Raytheon
BBN Technologies

10 Moulton Street
Cambridge, MA
02138
617.873.8000
www.bbn.com

20 March 2015

National Telecommunications and Information Administration
U.S. Department of Commerce
1401 Constitution Avenue NW, Room 4725
Attn: Cybersecurity RFC 2015
Washington DC 20230

Dear NTIA RFC Team:

In response to your request for comments regarding Stakeholder Engagement on Cybersecurity in the Digital Ecosystem, I wanted to pass along the attached report *Interdisciplinary Pathways towards a More Security Internet,* produced by the NSF-sponsored Cybersecurity Ideas Lab (workshop) held in Arlington, Virginia, on February 10-12, 2014. I was the chair of the workshop and write to you in that capacity.

The workshop was a multi-disciplinary event that sought to identify high-impact actions that could be taken to improve the security of the Internet ecosystem. In some ways the NSF workshop's goals complement those of your RFC. The workshop sought to reflect the interests of multiple stakeholders. The workshop also sought actionable ideas that would benefit the digital ecosystem.

In this light, I think you may find reading the report useful as a source of potential ideas and perspectives. If you have any questions about the report, the workshop process that engendered it, or feedback the workshop team has received since issuing the report, please do not hesitate to contact me.

Sincerely,

Dr. Craig Partridge

Chief Scientist, Raytheon BBN Technologies

Encl: Copy of report [The report is also available electronically at *http://www.nsf.gov/cise/news/CybersecurityIdeasLab_July2014.pdf*]

Interdisciplinary Pathways towards a More Secure Internet

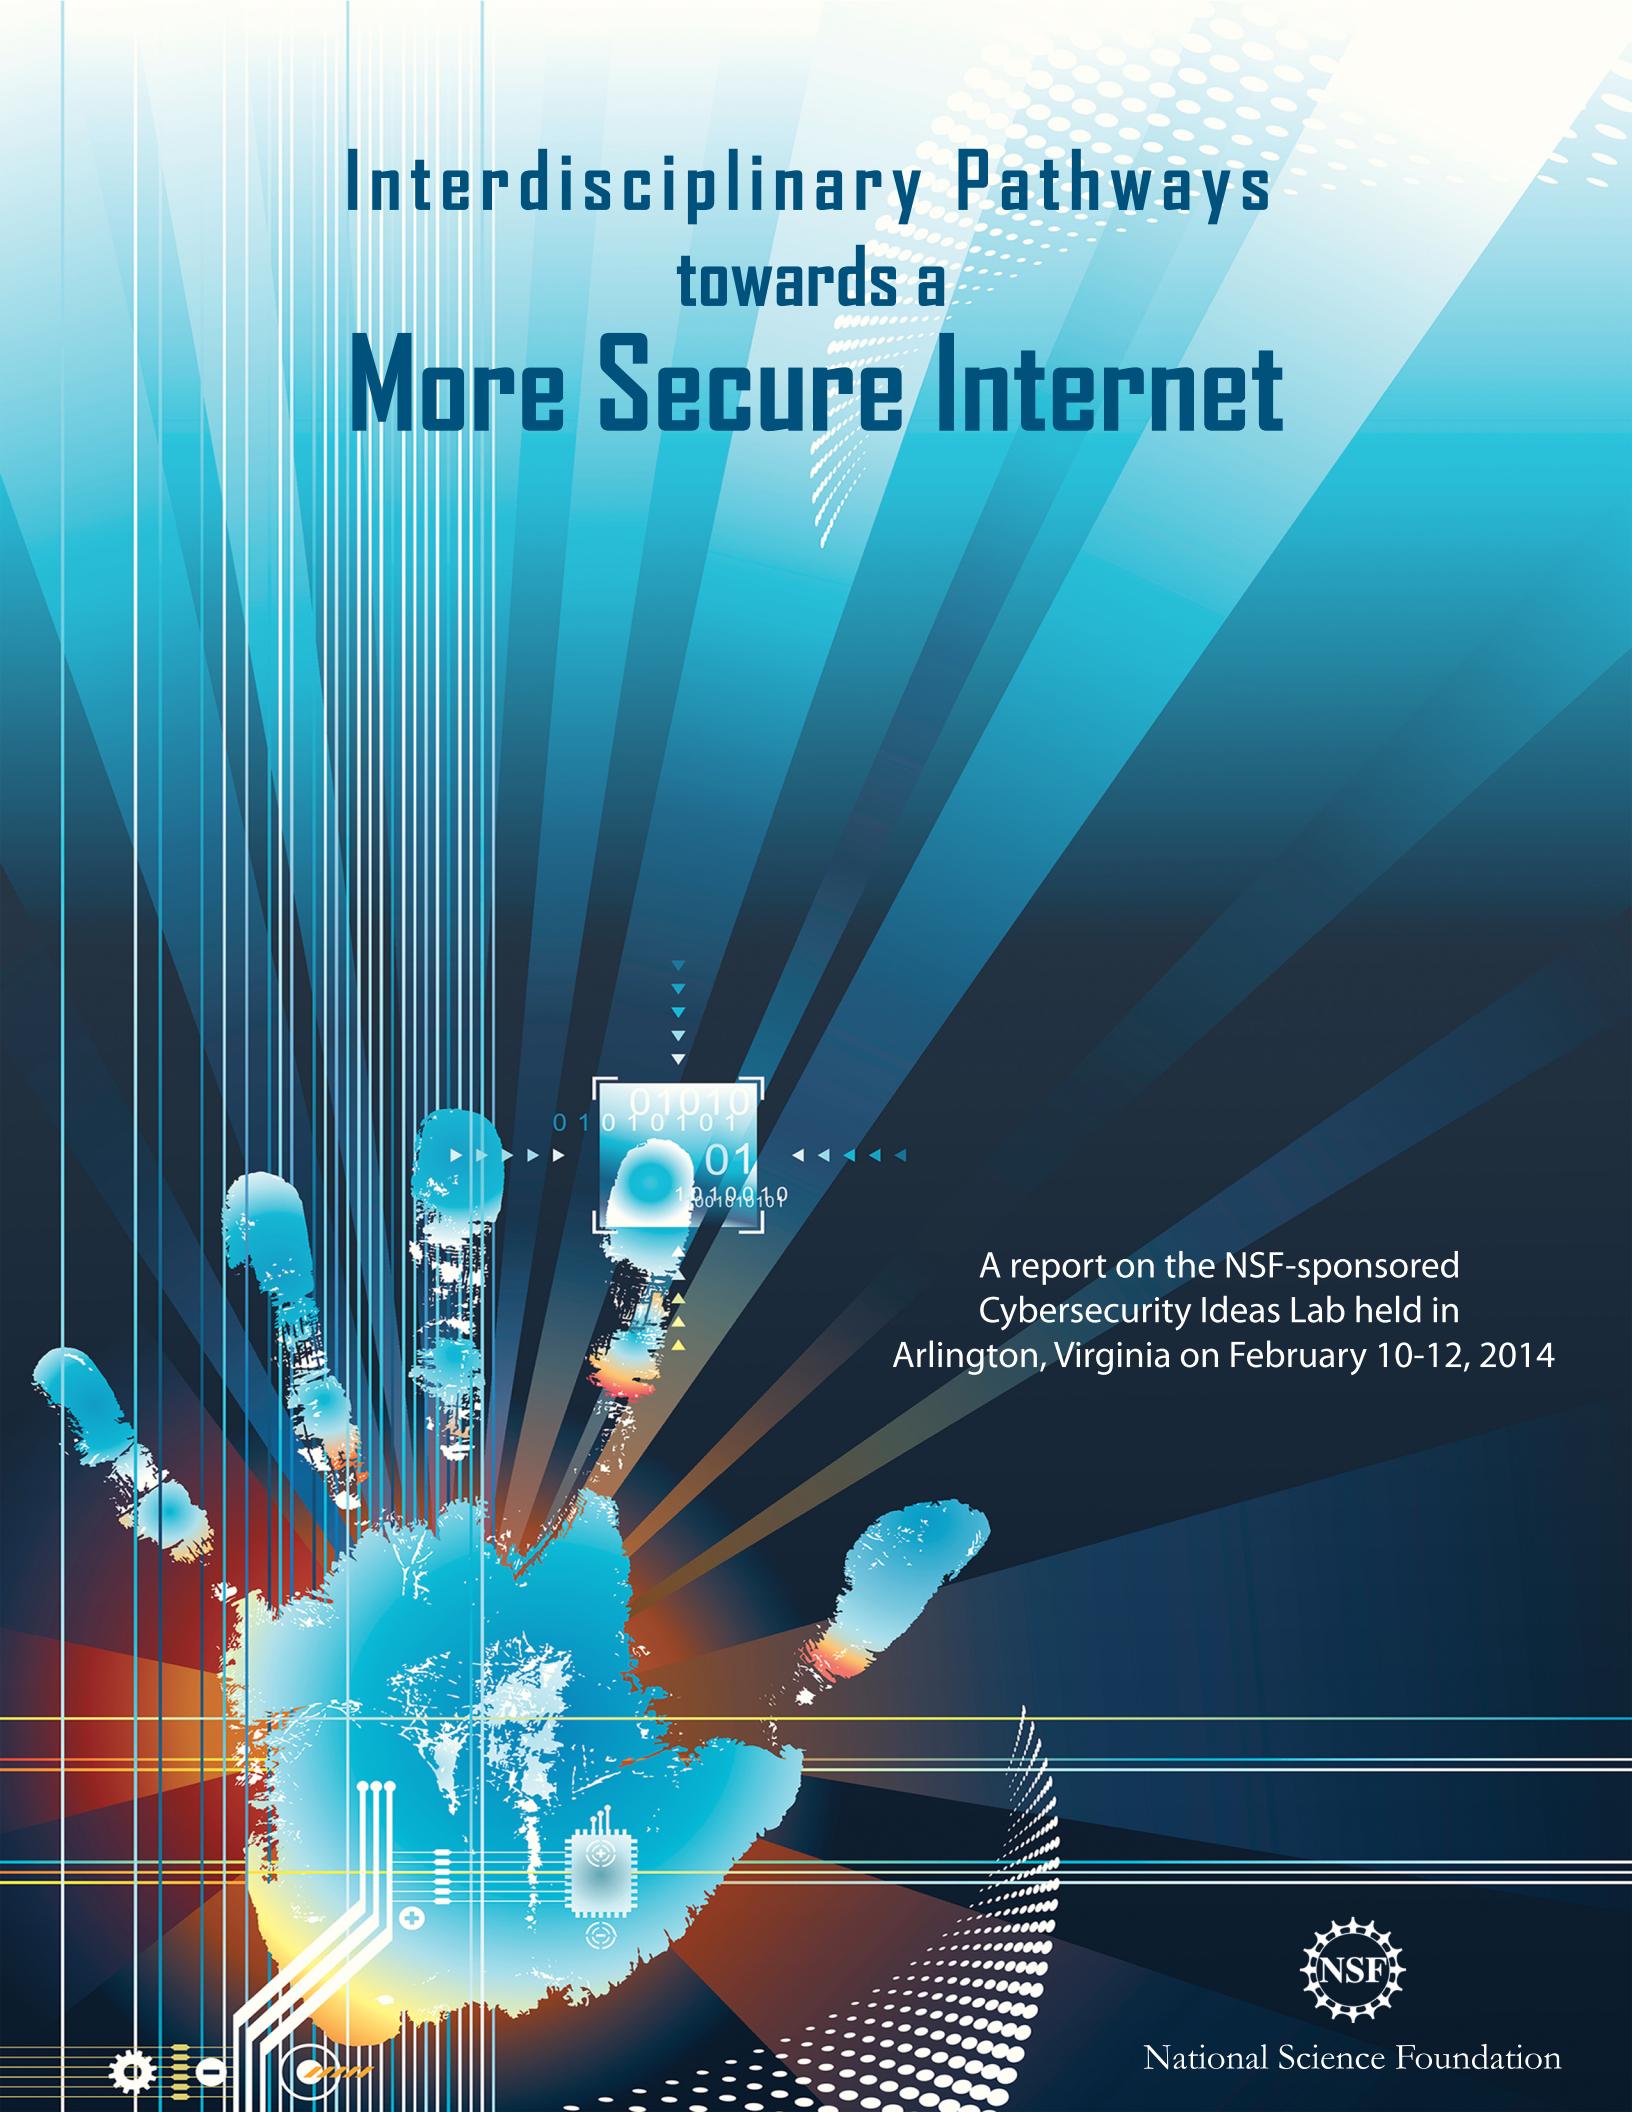

A report on the NSF-sponsored Cybersecurity Ideas Lab held in Arlington, Virginia on February 10-12, 2014

