict in Cyberspace, 1986 to 2012* (Arlington, VA: Cyber Conflict Studies Association, 2013), 14.

17. Lloyd's of London now offers a few policies, but they are exorbitantly priced. "At the keynote at Workshop on Economics of Information Security (WEIS) 2010, Tracey Vispoli, VP and head of CyberSecurity Infrastructure at Chubb Insurance, stated that the insurance industry has 'no expected loss data and no financial impact data.'" Florencio and Herley, "Sex, Lies and Cyber-crime Surveys," 9.

18. Deirdre Mulligan and Fred Schneider add that "Companies and individuals do not know how to value (i) confidentiality of information, (ii) integrity of information, or (iii) the pain of dealing with recovery from an attack's effects (e.g., bad credit ratings)." Deirdre K. Mulligan and Fred B. Schneider, "Doctrine for Cybersecurity," *Daedalus*, 140 no. 4 (Fall 2011), http://www.mitpressjournals.org/doi/abs/10.1162/DAED_a_00116#.U7EMoig7szl. Georgetown University's annual Workshop on the Economics of Information Security has, for the past dozen years, produced papers discussing these issues. See, for example, the 2013 agenda at <http://weis2013.econinfosec.org/program.html>.

19. The cybersituation was well-described (and the topic richly explored) by Andrew Jaquith in *Security Metrics: Replacing Fear, Uncertainty and Doubt* (Upper Saddle River, NJ: Pearson Education Inc., 2007), 20-21.

The ideal benchmarking data are cheap to gather, are expressed as numbers, contain units of measure, are objectively and consistently gathered, and are relevant to a decision-maker.

Information security has no equivalent . . . of the time-honored tradition of benchmarking organizational performance. Analytical rigor receives little attention, while nebulous, non-quantitative mantras rule: "defense in depth," "security is a process," and "there is no security by obscurity," to name a few. The numbers that do exist, such as those provided in vulnerability and threat reports from Symantec . . . and others, provide macro-level detail about the prevalence of malware or missing patches, but little else that enterprises can use to assess their effectiveness comparatively against others. Numbers provided by anti-malware, vulnerability management systems, and SIM/SEM

[Security Incident Management/Security Event Management, now commonly abbreviated as SIEM] systems certainly add value, but to date, no entity has yet attempted to aggregate and compare these data across enterprises.”

20. See, for example, MITRE’s efforts at compiling a database using consistent terminology described at <http://cve.mitre.org/> and <http://www.first.org/cvss>. Cigital’s “Building Security in Maturity Model” (BSIMM) is a leading effort to describe protective practices consistently across enterprises. Participating firms include Bank of America, Microsoft and Visa. “BSIMM describes the set of activities practiced by fifty-one of the most successful software security initiatives in the world. In that sense, it is a de facto standard because it’s what organizations actually do.” See <http://bsimm.com/facts/>.

21. I am grateful to Michael Assante for suggesting this phrase.

22. Interview, July 18, 2013.

About the Center for a New American Security

The mission of the Center for a New American Security (CNAS) is to develop strong, pragmatic and principled national security and defense policies. Building on the expertise and experience of its staff and advisors, CNAS engages policymakers, experts and the public with innovative, fact-based research, ideas and analysis to shape and elevate the national security debate. A key part of our mission is to inform and prepare the national security leaders of today and tomorrow.

CNAS is located in Washington, and was established in February 2007 by co-founders Kurt M. Campbell and Michèle A. Flournoy. CNAS is a 501(c)3 tax-exempt nonprofit organization. Its research is independent and non-partisan. CNAS does not take institutional positions on policy issues. Accordingly, all views, positions, and conclusions expressed in this publication should be understood to be solely those of the authors.

© 2014 Center for a New American Security.

All rights reserved.

Center for a New American Security

1152 15th Street, NW
Suite 950
Washington, DC 20005

TEL 202.457.9400
FAX 202.457.9401
EMAIL info@cnas.org
www.cnas.org

Production Notes

Paper recycling is reprocessing waste paper fibers back into a usable paper product.

Soy ink is a helpful component in paper recycling. It helps in this process because the soy ink can be removed more easily than regular ink and can be taken out of paper during the de-inking process of recycling. This allows the recycled paper to have less damage to its paper fibers and have a brighter appearance. The waste that is left from the soy ink during the de-inking process is not hazardous and it can be treated easily through the development of modern processes.





Center for a
New American
Security

STRONG, PRAGMATIC AND PRINCIPLED
NATIONAL SECURITY AND DEFENSE POLICIES

1152 15th Street, NW
Suite 950
Washington, DC 20005

TEL 202.457.9400 www.cnas.org
FAX 202.457.9401
EMAIL info@cnas.org

ISBN 978-1-935087-88-5

5 0 9 9 9 >



9 781935 087885



Printed on Post-Consumer Recycled paper with Soy Inks