Notes for White House 60-day Cyber-Policy Review

March 25, 2009

This material is culled from discussions among a collection of cyber-security researchers that NSF/CISE organized for telephone briefings with Melissa Hathaway as part of the White House initiated 60-day cyber policy review. The document does not give a trustworthy systems research agenda because (i) much has already been published on identifying important open problems and (ii) the relative importance of various research problems is more likely to be informed by the results of the 60-day than to play a central role in what that report says. Rather, the document focuses on both how our community can help the administration and how the administration can help our community be more effective participants in our nation's efforts to design, build, and deploy trustworthy systems

The academic/open cybersecurity research community has already contributed many important solutions to the trustworthy systems landscape:

- Public key cryptography and the algorithms that employ it (e.g., digital signatures, and all of secure online credit card transactions).
- Static program analysis that finds vulnerabilities (e.g., buffer overruns) in systems and applications.
- Machine learning and data mining for spam filtering, virus/worm detection, credit card fraud detection.
- Widely-used authentication systems, such as Kerberos.
- Firewalls, intrusion detection systems, and integrity monitoring tools.
- Legal reforms to facilitate research (e.g., the DMCA rule exemption), provide data about vulnerabilities, and create incentives for better security practices (e.g., Security Breach Notification statutes).

And the community continues to contribute, including to immediate and operational cybersecurity problems. For example, NSF-funded researchers are today collaborating on the operational response for the Conficker malware bot army, a threat with immediate and potentially dire consequences.

The community is well aware that the challenges we face as a nation are great. But these challenges are also opportunities for all of us—the academic, industry, and government sectors—to work together with the goal of deploying trustworthy systems that could empower our citizens, our economy, and the environment. We, the academic community, are excited about playing an active and direct role in finding solutions. We see ourselves not only continuing our traditional modes of operation—problem-driven and curiosity-driven research—but also embracing new roles, if given the opportunity. For example, the community would welcome a chance to establish a modern version of the Manhattan Project on authentication and identity management. However, we are not currently positioned to play such a larger role, and doing so would need some explicit empowerment and assistance.

The nature of cybersecurity problems is quite different than what one usually associates with the problems that drive most scientific disciplines. These differences influence how solutions ought to be pursued:

a. Cybersecurity is not purely a technology problem, nor is it purely a policy (economic or regulatory) problem. Too often, technologists today do technology without understanding law, investment policies, economics; policy wonks make policy without understanding technology. Building trustworthy systems will require combining technology and policy.[1] Bridges between these communities are being built, and the importance of fostering this collaboration cannot be overstated. Cybersecurity is also a social and political problem. For example, is it sensible to view the Internet as a "commons"? And to what extent do basic notions, such as identity, privacy, accountability, property, and good behavior, as well as our expectations for services, derive from such a choice for the underlying context?

b. The prevailing attitude that security is "a problem to be solved" is naïve. Medicine is an appropriate analogy, since despite enormous strides in medical research, new threats continually emerge and old defenses (e.g., antibiotics) are seen to lose their effectiveness. As the nation pursues opportunities for sustainability, healthcare, and commerce, there will be on-going needs for cybersecurity research or else the trustworthiness of these systems will erode as threats evolve.

c. We need to move from reactive deployment of defenses to a more proactive stance; otherwise, attackers will have the upper hand. This move will likely require investing in long-term research to develop a "science base" and a complementary set of new policy frameworks and institutions. NSF, NSA, and IARPA recently ran a joint workshop on this topic, and NSF's Trustworthy Computing (with 387 active PI's) program currently supports research into foundations of trustworthy computing.

After reflecting on how the community might have an even greater impact, both in the short term and the longer term, we note the following.

a. A good deal of useful cybersecurity research and technology is not being widely deployed. Effective transition needs to be incentivized. Any solution will involve two sets of participants: researchers, who are technology producers, and end-system users, who could be technology adopters.

    - Experts in law, economics, and the social sciences will undoubtedly have insight into how to create a climate that encourages adoptions of new technologies. They need to be involved.

    - We might aspire to embracing a culture that values "R+D+D" for Research, Development, and Deployment. By extending the researcher culture to value deployment, just as it now values publication, we would incentivize researchers to

[1] Sometimes policy initiatives can solve technology problems; sometimes new technology can potentially sidestep or render less difficult societal tradeoffs that are seen as fundamental and daunting.

make any additional investments needed for their innovations to be used in practical settings.

b. The importance of education and the proven synergy between research and education cannot be over-emphasized in light of the pressing need to expand the workforce through education and training.

c. A small Trustworthy Systems Research Advisory Board, populated by researchers in systems and software sitting side by side with experts in law, public policy, and economics, would be an effective conduit and sounding board for executive branch decision makers grappling with our nation's challenges. Civilian researchers will provide a useful and important different perspective on trustworthiness challenges and solutions.