National Science Foundation

Interdisciplinary Pathways towards a More Secure Internet

A report on the NSF-sponsored Cybersecurity Ideas Lab held in Arlington, Virginia on February 10-12, 2014

This report is the output from the Cybersecurity Ideas Lab held from February 10-12, 2014, in Arlington, Virginia, sponsored through NSF Grant Number 1418289. Any opinions, findings, conclusions, or recommendations expressed in this report are those of Ideas Lab participants and do not necessarily represent the official views, opinions, or policy of the National Science Foundation.

Foreword

The Internet ecosystem plays a vital role in tightly integrating the economic, political, and cultural fabric of society. These interdependencies leave the Nation vulnerable to a wide range of threats, with potential impacts extending into nearly every facet of daily life.

The security of cyberspace depends upon a range of interconnected factors, ranging from foundational knowledge in computing and communications to the policies that govern the use of the Internet and related technologies. Transformative advances in cybersecurity require a holistic approach.

To help accelerate this process, the National Science Foundation sponsored a unique, interdisciplinary workshop, the *Cybersecurity Ideas Lab*, from February 10-12, 2014, in Arlington, Virginia. This event convened participants with expertise ranging from computer science to economics to law, from academia, government and the private sector, to develop actionable ideas with the potential to significantly enhance the security of the Internet ecosystem.

The group was charged to think in both the long- and short-term; to consider both radical innovations and endorsement of existing ideas; and to recommend actions in any sector, not limited to research and development activities. Many exciting ideas emerged at the workshop; the most developed are presented in the form of recommendations in the following report.

On behalf of the National Science Foundation, I would like to thank the workshop participants for their contributions at this event. In particular, I'd like to acknowledge Craig Partridge for his leadership in serving as the workshop Director, as well as the rest of the workshop Steering Committee: Susan Landau, Damon McCoy, Deirdre Mulligan, Jennifer Rexford, Stefan Savage and Dave Ward. This team provided a starting point, charge and guidance for the workshop activities.

I would also like to acknowledge David Clark for serving as provocateur to stimulate new discussion pathways; Damon McCoy, the lead PI, for organizing the event; Keith Marzullo for his leadership and contributions in the planning process; and Emily Grumbling for many contributions to the coordination of the workshop.

The *Cybersecurity Ideas Lab* and its output will help to advance the national dialogue around cybersecurity. The recommendations that follow are intended to catalyze significant steps towards enhancing the security of the Internet ecosystem.

Farnam Jahanian
Assistant Director
Directorate for Computer and Information Science and Engineering
National Science Foundation
July, 2014

Table of Contents

EXECUTIVE SUMMARY

Today, our lives are intimately connected to the Internet. This integration of life and network is ongoing. Indeed, while networks were originally used for computers and their users to communicate, we now have commonplace objects (such as thermostats, light bulbs and kitchen appliances) that communicate with each other and with us. This trend is accelerating, leading to what is called the "Internet of Things."

Unfortunately, networked systems, as well as the networks themselves, are vulnerable to attack or disruption. Since the late 1980s, we have seen attacks that affect thousands – and now as many as hundreds of thousands – of people and cause many millions of dollars in damage.

Over time we have come to realize that the problem of security in our networks and networked systems is a multidisciplinary problem, touching on policy issues, economic incentives, and public and business awareness and education, along with new technical challenges.

In mid-February 2014, the National Science Foundation (NSF) sponsored a 2½-day workshop centered on identifying high-impact actions that could be taken to better secure the Internet. The workshop, the Cybersecurity Ideas Lab, brought together 35 invited experts in computer science, cybersecurity, economics, social science and policy (see Appendix). These experts were drawn from industry, academia and the government.

The workshop leveraged the NSF Ideas Lab process with a twist. Where the typical Ideas Lab seeks to identify research ideas or thrusts, this workshop combined professional facilitators with a diverse set of experts to identify actionable ideas that would lead to a more secure Internet. The Ideas Lab was coordinated by a seven-member Steering Committee (see Appendix).

The goal of the workshop was to bring forward a suite of meritorious ideas, not a consensus report. An individual participant should not be assumed to endorse all the ideas in this report.

The workshop participants generated a large number of ideas for high-impact actions that could potentially be adopted by cybersecurity researchers, policymakers or practitioners to advance cybersecurity. The workshop refined, combined and adapted these ideas to yield a practical working set of recommendations. The Steering Committee assembled and, through multiple review cycles, refined the ideas into the sixteen actionable recommendations presented below. While all of the recommendations are interdisciplinary, for the purpose of this report they are grouped into three categories: *technology*, *policy* and *leadership* (educational opportunities also regularly appear within individual recommendations).

Technology

Make Critical Subsystems Field-Updatable	Long-lived, often critical, systems such as embedded systems and industrial control systems have not typically been designed with updatable security capabilities. Mechanisms to update such systems, including current cyber-physical systems (CPS) and industrial control systems, are needed rather than the current system where some industrial control systems are designed to be locked down once deployed. One of the reasons such systems are locked down is the tremendous impact if they fail to work as expected. Any update system must address this challenge.
Enable Certificate Transparency and Security	Certificates are used to associate the entity that registered and owns a domain with a public key that can be used to protect communications with that domain. Certificates are the primary mechanism for authenticating that Internet connections are made with the entities they claim to represent. Thousands of certificate authorities issue and manage these certificates. These authorities can intentionally or unintentionally issue incorrect or false certificates, enabling eavesdropping on communications or even impersonation of the domain owner. Transparency and trustworthiness should be increased, either by endorsing efforts such as those by *certificate-transparency.org* or through other mechanisms.
Create a Framework for Managing Software Updates	The Internet of Things will challenge our current channels for distributing security updates. An environment must be developed for distributing security patches that scales to a world where almost everything is connected to the Internet and many "things" are largely unattended.
Make HTTPS the Least-Effort Scheme for Deploying a Website	The secure version of the Web's primary protocol should be made the usual way that servers interact with clients.
Cybersecurity Research Agenda	Key research thrusts are identified: certifying the security properties of future systems, security of software-defined networks, R&D on clean slate operating systems and system services, security and privacy on the cheap and consequences of aligning network structure with political jurisdiction.

Policy

Establish an Internet Rescue Squad	Implement a national cybersecurity response team to coordinate responses to cyber breaches and attacks, prompt secondary targets to initiate response programs and to support response activities.
Create a Cyber NTSB	Create a cyber analogue to the National Transportation Safety Board charged with analyzing cybersecurity incidents and providing public reports on the circumstances and causes of each incident.
"Standard" Impact Statement	Develop models for characterizing the cost and impacts, both positive and negative of cybersecurity frameworks and standards. Before new standards are implemented, their impact should be assessed in a process analogous to generation of an Environmental Impact Statement.
How Golden is Our Goose?	The computing and communications community has been a golden goose that has repeatedly produced economic and social gains. Security requirements are often seen to stifle innovation or usability of technology. Tradeoffs need to be better understood and used to assess whether the economic harm of securing or protecting from breaches may exceed its benefit.
Identity: A Problem That Doesn't Need Solving	Personally identifiable packet-level attribution will not address the most pressing cybersecurity issue: cross-jurisdictional cyber-exploitation. No new efforts should be undertaken in this area. On the other hand, the National Institute of Standards and Technology's efforts through the National Strategies for Trusted Identities in Cyberspace to improve authentication technologies are useful and should be funded.
Encourage the Adoption of Routing Security	A serious obstacle to better securing the Internet's routing and naming systems is the demonstrated desire of governments to leverage the routing and naming system to "take down" global internet sites for violations of national laws. Rather than being a prime offender in this regard, the U.S. Government should curb this behavior.
Enhance the Security of the Internet of Things by Identifying Enclaves	The security challenges posed by the emerging Internet of Things should be addressed now, to prepare before it is fully upon us. By identifying specific use segments, or "enclaves," Internet of Things infrastructure stakeholders can address the security requirements and devise event remediations for that enclave.

Leadership

Create a List of Top Priorities	A well-considered list of top priorities in cybersecurity is lacking. The Federal Government should leverage its leadership position to gather inputs from major businesses and enterprises to jumpstart the creation of such a list.
Lead By Example	With one of the largest IT infrastructures in the country, the Federal Government needs to embrace its role as a leading technology organization by adopting best cybersecurity practices in its IT systems.
Re-Establish trust in NIST's Cryptographic Standards Process	The U.S. should reestablish the credibility of NIST as an honest broker of cryptographic and security standards. This could be done via a rigorous external review of NIST's cryptographic standards process and a public commitment from the Government that NIST's security standards will not be subverted.
Develop Citizen and Small Business How-to-Guides for Implementing Security	Current cybersecurity guidance is primarily targeted towards large corporations and the technically savvy user. A set of clear, interactive guides aimed at individual citizens and small and medium businesses should be produced to demonstrate best cybersecurity practices.

While this list of recommendations is by no means comprehensive and addresses only some key areas of the problem space, it maps out concrete possible next steps that could have real impact across all sectors. It is hoped that these recommendations will stimulate new policies, drive new technical innovations, and encourage awareness and adoption of identified best practices for cybersecurity.

INTRODUCTION

Cybersecurity is a National Priority

With the rapid pace of technological advancement, daily life is now intimately connected to the Internet. Critical portions of business operations, financial systems, manufacturing supply chains and military systems are also networked. Indeed, while it was originally computers and their users that communicated over networks, we now have commonplace objects (such as thermostats, light bulbs and kitchen appliances) that communicate with each other and with us. This trend is accelerating, leading to what is called the "Internet of Things."

Unfortunately, many networked systems as well as the networks themselves are vulnerable to attack or disruption. Since the late 1980s, attacks have occurred that affect thousands – and now as many as tens or hundreds of thousands – of people and cause many millions of dollars in damage.

Cybersecurity has become relevant to nearly every aspect of today's society. Over time, experts in this field have come to realize that the problem of security in our networks and networked systems is a multidisciplinary problem, touching on policy issues, economic incentives and public and business awareness, along with purely technical challenges.

This report is the product of a 2 ½- day workshop, sponsored by the National Science Foundation (NSF), to identify actionable interdisciplinary pathways towards a more secure Internet. The workshop was planned and coordinated by a seven-member Steering Committee, including Susan Landau, Damon McCoy, Deirdre Mulligan, Craig Partridge, Jennifer Rexford, Stefan Savage and Dave Ward, with contributions from David Clark. The event itself was facilitated by a team of experts from the creativity firm *KnowInnovation*, who coordinated the activities in consultation with the Steering Committee, led by Craig Partridge, who served as workshop Director. The 35 invited participants in attendance represented academia, industry and government, with expertise ranging from computer science to policy. (See Appendix for list of participants.)

The Ideas Lab Process

An Ideas Lab is an intensive, immersive, multi-day retreat convening leading thinkers to scope a grand challenge and charter radically innovative paths forward. The participants, a diverse group of experts and stakeholders, provide a range of perspectives and experiences, and generate new, cross-disciplinary ideas that would not otherwise emerge. This result is achieved through a collaborative brainstorming, debate, critique, revision and crafting of high-risk, high-reward – or otherwise impactful – strategies and solutions. The most transformative ideas are chosen for further development or implementation.

A typical NSF-run Ideas Lab convenes a group of experts to brainstorm and refine new, innovative ideas for research around a specific topic or problem, with the best ideas selected for NSF funding. The *Cybersecurity Ideas Lab* was different. The workshop used elements of the Ideas Lab process to generate

recommendations for what could be done in any sector to make cyberspace more secure than we find it today; recommendations were not limited to new research efforts.

Workshop Goal

The goal of the Cybersecurity Ideas Lab was to generate actionable ideas for improving cybersecurity, in the form of both short- and long-term recommendations. The participants were encouraged to consider both brand-new ideas and new endorsements of existing tools or policies that would have significant impact if more broadly adopted. Recommendations could range from innovative and unconventional to practical and easy to implement.

Participants were encouraged to address specific problems, by proposing solutions or strategies grounded in fields as diverse as technology, policy, education, economics and sociotechnical systems. They were encouraged to identify stakeholders and key advantages, disadvantages, impacts and tradeoffs inherent in each solution or pathway, including barriers to implementation and mechanisms for overcoming these barriers.

Workshop Activities

The workshop convened on February 10 with welcome remarks from Farnam Jahanian, Assistant Director for Computer and Information Science and Engineering at the National Science Foundation, and Tom Kalil, Deputy Director for Technology and Innovation at the White House Office of Science and Technology Policy. Craig Partridge, the Ideas Lab Director detailed the workshop goal and structure. Several activities were conducted to familiarize the participants with each other's backgrounds, interests and expertise.

In order to seed discussions and brainstorming, the Steering Committee identified and presented five major themes: *identity management*, *adversaries*, *transition to practice*, *clean-slate thinking*, and *defining goals for cybersecurity*. These themes were used as the focus of several small-group roundtable discussions. Additional themes were identified, discussed and developed in a series of breakout sessions.

During the workshop, participants generated a large list of ideas, and then chose a smaller set for development into full recommendations. These recommendations were developed in parallel, by small groups at the workshop, and later compiled into this report under the guidance of the workshop Steering Committee.

The Purpose of this Report

This report is intended for use in disseminating the furthest-developed ideas and solutions that emerged at the workshop. Readers should keep in mind that rather than producing a comprehensive set of recommendations, the workshop's goal was to generate actionable ideas. Consistent with the focus on

enabling the expression of a range of viewpoints and thinking, the workshop did not seek to produce a consensus list of recommendations; the goal was to capture ideas that were actionable, well-articulated and intellectually robust. To this point, an individual's participation in the workshop does not imply that he or she endorses all or any of the specific recommendations contained herein. All contributions to the workshop and to this report were made by the participants as individuals; report content should in no way be taken as representative of the views of any organization with which any participant is affiliated.

RECOMMENDATIONS

The recommendations are presented in three categories: *technology*, *policy* and *leadership*. *Technology* recommendations include ideas for research, development, deployment and modification of cyberattack prevention and response tools. *Policy* recommendations require action by the Federal Government in the form of new laws or regulations, standards, governance or coordination models. *Leadership* pathways invite the Federal Government to leverage its existing leadership in information technology to improve cybersecurity practices, both inside and outside the Government. A theme of education and awareness appears throughout these three categories, including ways of informing user practices and raising awareness of the impact of breaches.

The recommendations are diverse, spanning a range of challenges, mechanisms and themes. They address various timeframes with respect to a cyber incident, including preventative or preemptive measures, tools for system maintenance and incident response and post-incident evaluation and reporting. Some aim to convene experts and gather or advance knowledge of risks. Others propose the development of new standards, laws or regulations, new education and research initiatives or deployment of specific tools and technologies. Some recommendations are cautionary and invite us to better understand the consequences of the choices we may make. Most envision the Federal Government playing a key role, as a technology leader, a policy leader or an enabler of new initiatives.

Technology

The technology recommendations fall into two groups: those that look to technology to solve specific problems, and more far-reaching ideas that would be best supported by research programs. The focus in the first group is on improving Web security and mechanisms for updating software and systems when vulnerabilities need to be removed. The second group covers a broad set of topics in computer science and the social sciences.

Make Critical Subsystems Field-Updatable

Long-lived, often critical, systems such as embedded systems and industrial control systems have not typically been designed with updatable security capabilities. Mechanisms to update such systems, including current cyber-physical systems (CPS) and industrial control systems, are needed, rather than the current system where some industrial control systems are designed to be locked down once deployed. One of the reasons such systems are locked down is the tremendous impact if they fail to work as expected. Any update system must address this challenge.

Problem Statement

Systems must be updated to address the evolving nature of threats over time. Embedded, cyber-physical, and industrial control systems, which are often quite long lived, have not typically been designed with updatable or replaceable security capabilities in mind, due to possible negative impacts of system disruption. This results in limited and costly ability to adjust to emerging threats and evolving uses of these systems.

There is opportunity here. Industry is designing and fielding higher-level computing systems with disaggregated security functionality, such as security coprocessors, “trusted” CPU modes, and key management subsystems. These disaggregated capabilities often include enough general-purpose capability to allow for additional functionality to be added over time.

Recommendation

Promote research and development that is specifically focused on the ability to field-update or replace security mechanisms in systems that have traditionally been locked down once deployed. Specific issues that will need to be addressed include:

1. Risks to any secure state maintained by components that are being updated.
2. Risk in changing components of a distributed system that other components might rely on.
3. Difficulties when a device cannot be physically accessed for repair.
4. Risk when update mechanism must be resilient to physical attack (e.g. when applied to Trusted Platform Modules).

Next steps

Actual deployment will require collaboration between industry and academia. Such collaboration would allow a path to drive updatable security into existing and emerging platforms that have not previously been exposed to the benefits of such capabilities. In addition, the practical issues arising in such a collaboration would also ease the formalization of disaggregation of security in higher-level computing systems.

Incentives will be necessary to encourage both communities to work together in this space, and to overcome a host of potential roadblocks, including historical and cultural biases as well as intellectual property issues. In the past, incentivizing such collaboration has succeeded in specific areas. For example, the Semiconductor Research Corporation (SRC) has successfully represented industry interests to academia through close collaboration with NSF and other Government funding agencies. They have served as conduit for bringing industry researchers together with academics for specific projects with both scientific and immediate practical benefit, and they have provided acceptable intellectual property terms for both communities.

Enable Certificate Transparency and Security

Certificates are used to associate the entity that registered and owns a domain with a public key that can be used to protect communications with that domain. Certificates are the primary mechanism for authenticating that internet connections are actually made with the entities they claim to represent. Thousands of certificate authorities issue and manage these certificates. These authorities can intentionally or unintentionally issue incorrect or false certificates, enabling eavesdropping on communications or even impersonation of the domain owner. Transparency and trustworthiness should be increased, either by endorsing efforts such as those by certificate-transparency.org or through other mechanisms.

Problem Statement

Website certificates are the primary mechanism for authenticating internet connections; they underlie key services and protocols such as secure sockets layers (SSL), hypertext transfer protocol secure (HTTPS), substitution-permutation networks (SPNs), and many more. With the Internet mediating a rapidly increasing scope of critical transactions, connection authentication has become essential. Certificates are issued and managed by ostensibly trustworthy third parties, known as certificate authorities (CAs). In the case of HTTPS, browser vendors include a default set of CAs that the browser's developers believe to be trustworthy; these CAs can issue certificates for any website. For example, Firefox includes over a thousand default CAs, including governments (e.g. China and Russia). This creates significant vulnerabilities, because it is possible for a CA to intentionally issue incorrect certificates (thereby enabling malicious behavior), to be tricked into issuing incorrect certificates, or to be penetrated by an adversary who may issue false certificates on their behalf.

In principle, any CA can issue a certificate for any website. Certificate authorities are often organized in a hierarchy with one CA at the root; such a structure establishes a chain of trust.

Recommendations

Given the importance of certificates, and the difficulty users have in knowing which CAs to trust, it is important to increase the transparency and trustworthiness of the certificate process. Independent of Government action, there is already considerable momentum in industry towards this goal (e.g., certificate-transparency.org), but more can and should be done.

Next steps

One proposed approach is as follows.

1. When a CA issues a certificate, it also sends a copy to a public log, which uses cryptographic means to preserve the integrity of the copy.
2. Independent auditors can review the log's contents and verify or identify fraudulent certificates.
3. Browsers, as well as other users, can independently verify that a certificate appears in the log. In addition, they can check with an auditor to verify any given certificate.

This approach can be generalized to use multiple logs and multiple auditors.

Other transparency measures can be considered. One useful approach would be for browsers to report when an unexpected CA issues a certificate for a given website, such as a CA in Burundi signing a certificate for google.com. Other reputation mechanisms, such as a rating system for CAs and more aggressive vetting of CAs, could be implemented. One current step in this direction is the use of cert-pinning, a technique where a browser checks a certificate against some vetting function, for example comparing it to the first received certificate or using a secure hash of the certificate embedded in a Web application. Finally, there could be efforts to develop another authentication mechanism that does not require trusted third parties like CAs.

Other nontechnical approaches include efforts to educate users to prefer browsers that are more judicious in their acceptance of CAs.

For such efforts to succeed, industry and government would need to be aware of any emerging mechanisms to increase certificate transparency and adjust procedures accordingly. Transparency may disrupt some workarounds that are relied upon today. This would also change how CAs are used, which could put pressure on CAs. The logging approach outlined above still has an issue of trust associated with CAs that might be hard to avoid – it could be misused by, say, state actors. Finally, implementing this new certificate validating infrastructure would require additional resources, organizations, and skill development. The cost of doing this may be too high for some countries with smaller economies.