With regard to orchestrating and supporting the research enterprise, we would note:

a. A commitment is needed not just to adequate levels of research funding but also to continuity of research funding. Only with this kind of commitment can a community mature and have the confidence to attack the really hard problems (i.e., those problems that require a long-term research).

b. There must be a diverse ecology of research funding opportunities.

- NSF shouldn't be the sole significant source of research funding in cybersecurity. Opportunities should exist for other styles of research to be supported, such as research that is more closely aligned with specific problems, research that is better coordinated amongst larger numbers of investigators, research that involves significant numbers of supporting staff beyond the PI's. NITRD could provide far more effective collaboration.

- New federal initiatives—electronic health records, sustainability and the smart power grid, new transportation infrastructure etc.—that leverage trustworthy systems must include a commitment to engage cybersecurity researchers in the application domain and in understanding the unique needs of those systems.

c. Excessive classification works against the nation's interests.

- Classified research does not engage many of the nation's most capable cybersecurity researchers, is necessarily less likely to receive broad scrutiny by a diverse community of experts, and does not contribute to the education the next generation of cybersecurity researchers and practitioners. Classified research programs are also slow to impact the civilian cyberinfrastructure and its equipment, on which much of our nation's and our government's and even our military's critical infrastructure depends.

- Secrecy regarding cyberattacks shields the research community from the very data they need in order to understand the real problems. Greater transparency about

current defensive efforts (and what they are intended to achieve) would further help ensure that researchers are working on the right problems and are making sensible assumptions about the environment.

- Access to real data about operational networks is important. Significant threads of cybersecurity research critically depend on empirical insights into how different forms of activity actually play out in practice. How does malicious behavior manifest itself in large-scale systems? In addition, data sources not generally available today to researchers, such as Google's global perspective on malicious web activity and Verisign's visibility into domain registrations and their associated changes and infrastructure, could be invaluable for researchers. There are major legal, ethical, and business-sensitivity issues to overcome for researchers to obtain access to such data; but without the data, important research can not be pursued.

d. Academia has much to offer for evaluating cybersecurity technologies in a quantitative manner, and an "evidence-based" security doctrine backed-up by quantitative data would be a big step forward in providing a scientific basis for government and industry to make security investment and deployment decisions. Moreover, there is little incentive for another community to pursue this line of inquiry (just as there is little incentive today for vendors or service providers to provide this access to academia). Beyond the obvious short-term pay-off, the entire research community would be better positioned to produce work that is relevant if they can witness real adversaries and real users in a realistic context.

Fred Schneider (fbs@cs.cornell.edu) serves as the point of contact for questions and comments on this document.

This document was edited by Ed Lazowska (University of Washington) and Fred B. Schneider (Cornell University. It is endorsed by Steve Bellovin (Columbia University), Jean Camp (Indiana University), Fed Cate (Indiana University), David Clark (MIT), Kay Connelly (Indiana University), Lorrie Cranor (Carnegie Mellon University), Michael Cuiker (University of Maryland), David Dill (Stanford), Joan Feigenbaum (Yale), Stephanie Forrest (University of New Mexico), Johannes Gehrke (Cornell), Susan Hohenberger (Johns Hopkins University), Eric Johnson (Dartmouth), David Kotz (Dartmouth), Ed Lazowska (University of Washington), Pat Lincoln (SRI), Deirdre Mulligan (UC Berkeley), Andrew Myers (Cornell), Steven Myers (Indiana University), Helen Nissenbaum (NYU), Charlie Palmer (Dartmouth), Vern Paxson (UC Berkeley), Michael Reiter (UNC-Chapel Hill), Avi Rubin (Johns Hopkins University), William Sanders (UIUC), Stefan Savage (UCSD), John Savage (Brown), Fred Schneider (Cornell), Scott Shenker (UC Berkeley), Sean Smith (Dartmouth), Salvatore Stolfo (Columbia University), Roberto Tamassia (Brown), Paul Thompson (Dartmouth), Giovanni Vigna (UC Santa Barbara), David Wagner (UC Berkeley), Dan Wallach (Rice University), John Wroclawski (USC), Martin Wybourne (Dartmouth).