Create a Framework for Managing Software Updates

The Internet of Things will challenge our current channels for distributing security updates. An environment needs to be developed for distributing security patches that scales to a world where almost everything is connected to the Internet and many "things" are largely unattended.

Problem Statement

The Internet of Things (IoT) will challenge our current channels for distributing security updates. Years of experience in the development of desktop and server software have shown that even when developed with the best-known development processes and by highly skilled programmers, there are inherent vulnerabilities in code of any reasonable size. Thus an essential component of deploying software is a channel for distributing security updates and patches.

While security updates in the Cloud are often more easily managed due to both management and configuration issues, the highly distributed and heterogeneous environment of embedded devices presents significant challenges for patching vulnerabilities. Those barriers include a diverse developer base that is often less well-resourced and knowledgeable in security; varied risks posed by devices and deployment contexts (for example thermostat v. car); under-appreciation of the cumulative risk posed by the presence of an increasing number of inconsistently managed networked devices. Given the diversity of contexts in which these devices are used, they present new risks, including risks to human life and health, potential to invade privacy, and risks to third parties, where embedded devices can be used as launch pads for attacks.

Because many devices will be unattended, extra attention needs to be paid to the problem of spoofed updates, where the software update system is subverted to get devices to install malware. Given the devices will not be supervised, uninstalling or repairing the malware will be even more difficult than it is today.

Recommendation

Build a security patch management system for the emerging IoT environment that is commensurate with the risks it brings. The security download mechanism includes a framework that is both trustworthy and instills trust in users.

Next steps

There are myriad models for software update management, supporting off-the-shelf software, enterprise software, cloud services, and embedded systems. It is recommended that the White House facilitate a learning forum that brings together established software industries that have developed infrastructures and processes for security patch management to share information and experiences with stakeholders of the IoT. This could create an appetite for and understanding of the value of a streamlined approach to security updates.

Subsequently, we recommend that the Federal Trade Commission convene a workshop as part of their ongoing exploration of the IoT that focuses on the consumer protection framework for patch management processes for addressing security vulnerabilities. Such a workshop can explore the ways in which variations in risk profiles may influence the mandatory or voluntary nature of patching; the appropriate level of user involvement; patch authenticity; and the mechanics of distribution and timing of the downloads.

This recommendation would be relatively easy to launch, since it primarily requires bringing stakeholders, technology experts, and policy experts together. It should be done quickly because IoT is already expected to have a large role in the not-too-future Internet. Acting on the recommendations of these groups could be more difficult, since they would involve both technical advances, changes to

business models influenced by law and policy, and consumer education.

Make HTTPS the Least-Effort Scheme for Deploying a Website

The secure version of the Web's primary protocol should be made the usual way that servers interact with clients.

Problem Statement

Today, the easiest way to deploy a website is to forgo the encryption and site-authentication available via hypertext transfer protocol secure (HTTPS, the secure version of the Web's Hypertext Transfer Protocol). The result is that tremendous amounts of data, including personal information, are sent unencrypted over the Internet.

To add HTTPS, a website creator needs to:

1. Buy a certificate to prove that (s)he owns the domain name for his/her website, which often costs more than the domain name itself.
2. Make extensive configuration changes to the website.
3. Forgo use of third-party plugins that do not support HTTPS.
4. Debug many failures that often result from turning HTTPS on for software that does not anticipate its use.

In some cases, these challenges can add many days to the deployment of even simple websites. This is unfortunate. While not all websites require data being exchanged to be encrypted, website creators may decide against encrypting transmitted data for poor reasons.

Recommendation

Using HTTPS should be made the least-effort path to deploying a website.

Next Steps

First, the .com registry and its affiliated registrars, along with other US-influenced domains, could issue a certificate, signed via a DNS-Based Authentication of Named Entities (DANE) with every domain registration. It would be best if the Internet Corporation for Assigned Names and Numbers (ICANN) would do this on its own initiative, but if not, the Federal Government may need to require it, at least for US-based sites, perhaps via legislation.

Next, registrars could adopt a well-defined automated process for obtaining certificates that allow automated Web server configuration tools to obtain a certificate. The Federal Government can only require this of US-based registrars; note that represents a substantial fraction of registrars. ICANN could develop requirements for the wider world.

Fund efforts to update popular open-source software to automate the steps required to add HTTPS support to websites.

Fund academic research into the usability of the tools needed to deploy HTTPS into Web services and

maintain security (key management, key cycling if servers become compromised, and so on).

Cybersecurity Research Agenda

Key research thrusts include: certifying the security properties of future systems, security of software-defined networks, R&D on clean slate operating systems and system services, security and privacy on the cheap, and consequences of aligning network structure with political jurisdiction.

Problem Statement

Cybersecurity is an ongoing, critical and active area of research and development. As new technologies and practices are adopted and deployed throughout our networks, the Nation's research agenda must expand to address and anticipate the changing risk landscape, yielding new near-term solutions and long-term capabilities.

Recommendation

It is recommended that the following topics, which are not comprehensive but will significantly enhance the horizons of the cybersecurity-relevant R&D, be incorporated into our national research agenda.

A. Certifying the security properties of future systems

This is a long-term research agenda that is both high risk (it sets a high bar) but also has high rewards if successful.

Adversaries routinely exploit design flaws and implementation bugs in today's systems (e.g., cryptographic primitives, network protocols, and software systems) to compromise integrity, confidentiality, and availability. More widespread use of tools and techniques for proving these systems correct and secure (under a particular threat model and deployment environment) could lead to better systems security, and reduce the reliance on case-based testing to uncover bugs. Standards bodies and relevant government agencies could have a certification process or competitions for provably secure artifacts. The National Institute of Standards and Technology (NIST) could play a much larger role in validating designs and implementations before they are used in the field, particularly for critical infrastructure. Research in this area could explore a number of topics including (i) the policy issues of having a government agency play this role, (ii) translation of existing verification tools and proof techniques into practice (including education and training), and (iii) foundational research on the scalability and usability of tools and techniques for validating ever more sophisticated artifacts.

This research agenda is consistent with programs funded by Defense Advanced Research Projects Agency (DARPA), other Department of Defense (DoD) agencies, and NSF.

B. Security of Software-Defined Networks

Because of the rapid development and deployment of network virtualization, this is a high priority research direction.

Bugs in router and switch software introduce vulnerabilities that adversaries can exploit to control computer networks. The emergence of Software-Defined Networking (SDN) represents a unique opportunity to get security "right" from the beginning, and refactor or even eliminate the "dusty deck" of today's router and switch software. Standard interfaces (e.g., OpenFlow) between an SDN controller and the underlying switches can enable general, reusable techniques for verification, testing, and software synthesis. However, SDN may also introduce new security risks, such as the reliance on a logically centralized controller and the much wider range of (possibly untrusted) third-party software developers. Industry is moving quickly in the SDN space, making it all the more important to address the security challenges before major design decisions are made, and software artifacts built. Government funding agencies could encourage more research on SDN security, including (i) an analysis of the security threats, (ii) the design of provably correct controller platforms, (iii) techniques for protecting an SDN from rogue controller applications, and (iv) programming languages raise the level of abstraction for writing controller applications. Existing SDN deployments focus on a single administrative domain, but supporting multi-domain deployment experiments (e.g., across multiple research or government networks) would enable the research community to start investigating the security and trust issues that will arise in future commercial deployments.

This research agenda is consistent with programs funded by DARPA, other DoD agencies, and NSF. There is a potential for joint programs with industry.

C. Evaluating the Deployability of New Security Technologies

This is highly multidisciplinary research topic that, if successful, could have a large impact on new and effective deployed security solutions.

The vast majority of security solutions, such as new protocols or primitives, are never deployed in practice. Even the solutions that are ultimately deployed take years of iteration to arrive at a design that is amenable to deployment in practice. Examples include IPv6, DNS-SEC, and BGP security that have taken years to standardize and (partially) deploy. The community still has little ability to assess what kinds of protocol enhancements can be readily deployed in practice, and how to best encourage deployment. Achieving deployment traction requires a deep understanding of deployment costs, the incentives for adoption, strategies for incremental deployment, and how to evaluate the security benefits of a partial deployment. Creating rigorous techniques for analyzing deployability, and using them to evaluate existing proposals, could significantly improve the state of the art. This research area is inherently interdisciplinary, and could draw on policy analysis (to understand who all of the players are in the ecosystem), game theory (to analyze and structure the incentives for

deployment), information-technology policy (to understand the role the government could play in encouraging or mandating deployment), and security (to identify new threats these protocols may introduce). This research area would also benefit from a stronger relationship between researchers and industry, including equipment vendors (to understand development costs and performance challenges) and network operators (to understand deployment and management challenges).

This research agenda is consistent with programs funded by DARPA, other DoD agencies, and NSF. There is a potential for joint programs with industry.

D. Research and Development for a Clean Slate Operating System

This research agenda is high impact, and focused with the well-defined goal of a new, clean slate operating system (OS) and related system services.

Current commercial operating systems have two significant failings: they are inherently insecure, constituting single points of failure, and they serve both application and human users poorly. Three decades of operating system research has, until recently, mostly focused on specific technologies, like hypervisors, minimal kernels and sandboxing that could be incorporated into any OS. These point technology solutions do very little to improve the users' experience.

Wide-scale adoption of a clean slate redesign poses serious challenges, but past efforts by government to produce entire operating systems have, in fact, led to existing well-established systems, such as Unix systems that underlie both the Berkeley Software Distribution and the Mac OS. Rapid changes in mobile device industry and technology also point to a potential for adoption of new, government funded system software that combines the goals of increased security and improved user experience.

A specific goal of this research agenda is the identification of specific requirements on the new OS, such as security partitioning that prevents an adversary's ability to compromise an entire OS, as well as performance and usability features required to improve the user experience. Identification of research gaps on usability, and support of parallel, additional research in these areas would also be useful.

This research agenda is consistent with the Clean-slate design of Resilient Adaptive Secure Hosts (CRASH) program funded by DARPA as well as projects funded under the cybersecurity program of NSF, and there are a few ongoing attempts at such clean-slate redesign. While this is a good start, none have yet reached the level warranting wide scale adoption. A suitable clean slate operating system is still three to six years away, assuming sustained funding. This clean-slate OS would be available for adoption by then current device and OS providers and microkernel and hypervisor technologies incorporated in the new OS could ensure that backward application compatibility could be supported with increased, if not perfect, security.

Work to affect transition can begin in parallel with development of the new system.

Niche markets could be identified that combine the need for high security with high government support and regulation, and a combination of regulatory push and economic pull could support the inherent attraction of combined high security and enhanced usability in the new OS. Current research suggests that even a ten percent penetration of current markets by a significantly more secure OS could significantly enhance overall security of the computing population, by providing anchor points for trust management and recovery. Additional research on this topic (computing ecology in an adversarial environment) should also be considered.

E. Security (and Privacy) on the Cheap

This research agenda has moderately high risk and high reward.

Many industries (for example, automotive and those arising in the growing area of "Internet of Things") need methods of security development that can be used for smaller and lower resourced products that are now being connected to the public Internet. Such products have relatively low budgets and smaller staffing with limited expertise. Unfortunately, such products, once deployed, will provide new attack surfaces of perhaps very large scale. This research program focuses on the development of low cost (both in terms of tools and in terms of training) methods for developing secure systems. Example outcomes would be design patterns, software engineering tools, middleware, and automated code generation that could be easily used by engineers.

This is a complex research agenda, spanning computer science, computer engineering, sociotechnical systems, cultural anthropology, and economics. DARPA and NSF would best support it, although later developments might be supported by NIST as well as by the Department of Transportation (DoT), the Advanced Research Projects Agency-Energy (ARPA-E), the National Institute of Food and Agriculture and other specific mission agencies.

F. Consequences of Aligning Network Structure with Political Jurisdiction

The Internet is becoming Balkanized as several countries have found it part of their national agenda to exert stronger control over their network's use. This trend has received an additional push with the recent revelations that have shown the dangers of the implicit control the US has had in the Internet governance. While there are many who wish to fight against such Balkanization, the consequences of the process continuing should be studied now. Doing so can both allow all to benefit from this trend, and to provide informed back pressure - if it is indeed possible to do so.

The consequences appear to be complex. On the positive side, Balkanization allows better accountability for enforcement. This is analogous to the power countries obtain by having their own currency - it allows them to exert better control over their own economies, and to make tradeoffs in internal and external markets to counteract the

particular circumstances of the region. Of course, it also increases the power a country has to oppress its citizens - again, the analogy of local currency applies here as well, although with a different kind of impact. Thus, the study should consider impact to foreign policy. In addition, there is the potential of a significant economic impact, since some regions may wish to use particular infrastructure. The impact is not easy to predict - local industries may benefit, or new industries may develop to accommodate parts of the infrastructure that become common across a large number of jurisdictions. At the same time, internet companies that have benefitted from "one" Internet (e.g., Google, Facebook) could find such a Balkanization disruptive and costly.

Clearly this research agenda is multidisciplinary. There are also prominent stakeholders: simply within the US Government, they include the Department of State, Department of Commerce, DoD, the Federal Communications Commission, and the Federal Trade Commission. The results would have a broad audience outside of the United States. There are also a number of intellectual communities and technology industries that have staked out strong positions on this topic.

It will be important to frame this research work carefully and to properly recognize the range of skills required to conduct the research and to generate results that are informative and credible. One possibility is that the National Research Council should be asked to convene a study to recommend how best to make progress in this area.

G. Intent-Based Forwarding: Using Identity and Service to Route Packets

Networking today can be described as universal reachability with exceptions. Transparent bridging with spanning tree and IP routing both provide a reachability layer that connects all network participants along the shortest path. In order for a network user or network owner to impose policy, whether that is reachability constraints for security purposes or path selection for performance (or cost), they must configure layer-violating features on a hop-by-hop basis. These features then are operating at odds with the underlying universal reachability layer. Any errors in the network configuration can lead to unintended reachability.

Intent based forwarding removes the reliance on L2/L3 connectivity. Instead, each network request is examined for its intent. This includes the identity of the requester, service being requested, location and posture of the device making the request, time of day, and so on. The user intent is mapped against the network policy, whereupon the security, topology, and other services required to satisfy the user intent are found (or the request is denied if the intent doesn't match any policy).

The key to usability will be to express the user identification and intent in terms of a hierarchical class structure. The policy can then be expressed at different levels of the class hierarchy. This will allow for example, the network operator to easily map the network policy into their desired application experience, easing the task of writing and updating the network policies.

Having a clear expression of policy and intent increases the odds that the implemented policy will match the administrative policy. By having the network fail-safe, not allowing communication outside of the allowed policy, the security of the network is improved. The usability of the network is similarly improved for all other policies.

Possible implementation paths for intent based forwarding have been identified. There are also a number of ancillary issues requiring further investigation.

Finally, there is a possibility to combine this work with Named Data Networking (NDN) in the future. This could allow a much more precise determination of the user intent to be used, and could also resolve some of the complicated issues in NDN to do with service location and mapping to L3 addresses.

This research agenda could be supported by DARPA and NSF, and has the potential of industrial support.

Policy

The policy recommendations cover a range of topics, but all require Federal Government action in some form. They also address two key cybersecurity challenges identified in workshop discussions:

1. Reliable data about cybersecurity incidents is lacking. Due to a range of sensitivities about revealing vulnerabilities, concerns about liabilities, and the like, information about incidents (and vulnerabilities) is incompletely shared and often, poorly analyzed. Our ability to respond to incidents and learn from them is thereby compromised.
2. Multiple organizations have been tasked to coordinate on cybersecurity. The result is that some needs are well covered, some needs are barely covered, and some needs are partially met. Further, the degree to which government, business, or the public gets the cybersecurity assistance they need varies widely.

Establish an Internet Rescue Squad

Implement a national cybersecurity response team to coordinate responses to breaches and attacks, prompt secondary targets to initiate response programs, and support response activities.

Problem Statement

Many national-level agencies are currently chartered to provide clearinghouse and industry-wide advisory services. Others, often in the law enforcement arena, take information that could be of use to target entities and sequester that information with the goal of achieving successful prosecution. These drivers result in a world in which victims struggle to find help in responding to an attack (for which they are often ill-prepared) and there is no one charged to help secondary target and victim enterprises take the remediation steps to avoid becoming attack victims themselves.

Recommendation

Implement (or add to the charter of an existing entity) a team of incident responders who, when provided information about an ongoing or recent problem, will analyze the incident; identify entities who are or might be directly impacted; establish outreach to those entities with clear information and actions; and collect secondary information that the entities are willing to share back out. This will be provided as a public service to impacted entities.

Next Steps

Existing entities can be leveraged in the establishment of this response team. The National Cyber-Forensics and Training Alliance and the Advanced Cyber Security Center (a Massachusetts-based industry sharing alliance) may provide useful models of small-scale, distributed versions; the National Cybersecurity and Communications Integration Center, United States Computer Emergency Readiness Team, and the Financial Services Information Sharing and Analysis Center are models of clearinghouse and analysis centers that can be of use. Security vendors often have cross-customer response teams that

can be used as force multipliers, as well as tactical incident response services that can provide deeper support than those with a national focus).

Potential complications:

- Liability could become an issue if the response team makes incorrect assessments. This could be ameliorated through transparency of operations.
- Targeted enterprises may be unwilling to share information for fear of disclosure.
- Placing the team in an agency that will support a strong level of technical and operational excellence with an emphasis on rapid response may prove challenging. As this is not enforcement, but emergency response, existing technical teams in the FBI may be inappropriate to scale to this function.
- Maintaining and implementing a coherent CRM system to manage points of contact across industries will be a necessity. This may also be synergistic with law enforcement needs in this area.
- Will information in the possession of this team be subject to FOIA? Can it be protected to incentivize companies to share information?
- What should the steady state size be, and how will its initial instantiation affect its culture and growth opportunities?
- How will the response team prioritize its limited resources, and ensure maximal benefit?

Create a Cyber NTSB

Create a cyber analogue to the National Transportation Safety Board charged with analyzing cybersecurity incidents and providing public reports on the circumstances and causes of each incident.

Problem Statement

A critical problem in cyber security is a lack of reliable, consistently reported, data about security incidents. The lack of data makes it difficult for others to learn from these attacks, and is leading to misplaced priorities.

Recommendation

The government should create an organization charged with investigating cybersecurity incidents. Such investigations would not be for law enforcement or immediate response purposes, but rather to carefully analyze each incident and publically report the who, what, where, when, how and [perhaps] why behind an incident. One can think of this service as analogous to what the National Transportation Safety Board (NTSB) does for the transportation industry, or as an improvement upon the Computer Emergency Response Team (CERT) or Computer Security Incident Response Team (CSIRT) approach.

Done right, such an organization could make tremendous contributions, by providing a common base of information about what types of incidents occur, who is affected, who is attacking, the methods of attacks, and the vulnerabilities that are exploited, both at a given point in time and as a way of identifying and characterizing trends. The reports could be mined to guide research and policy, much as NTSB reports lead to improvements in transportation safety.

The cyber NTSB would also make careful, data-based, policy recommendations in response to incidents. (As an example, there's reason to believe the wrong lessons are being drawn from the recent Target incident: the details of the attack are not terribly interesting, but the need to give the same protections to debit card holders as credit card holders in cases of identity theft is a clear policy lesson.)

This data-driven role, both in driving improvements and driving policy recommendations, clearly differentiates the cyber NTSB from any existing organization.

Furthermore the cyber NTSB could be the center of expertise that other agencies impacted by cyber security issues, agencies covering bank safety, transportation safety, and programmable medical devices, call upon to assist in their investigations. The cyber NTSB would also cooperate with law enforcement and national security organizations.

Next Steps

This recommendation could potentially span industry, government and academia in a variety of overlapping roles.

The board could be created by Executive Action or by the Congress. Federal legislation may be needed to provide the incentives necessary for the board to succeed. Both protecting participants and requiring participation would appear to require legislation.

Pitfalls

Implemented incorrectly, a cyber NTSB could create for more problems than it solves. Concerns include:

- *How will the right expertise for each investigation be assembled?* Having the wrong experts will lead to bad or possibly even harmful, analysis. Expertise along all dimensions (domain, system, security, and so on) varies considerably.
- *How will a cyber NTSB fit in with existing industry practices?* Industry has some methods for privately sharing information about cyber security incidents. Can this process be made to work cooperatively? What can we learn from the history of Information Sharing and Analysis Centers (ISACS; see http://www.isaccouncil.org/home.html.)? How do we interact with existing CERTS?
- *Who would have access to results?* Some incidents will be sensitive (in that documenting an exploit may make others vulnerable), yet the value of the reports is rooted in their being public. How will these two concerns be balanced? Are there different degrees of being public, or different kinds of public that could, for example, lead to different degrees of discussion?
- *How do should terminology and practices be standardized?* The security industry lacks a clear taxonomy of attacks and words to describe incidents, and there is concern a new board could be overwhelmed by simply trying to develop a common set of terms (and there are consequences if the board gets it wrong). Converging on standard terminology is a function of the National Institute of Standards and Technology (NIST), which has a history of addressing Federal cybersecurity and also of bridging to industry.

Creating a cyber NTSB would require considerable care if the cyber NTSB is to be a success. Here are some possible approaches:

- *Provide incentives to participate*. Participants in NTSB investigations benefit from the fact that it

is separate from the key transportation regulatory agencies and from law enforcement. Similar incentives may be required for a cyber NTSB. One might, for instance, limit the liability of corporations if they participate fully in investigating their data breaches. Some level of mandate/legal authority may be required, such as requiring cyber NTSB investigations for incidents of a certain size or larger. This would overcome some of the practical limitations experienced with organizations such as CERT/CSIRTs.

- *Don't overwhelm the board with investigations*. There are currently far more incidents than a new investigatory organization can fully investigate. Set some thresholds for incidents that **must** be investigated (e.g. > 10,000 users affected or > $2M in damage) and make investigation of other incidents discretionary.
- *Divide and conquer?* Consider whether there are meaningful differences among entities that could/should translate into substructures.
- *Learn from past failures and ambiguities*. The Securities and Exchange Commission requires certain kinds of disclosures, those that relate to material risk to investors. Their 2011 guidance has led to some limited reporting of cyber incidents, constrained by a compliance orientation in the reporting and the voluntary nature of "guidance." There is a possibility of a requirement to report, and whether that will lead to more meaningful reporting is open to debate.

"Standard" Impact Statement

Develop models for characterizing the cost and impacts, both positive and negative, of cybersecurity frameworks and standards. Before new standards are implemented, their impact should be assessed in a process analogous to generation of an Environmental Impact Statement.

Problem Statement

New standards offer the potential for benefits but can add costs --- some understood and others not. Furthermore, many common security practices are known to have serious limitations such that they should not be adopted without a clear analysis of their utility. The decision to deploy a standard can have widespread consequences, and all stakeholders benefit when these consequences can be accurately assessed in order to understand whether, how, and where to deploy a new standard or framework for cybersecurity. Creating a basis for making such assessments is the intent of this recommendation.

Recommendation

1. The Federal Government should develop expertise and models for characterizing the cost and impact of cyber security frameworks and standards, including incremental and partial deployments. [Proposed time frame 12-18 months]
2. For each new standard written or incorporated by reference, the cognizant agency must
    - Rely on wide spread direct experience with what's being recommended prior to mandating the standard
    - Expect to develop an impact statement based on the experience
    - The comment period should explicitly solicit input on the impact statement
    - [Timeframe: ongoing]

Next Steps

There are two next steps, one near-term and the other longer-term.

The near-term step is to develop the expertise and models for characterizing the cost and impact of cybersecurity frameworks and standards. This step overlaps the recommended next steps for the "Golden Goose" idea above. The likely impacts that imposing standards will have on the stakeholders – public or private – must be understood.

The second step, which may be achievable with an Executive Order, is to require agencies to use the frameworks to evaluate the impact of their proposed standards. This step may encounter some resistance. Imposing additional obligations on those who are formulating new cybersecurity frameworks and standards does mean additional time and effort must be devoted, delaying any ultimate benefits from deployment. Some of these risks are seen with "environmental impact statements" that are already required in connection with large-scale public works projects.

Finding an appropriate balance between understanding impact and not creating excessive time, effort or delay is an important initial step. Working with sector-specific agencies that have regulatory oversight and/or rule-making authority will be an important first step.

How Golden is Our Goose?

The computing and communications community has been a golden goose that has repeatedly produced economic and social gains. Security requirements are often seen to stifle innovation or usability of technology. Tradeoffs need to be better understood and used to assess whether the economic or social harm of securing or protecting from breaches may exceed its benefit.

Problem Statement

The title of this recommendation comes from the story of the goose that laid golden eggs, and its moral of the cost of poorly-informed choices. It is frequently asserted that "Security is bad and getting worse" or "We are losing the security battle." Often, these assertions are supported by anecdotal stories, without specific measurements of the costs incurred by given flaws, or the benefits provided by the systems vulnerable to those flaws. Similarly, it is widely believed, again largely with anecdotal examples, that imposing security requirements hinders technological innovation. An informed discussion of necessary cybersecurity improvement steps would require an understanding of the benefits that the Internet provides to the United States, the actual reduced rate of innovation due to security requirements and the amount of reduction that society would consider unacceptable, the potential for losses due to lack of security requirements, the controls that might provide better resilience against such hazards, and the costs that those controls impose on the benefits society receives.

Recommendation

An evaluation of the benefits of the Internet, in GDP, American quality of life, and other measures should be conducted. The costs of increased security requirements should then be evaluated, as well as the potential losses due to security breaches that could result if required security measures are not taken

(including loss of life, loss of time, and impact on economic activity). Risks and proposed mitigations should be evaluated in the context of the losses that society deems unacceptable.

Next Steps

The obvious next step is to fund multiple studies to evaluate the benefits of the Internet and to study the potential costs of imposing certain security requirements. This work could be funded through a single Government agency (which may improve focus) or multiple cooperating agencies (to increase in-Government awareness of tradeoffs and reflect per-agency needs). Regardless of the funding approach, the performers doing the research should be interacting with each other, through group program meetings or a similar mechanism, as such meetings should lead to rapid refinements of the frameworks and models.

Once the frameworks and models (the fruits of these studies) have been established, they can be used to help prioritize which security requirements to encourage and discourage. Research should also be encouraged in areas where it is believed that security is required but the known approaches have too high a cost. We may be able to take as inputs the results from some other ideas in this report (such as the list of high impact opportunities described in the Leadership section below).

Identity: A Problem That Doesn't Need Solving

Personally identifiable packet-level attribution will not address the most pressing cybersecurity issue: cross-jurisdictional cyberexploitation. No new efforts should be undertaken in this area. On the other hand, the National Institute of Standards and Technology's efforts through the National Strategies for Trusted Identities in Cyberspace to improve authentication technologies are useful and should be funded.

Background

Depending on circumstances, users on the Internet can be relatively anonymous. As a result it is often suggested that attribution and better identity management would effectively deter exploits and attacks. Close examination shows, however, that attackers hide through "multi-stage" attacks—attacks where an attacker infiltrates one computer as a platform to attack a second, infiltrates the second to attack the third, etc. As long as one can examine the ISP or machine from which the attack has been launched, it is relatively easy to trace the attacker. Similarly, a multihop attack can sometimes be traced across hops within a particular jurisdiction. That ability ceases once an attack crosses jurisdictional boundaries in cases where jurisdictions do not cooperate with the US—exactly where the problem is most pressing.

Redesigning the Internet to enable tracking of personally identifiable packets would not solve multi-stage, multi-jurisdictional attacks, but would disrupt many important internet values, including privacy, freedom of expression, and freedom of action (the latter two are important to US foreign policy and national security). It should be noted that network-level addresses are, on their own, often quite valuable for enabling investigations, at least initially.

In contrast to packet-level attribution, application-level identity management can be quite useful in securing many types of transactions, including in many of the following situations:

- A person or device needs to be authenticated to a resource for a fixed, limited period of time.

- A user serving multiple—and perhaps changing—roles within an enterprise needs an identity that can enforce separation of duties (as well as reflect changing roles).
- An anonymous or pseudonymous user, being tracked— perhaps across multiple devices—for customer-management purposes (including being served with personally tailored advertising).
- A single user with multiple identities—as an employee (say of a medium-security government facility), a member of a private-sector technical working group, a member of an online social network, an officer of a parent-teacher organization, and a member of the local dog club—with varying levels of authentication required for access.
- (One use case that appears less amenable to current identity-management solutions is the far less critical one of a user seeking access to resources over the "open" Internet, that is, to resources for which there are no trust relationships such as business contracts already in place.)

Recommendation

There are many reasons not to embark on efforts for personally identifiable packet-level attribution, and no justification for doing so; a solution will not help with the most pressing issue, cross-jurisdictional cyberexploitation. It is recommended that **no** work be done in this area. In contrast, improving technologies for authentication, including for device-to-device authentication, would be useful, and the National Institute for Standards and Technology's efforts in this direction should be supported and perhaps expanded.

Encourage the Adoption of Routing Security

A serious obstacle to better securing the Internet's routing and naming systems is the demonstrated desire of governments to leverage the routing and naming system to "take down" global internet sites for violations of national laws. Rather than being a prime offender in this regard, the U.S. Government should curb this behavior.

Problem Statement

The Internet is a loosely federated network of networks that is held together by the global name and number space. Organizations (e.g., Google, Telefonica, Verisign, NTT, Bank of India, University of Cambridge, etc.) run their own autonomous networks that are then interconnected using the global routing system. The routing system is based on a set of machine-interpretable addresses, called IP addresses (e.g. 192.168.1.1). Each organization is allocated one or more blocks of IP addresses, which are then used to transact with that organization on the global Internet. Global communication is further enabled by the domain name system (DNS) that provides a mapping from familiar human-readable addresses used on (e.g. www.google.com) to these machine-interpretable IP addresses (e.g. 192.168.1.1).

The routing and domain name systems enable global communication. Regardless of where one connects to the Internet, its use is based on globally-resolvable stable identifiers for online entities, which allow users to locate the entity they wish to communicate with, and to have a predictable communications experience. As such, the integrity and availability of the information provided by the domain name system, and the Internet routing system is crucial to preserving a global Internet that is not bound or otherwise constrained by national or other geopolitical borders.

Despite the crucial importance of the routing and domain name systems, they remain highly vulnerable to simple attacks. Today, an attacker located anywhere on the global Internet can manipulate routing information in order to intercept network traffic destined to any IP-address block allocated to any organization on the Internet. Such attacks occur with surprising regularity. The integrity of DNS information can be easily manipulated to redirect Internet traffic to maliciously controlled servers.

Fortunately, the global network operations community has converged on a number of solutions that significantly reduce the vulnerability of these systems to attack. DNS Security Extensions (DNSSEC**)** suite of specifications provides authentication and integrity for the domain name system, and the Resource Public Key Infrastructure (RPKI) provides authentication for the routing system. Deployment of these security systems began in 2005 and 2011 respectively, and they continue to be incrementally rolled out by network operators today.

Both DNSSEC and the RPKI are built upon hierarchical allocation structures; each authority in the hierarchy is allocated a set of domain names (or IP addresses blocks) which they may either (1) suballocate/delegate to other entities, or (2) authorize for use in the domain name system (or the routing system). The use of hierarchy enforces global uniqueness, preventing conflicts that could occur should the same domain name (or IP address) be allocated to multiple organizations. This global uniqueness helps preserve a global Internet, which is not bound or otherwise constrained by national or other geopolitical borders.

While these hierarchical systems and their corresponding security enhancements help create a secure and global Internet, they also introduce control points for domain names and IP addresses. More specifically, if the authority that authorizes use of a domain name (or IP address) is required to take down or redirect an address, the organization that holds the address will go offline for *all Internet users across the globe.*

This situation creates a conflict. On one hand, the security of global routing and domain name system is crucial to the continued success of the Internet. On the other hand, there is the opportunity/risk that governments will use these security mechanisms to enforce local laws (i.e., to remove objectionable content) and perhaps even enable offensive cyber capabilities (i.e. degrading an adversary's internet availability).

Such actions may be legal within a local jurisdiction, but the impact of such actions would has global repercussions, potentially affecting jurisdictions that are not subject to the laws being enforced. Coopting the security mechanisms afforded by DNSSEC and RPKI in this manner could force different entities to set up their own local systems as a means of avoiding the impact of takedowns that leverage DNSSEC and RPKI. This would in turn lead to the fragmentation of the Internet, which the U.S. seeks to avoid.

Thus, the security of the domain name and routing system is directly at odds with the goal of preserving a global Internet not bound to national or geopolitical boundaries.

The risk is real. The U.S. Government uses the global namespace to enforce local laws. For instance, in 2011, U.S. Immigration and Customs Enforcement (ICE) used the DNS to take down the seized domains Rojadirecta.com and Rojadirecta.org for alleged copyright violations, even though Spanish courts found that this Spanish company had not violated Spanish law. Although the specific DNS entities might be physically located within U.S. jurisdiction, the namespace they authorize is global and relied on by parties worldwide. Thus, the scope of the takedown can easily extend beyond the intended jurisdiction.

Recommendation

Given the unintended consequences of using DNSSEC and RPKI to enforce local laws, the U.S. Government should both refrain from using such mechanisms and publically endorse this principle in its Internet governance strategy. It should make clear that failure do so could lead to:

1. Lack of widespread deployment of security solutions such as DNSSEC and the RPKI.
2. Fragmentation of the global Internet, if nation-states migrate away from using the global domain name and routing systems and move towards locally-controlled namespaces that avoid the risk of being subject to the laws and policies of other jurisdictions
3. Multiple countries enforce conflicting local laws (e.g. censorship, information control, copyright issues) on these global systems, which can violate the integrity of these system and potentially lead to broader instability in the network and worldwide.

Next steps

This policy recommendation would need to be adopted by the White House as a priority and would require broad support both nationally and internationally. A first step in this direction is broad dissemination of the problem and risks. For example, non-governmental organizations (NGOs) and advocacy groups might be a vehicle, e.g. participating in high-level international ICT policy discussions under the United Nations Economic and Social Council. More locally, a clear and strong statement from a group of scientists, such as the Computer Professionals for Social Responsibility (CPSR) or the ACM could be used by Congressional staffers and by OSTP to forward the issue. This recommendation is timely because of the recent NTIA announcement of plans to transition certain aspects of DNS administration (the IANA functions) away from the U.S. Government, thus encouraging a more globalized DNS. Adopting this recommendation could provide important leverage for the U.S. to encourage other countries to adopt similar policies, thereby helping to preserve an open Internet. As a practical matter, takedowns are an effective tool for law enforcement, so an important next step is also to fund and conduct research into alternative technical methods for law enforcement activities in this space.

Enhance the Security of the Internet of Things by Identifying Enclaves

The security challenges posed by the emerging Internet of Things should be addressed now, to prepare before it is fully upon us. By identifying specific use segments, or "enclaves," of the Internet of Things infrastructure stakeholders can address the security requirements and devise event remediations for that enclave.

Problem Statement

The Internet of Things (IoT) is staged to become a dominant part of the Internet over the next decade. Companies have already made investments and business plans in this direction, such as Cisco's and Qualcomm's *Internet of Everything*, GE's *Industrial Internet*, and IBM's *Smart Planet*. The governments of many countries are mobilizing for the change; the U.S. Federal Trade Commission (FTC) is holding workshops on consumer concerns, South Korea is investing in the smart city Songdo, and partnerships are being created between countries as they look to create an early foothold in what is estimated to be a growing market over the next several years. Experience with earlier waves (computers, smartphones, etc.)

has shown that putting security last in the development process means that developers find themselves trying to patch up existing systems rather than doing things right from the beginning.

Recommendation

Face the issues involved in enhancing the security of Internet of Things deployments now. This includes not only those connected to or controlling critical infrastructure (smart cities, mHealth, etc.), but also less critical, but still connected infrastructure.

Next Steps

The issues can be fruitfully addressed by first identifying categories – enclaves – of devices that have similar functionality, deployment environments, threats, etc. Some example enclaves might include medical devices, automotive systems, voting machines, building controls, vertical farming sites, and so on.

For each enclave, one can perform a sector survey and gap analysis. Issues include:

- Any existing regulatory authorities, if any (the Food and Drug Administration, the Federal Highway Administration, FTC, etc.)
- The existing organizations that are setting standards (Institute for Electrical and Electronics Engineers, International Telecommunication Union, Industrial Internet Consortium, etc.)
- The existing security standards at all levels (process, software, hardware, management, etc.)

Then, an effort needs to be launched into developing and assessing the minimal security requirements for devices and software. This might be provided through the formation of an independent Underwriter's Laboratory (UL) like organization that independently assesses the security characteristics of products in the enclave. Any organization or approach should be set up so that the incentives of the assessment are aligned with the public benefit of security.

Finally, issues having to do with the establishment and support of industry-organized communities that support notification of security incidents, impact, and remediation within the security architecture. In some cases, there may already be such organizations; in many others, their creation will need to be fostered.

Of course, the interaction across enclaves is significant. Hence, in the long term, there will need to be cross-segment alignment of structures and practices.

This is a very expansive recommendation – it is essentially a high level roadmap that involves issues in technology, policy, education and awareness, economics, and sociotechnical organization. The barriers are large and numerous. For example,

- It is unclear who has the responsibility to address these issues from an over-arching perspective. Who, if anyone is/should be responsible for security across all sectors? There may very well be different responsibilities for different sectors.
- It is also unclear who should be responsible for performing the sector survey and gap analysis. One possibility would be for the National Institute of Standards and Technology to lead the development of reference architecture and technical roadmaps for a future IoT, similar to what it has done for Cloud Computing and for Big Data.

- There are significant policy issues that would need to be addressed. For example, Existing laws and regulations define requirements with respect to safety, but often do not directly address requirements with respect to security. Indeed, suppliers of software and of IoT devices are not currently required to disclose vulnerabilities or breaches, and may resist any requirements to do so.
- The maturity of various segments with respect to standardization of software used to produce IoT devices varies widely. In some industries there is significant standardization, in others there is almost none. In addition, some industries have strong regulatory oversight, while others have almost none. At the same time, there is a large category of consumer devices that do not have a well-established cohesive sector. For example, what industry organization would a company producing internet-connected thermostats or wearable computers belong to?

Leadership

The leadership ideas all involve leveraging the US Government's existing or historical role as a leader in information technology. As a leader, the US Government is in a position to set an example of how to properly secure systems. It is also in a position to assemble security expertise from government and industry and leverage that expertise. These ideas all involve leveraging that leadership.

These ideas also highlight a recurring theme of the workshop: we need to be careful in how we set priorities. Making intelligent security decisions requires us to weigh a thicket of policy and technology and social issues. One feature of several ideas here is finding ways for the Government to lead through that thicket.

Another recurring theme was the opportunity to educate the larger community. There are opportunities to educate by example (e.g. by having the Federal Government adopt best security practices), to educate by assembling expertise and publishing the result (e.g. assembling the top priorities) or by directly teaching (e.g. how-to guides for citizens and small businesses).

Create a List of Top Priorities

A well-considered list of top priorities in cybersecurity is lacking. The Federal Government should leverage its leadership position to gather inputs from major businesses and enterprises to jumpstart the creation of such a list.

Background

There is no shortage of cybersecurity threats, but a sound basis for prioritizing them is lacking, as is a basis for determining which would benefit from executive-branch support, efforts or action and which do not require it. Absent such a basis, government action risks wasting effort, alienating stakeholders, and losing good will due to these actions addressing the wrong problems

A wide range of stakeholders has a correspondingly wide range of perspectives on cybersecurity threats. These include US business and enterprises, Internet service providers, privacy advocates, researchers, individual users, and those in the security industry. Here we focus on the first group, as a way to illuminate an important subset of the problem space of very high relevance for the Nation's prosperity.

Recommendation

Concretely, we recommend engaging in a near-term time frame with a broad set of Corporate Security Officers (CSOs) or their equivalents from US businesses and enterprises, in order to solicit from them a prioritized list of specific cybersecurity concerns that in their view would benefit from government leadership, coordination, and/or assistance. We view CSOs as striking the right balance between technical depth and operational experience, as well as being cognizant of the constraints and opportunities that decision-makers face. These inputs would then be analyzed to distill out predominant themes and prioritize actions in view of the government's actual abilities and resources.

Next Steps

Request that public companies charter their CSOs to identify the Top Nine cybersecurity risks/concerns that affect their business, and communicate this to their Board of Directors. Then, identify the actions or issues where government involvement is applicable, and communicate those issues for inclusion in the prioritization process. This addresses co-option, ensures that the needs are tied to actual company concerns, and ensures that public company boards are engaged in the process of assessing risks and concerns.

This process faces several potential difficulties. A major one is the inputs or analysis becoming co-opted by parties interested in prioritizing a particular concern that is not in fact widely perceived as significant. In addition, CSOs may be reluctant to provide frank input due to (1) disclosing business-sensitive information about their operations, (2) informing attackers of the nature of the most significant vulnerabilities, or (3) a fear that they will provide government justification for unduly broad interference in their operations, which may be rooted in poor past experiences with government engagement. In light of this last issue, the effort must avoid mission creep or diffusion in which the process goes from prioritizing threats to an information sharing and data-collection exercise. Finally, the process will also inevitably disappoint some participants who do not find their concerns reflected in the final priorities / efforts, or who believe that the gravest threats are those not presently perceived.

Lead by Example

With one of the largest IT infrastructures in the country, the Federal Government needs to embrace its role as a leading technology organization and adopt best cybersecurity practices in its IT systems.

Background

The US Government IT infrastructure is one of the largest (and sometimes oldest) in the country. Government Agencies often use legacy and out-of-date systems that expose them to vulnerabilities and exploits. The PCAST report of November, 2013 "Immediate Opportunities for Strengthening the Nation's Cybersecurity" emphasized "Finding 1: The Federal Government rarely follows accepted best practices. It needs to lead by example and accelerate its efforts to make routine cyberattacks more difficult by implementing best practices for its own systems."

Recommendations

PCAST's top recommendation, which is endorsed here, was to urge the Government to comprehensively **update all operating systems** (and other software) to the latest, fully-security-patched versions. As they noted, that means no further use of Windows XP. But it also means more: across the diverse government-used closed and open source operating systems, browsers, programming environments, databases, and across mobile devices, desktops, and servers, nothing is ever more than one major version behind, with automated, timely installation of all security patches. This is not a one-time upgrade; this is a change from static installation to continual monitoring and update.

PCAST's second recommendation, that the Government employ Trusted Platform Modules (TPM), might better be phrased as "**use hardware-protected crypto** for especially sensitive credentials and to protect

the integrity of the system boot process." The details of how this is done will vary by system, since the technique is well established but not yet as broadly adopted as the first recommendation. The core idea is that malware on the system should not be able to exfiltrate important long-term keys or attain persistent access that survives a computer reboot.

Next Steps

1. The Office of Management and Budget (OMB) would require agencies to act on above recommendations in the next (Fiscal Year 2015) annual reporting memo to agencies.
2. The White House National Security Staff (NSS) would require agencies to report progress on meeting or achieving the above recommendations in Cross Agency Priorities reporting.
3. NSS and OMB would prioritize Agency funding, depreciation for needed equipment (software and hardware) and additional resources for government off-the-shelf systems replacement/upgrade when needed.
4. The Committee on National Security Systems (CNSS) and the Information Assurance Directorate (IAD) would issue instructions to the Department of Defense (DoD) and Intelligence Community (IC) to act and report on recommendations.
5. Agency Inspectors General would prioritize this activity for assessment in annual independent assessments and reports.
6. Agency Leadership and CIOs would issue instructions to system owners that priority is to protecting mission risks from cybersecurity threats, rather than interoperability risks in maintaining updated systems.

Where users and IT support organizations are distant, there may need to be explicit recognition from a higher level that guarantees of backwards compatibility are not achievable. Better that a few edge cases break, at planned update times, than continue to risk catastrophic compromise and outage of the entire system.

Government off-the-shelf systems (GOTSS) may seem to require XP, but it is possible that the software may actually work on current versions of the Windows operating system. If not, any systems that must use XP will have to be isolated from the rest of the network.

Re-Establish Trust in NIST's Cryptographic Standards Process

The US should reestablish the credibility of the National Institute of Standards and Technology as an honest broker of cryptographic and security standards. This could be done via a rigorous external review of NIST's cryptographic standards process and a public commitment from the Government that NIST's standards will not be subverted.

Background

With the development of the Advanced Encryption Standard (AES) in the 1990s, the National Institute of Standards and Technology (NIST) became a highly respected provider of strong, secure cryptographic standards for Federal non national-security systems. These standards were arrived at through a public, open process that was seen as careful and credible. AES were widely adopted by both domestic and foreign industry. This adoption improved security.

The revelation that NSA had corrupted NIST's standards process through convincing NIST to approve Dual EC-DRBG, a random number generator, as a Federal standard, has badly damaged trust in NIST. (The standard is widely believed to have an NSA backdoor; any system that uses it, including systems dependent on RSA's BSAFE toolkit, is not secure against the NSA.) NIST's credibility has badly suffered; there appears to be great unwillingness to use NIST standards.

There are various international efforts to find substitutes, yet currently there is no credible alternative to NIST. Various standards organizations, including IETF, W3C, etc., lack the depth to conduct the long-term, international standards competitions that NIST has organized to develop the Advanced Encryption Standard, the hash competition, etc. Should other nations such as Russia or China develop competing efforts, the end result is likely to be non-interoperable and less secure systems. This would be bad for US and international security as well as for US industry.

NIST has deprecated the subverted standard and embarked on an internal review to improve its process for approving cryptographic standards. It will also be having an external review process. Restoring credibility to a damaged institution is nonetheless is very difficult.

Recommendation

The US must reestablish the credibility of NIST as an honest broker of cryptographic and security standards.

Next Steps

Several things can be done. First, as NIST evaluates its cryptographic standardization processes, the agency must ensure that all aspects of the process are public. As part of this, it is critical that NIST's outside review be done by an independent and credible outside authority. Second, a public commitment by the White House—or NSA—that there will be no subversion of NIST standards and process, similar to the public commitment in PDD-28 regarding no economic espionage, would be a very strong statement that could go some ways to repairing the damage. Finally, increasing the number of technical experts within NIST's Computer Security Division (CSD) from the current dozen to approximately thirty (fifteen cryptographers and fifteen protocol analysts) would better enable NIST to carry out its role (note, though, that CSD will never have the strength or depth of NSA).

Develop Citizen and Small Business How-to Guides for Implementing Security

Current cybersecurity guidance is primarily targeted towards large corporations and the technically savvy user. A set of clear, interactive guides aimed at individual citizens and small and medium businesses should be produced to demonstrate best cybersecurity practices.

Problem Statement

Authoritative technical documents produced by the National Institute of Standards and Technology (NIST), i.e. SP-800 series, do not provide actionable guidance for accomplishing the simplest tasks for the average internet user or small business. The existing forest of YouTube videos and how-to blogs is

difficult to discover and navigate, and lacks authority. A comprehensive solution addressing the target audience is required.

Recommendation

An authoritative, easy-to use set of how-to guides should be created to demonstrate key security implementation tasks for the lay-audience, including private citizens and small and medium-sized businesses (SMBs).

These cookbook-like guides should be created with a user-friendly interface, perhaps using wiki-like technology with embedded videos, scripts and widgets to make them easy to follow. They should be comprehensive, addressing a large set of available technologies (e.g., Wordpress, Drupal, Joomla, etc. for website security guides).

Next Steps

Guide creation and rollout could occur in the following steps:

1. Identify the best practices and distill them into workable, affordable, solutions. Convert into search engine-friendly course content with an expert design team.
2. Overcome government reluctance and restrictions on recommending technologies, perhaps through crowd sourced input.
3. Six months of planning and production would lead to the program launch, with an incremental and timely rollout of content as it is developed, with periodic updates as technologies change.
4. Publicize the availability of the how-to guides.
5. The first guide, on website security (how to protect your website from DDoS, SQL injection, takeover, etc.), would be published within first year, followed by a guide on email security, etc.
6. An audience rating system similar to those used on Amazon or Netflix could be included, as well as carefully monitored public feedback forums.

This project would best fit within the Department of Homeland Security's mandate for improving cybersecurity for citizens and small organizations without dedicated IT staff. While NIST could provide resources there may be problems with NIST's requirement to remain technology agnostic. The National Initiative for Cybersecurity Education (NICE) could be the proper starting point. Alternately, these efforts could be driven by the Federal Trade Commission (FTC), which operates OnGuardOnline.gov, or the Small Business Administration (SBA).

The project would extend the traditional government role. New thinking on how to allow a government entity to provide actionable guidance without appearing to be picking winners might be needed. It is likely that vendors will push back and attempt to influence guide content; they should be encouraged to participate, but not given undue influence.

MOVING FORWARD

While the list of recommendations presented in this report is by no means comprehensive, and addresses only some key areas of the problem space, it does map out concrete possible actions that could have a real impact on the security of the Internet ecosystem. This section outlines the next step(s) on the path to implementing each recommendation to fruition and also discuss how long we believe it would take each recommendation, if implemented, to begin to yield visible results.

Next Steps

There are usually multiple paths to any goal and multiple steps along any path. The paths and steps taken are usually chosen based on a variety of factors including the person or persons leading the effort and the political and organization challenges those leaders face.

At the same time, there's a benefit to having a plan, as a prototype against which to compare alternatives. To that end, for each recommendation above, a possible next step has been sketched. Those steps are summarized briefly in the table below.

Recommendation	**Next Steps**
Make Critical Subsystems Field-Updatable	Develop incentives for industry and the research community to work jointly on this problem.
Enable Certificate Transparency and Security	Develop transparency mechanisms for certificate authorities.
Create a Framework for Managing Software Updates	Convene a workshop for the software update community to exchange ideas about how to support updates for unattended objects.
Make HTTPS the Least-Effort Scheme for Deploying a Website	US-based DNS registries (such as .com) should be required to issue a certificate with each domain name.
Cybersecurity Research Agenda	Incorporate the listed research topics into the national research agenda.
Establish an Internet Rescue Squad	Study whether existing entities can be leveraged, either as examples of how such an organization can work or as the base for such an organization.
Create a Cyber NTSB	Convene a study to recommend how best to create and operate a Cyber NTSB.
Standard Impact Statement	Develop models that can be used to inform an impact statement.
How Golden is Our Goose?	Fund multiple studies to evaluate the benefits of the Internet and the costs of imposing additional security requirements.
Identity: A Problem That Doesn't Need Solving	Decline to support efforts to associate an identity with each packet.
Encourage the Adoption of Routing Security	Implement a US Government policy of refraining from using domain name takedowns as an enforcement mechanism.
Enhancing the Security of the Internet of Things by Enclaves	Determine who is best placed to address the multi-industry issues this recommendation raises.
Create a List of Top Priorities	Solicit lists of top cybersecurity risks from public

	companies and collate them into a list of top priorities.
Lead by Example	Put in place mechanisms to implement PCAST's top two recommendations by FY 2015.
Re-Establish Trust in NIST's Cryptographic Standards Process	Have NIST update its procedures to be more open. Increase NIST CSD headcount. Declare NIST standards are not to be subverted by other Government agencies.
Develop Citizen and Small Business How-to Guides for Implementing Security	Task the proper agency to compile such a guide.

How soon might we see an impact?

Both what is achieved, and when it is achieved matter. In that light, this section briefly seeks to estimate how soon, if we moved to implement the recommendations, one might expect to see signs of improvement in cybersecurity.

For simplicity, for each recommendation we have estimated when we might see an impact if we chose to act on the recommendation. For simplicity, the impacts are divided into immediate (that is to say, some impact happens quickly) and medium-term. While many recommendations will also have long-term impacts, trying to fully characterize those impacts was too speculative.

Recommendation	**Immediate**	**Medium-Term**
Make Critical Subsystems Field-Updatable	Workshop of impacted communities and research community.	Beginning of standards activities for updating critical systems.
Enable Certificate Transparency and Security		
Create a Framework for Managing Software Updates	Workshop of software update community.	Initial software update systems for unattended devices.
Make HTTPS the Least-effort Scheme for Deploying a Website	New domains have certificates and new sites begin to use them.	A higher percentage of sites using HTTPS.
Cybersecurity Research Agenda	Research called out in solicitations.	
Establish an Internet Rescue Squad	Plan to create an internet rescue squad released for public comment.	Funding for Internet Rescue Squad budgeted.
Create a Cyber NTSB		Report with implementation recommendations. Possibly enough to budget for creation.
Standard Impact Statement		Draft impact statement released for public comment.
How Golden is Our Goose?		Studies funded on the tradeoffs between innovation and greater security in our

		systems.
Identity: A Problem That Doesn't Need Solving		
Encourage the Adoption of Routing Security	US Government announces it will not seek domain takedowns.	
Enhancing the Security of the Internet of Things by Enclaves	Identification of agency to coordinate enclave meetings.	Meetings of various enclaves of interest.
List of Top Priorities	A published list of top priorities.	Change in Government and industry priorities based on list.
Lead by Example	The US Government has implemented the PCAST recommendations.	
Re-Establish Trust in NIST's Cryptographic Standards Process	NIST given authority to enhance capabilities of CSD; official Government statement demonstrates respect for and value of NIST's standards process.	NIST in a position to reassure community of its increased ability to continue to work as a trusted independent authority in security.
Develop Citizen and Small Business How-to Guides for Implementing Security		Guides released for public comment.

It is hoped that this report will spur action to increase knowledge and affect new policies and practices that will significantly increase the security of the Internet ecosystem.

APPENDIX: Ideas Lab Participants

Ideas Lab Steering Committee: Susan Landau, Damon McCoy, Deirdre Mulligan, Craig Partridge, Jennifer Rexford, Stefan Savage, Dave Ward

Ideas Lab Participants:
Marjory Blumenthal
David Clark
Andy Ellis
Tom Forest
Stephanie Forrest
Sharon Goldberg
Eric Grosse
Emily Grumbling
Alex Halderman
Farnam Jahanian
Erin Kenneally
Robert Laddaga
Susan Landau
Wenke Lee
Damon McCoy
Danny McPherson
Keith Marzullo
Tyler Moore
Deirdre Mulligan
Andy Ozment
Craig Partridge
Vern Paxson
Ron Perez
Jennifer Rexford
Avi Rubin
Stefan Savage
Stuart Schechter
Fred Schneider
Matthew Scholl
Howie Shrobe
Richard Stiennon
Dave Ward
Detlof von Winterfeldt
Nicole Wong
Rebecca Wright