Security is Not a Commodity:
The Road Forward for Cybersecurity Research

Stefan Savage
UC San Diego

Fred B. Schneider
Cornell University

Version 4: February 3, 2009[1]

Computers at Risk[2], a 1991 report by the Computer Science & Telecommunications Board of the National Research Council begins:

> *We are at risk. Increasingly, America depends on computers. They control power delivery, communications, aviation, and financial services. They are used to store vital information, from medical records to business plans to criminal records. Although we trust them, they are vulnerable—to the effects of poor design and insufficient quality control, to accident, and perhaps more alarmingly, to deliberate attack. The modern thief can steal more with a computer than with a gun. Tomorrow's terrorist may be able to do more damage with a keyboard than with a bomb.*

When these words were written, the first web browsers were still on the drawing board and the Internet was a place for high-tech aficionados. How far we have come! Today, our dependence on inter-networked computing systems means that virtually every walk of American life—whether personal or commercial, public or private, civilian or military—is intermediated by computer systems. But virtually none of these systems are trustworthy; all are subject to attack; in fact, many are actively under attack today.

It is 2009 and we very much remain a nation at risk. Moreover, we are embarking on a trajectory that will only put us further at risk as we revamp energy distribution, revitalize our transportation systems, and computerize our health care records. We have no basis to place our trust in such systems. We must learn to engineer our inter-networked systems and defend them, appreciating that they will constitute a new battlefront as they are deployed.

Today's landscape

By virtually any metric—number of attacks, number of computers compromised, number of distinct malware variants, dollars lost, etc.—our existing engineering practices and defenses are failing us. Despite huge increases in direct IT security expenditures, which are estimated to reach $79 billion annually by 2010[3], the impact of cyberattacks has not been diminished. In

---

[1] For the most current version of this essay, as well as related essays, visit http://www.cra.org/ccc/initiatives

[2] *Computers at Risk: Safe Computing in the Information Age*. National Academies Press, Washington DC, 1991. http://books.nap.edu/catalog.php?record_id=1581

[3] *Information Security Products & Services – Global Strategic Business Report,* Global Industry Analysts, Inc., July 2007.

fact, the direct costs of cybercrime to US business are estimated at over $67 billion annually[4]. This is to say nothing of indirect costs, of costs to individual consumers, or of the unquantified impacts due to undetected attacks, industrial espionage, impacts to national security and so on. **We are losing the cybersecurity war, and our most potent adversaries are not yet even on the battlefield.**

The core of the problem is inherent in the nature of security itself. Unlike computer and communications hardware and software, security is not a commodity. It cannot be scaled simply by doing more. Security is holistic—a property of a system and not just of its components. Even a small change to a system can have catastrophic consequences for its security.

There are familiar and predictable technology curves by which computer processing performance, storage, and communication all scale over time. Security does not follow such a model. Instead, security is characterized by asymmetries:

- **Defenders are reactive, attackers are proactive.** Defenders must defend all places at all times, against all possible attacks (including those not known about by the defender); attackers need only find one vulnerability, and they have the luxury of inventing and testing new attacks in private as well as selecting the place and time of attack at their convenience.

- **New defenses are expensive, new attacks are cheap**. Defenders have significant investments in their approaches and business models, while attackers have minimal sunk costs and thus can be quite agile.

- **Defenses can't be measured, but attacks can**. Since we cannot currently measure how a given security technology or approach reduces risk from attack, there are few strong competitive pressures to improve these technical qualities. So vendors frequently compete on the basis of ancillary factors (e.g., speed, integration, brand development, etc.). Attackers can directly measure their return-on-investment and are strongly incentivized to improve their offerings.

The result is a cybersecurity industry built around defending against known attacks. Such a reactive stance might have made sense decades ago, when the numbers of attackers and their rate of innovation were relatively slow and unfocused. Today, however, the economic engine of profit-driven cyberattacks has transformed the threat landscape, and a plethora of well-funded attackers are now carefully and diligently exploiting the fundamental asymmetries.

Whereas attackers once actually had to understand how their attacks worked, today attacks come pre-packaged. It is easy enough to download or purchase malicious software on-line that is undetected by existing commercial security products, to buy and sell compromised hosts by

[4] *2005 FBI/CSI Computer Crime Survey*, January 2006. http://www.gocsi.com

the thousands, and to traffic in sensitive information extracted from those workstations and servers alike. Looking ahead, so-called embedded computing promises to expand the landscape of targets to the SCADA systems that control our utilities and industrial processes, the myriad computer systems controlling our automobiles and transportation infrastructure, and the mobile applications built upon the next generation of cell phones and portable devices. The new threats focused on these systems are unlikely to be addressed by simply re-purposing our existing defenses.

Today's security industry—serving civilian and military customers alike—has simply not been able to develop the "game changing" approaches necessary to address the challenges we now face, and there is no reason to believe this will change in the future. We can't succeed by focusing our defenses on past attacks, and we must move from a reactive stance to a proactive one.

Where we need to go

The natural source for risky, transformative, game-changing ideas is the research community, which specializes in technological innovation and is less driven by the short-term competitive pressures that tend to produce incremental efforts. Unfortunately, research in cybersecurity has been hamstrung on multiple fronts and, therefore, has not been able to develop a much needed science base to anchor cybersecurity innovations, nor has it been able to produce the range of solutions we require.

The first challenge has been financial. Federal expenditures for basic and applied cybersecurity research are tiny compared to the severity of the threat. And the absence of continuity in what little funding there is has stymied the development of a research community, which in turn leads to a national shortage in cybersecurity expertise. Indeed, *Toward a Safer and More Secure Cyberspace*, a very recent National Research Council CSTB report[5], reports that funding levels for cybersecurity research are low, preventing researchers from pursuing their promising research ideas and that, excepting the National Science Foundation (NSF), Federal funding agencies predominantly targeted short-term problems rather than addressing the harder, longer-term challenges that constitute our only hope to win this war. This echoes the findings in the President's Information Technology Advisory Committee's independent report *Cyber Security: A Crisis of Prioritization*[6], which stated that (i) cybersecurity solutions would emerge only from a vigorous and well funded program of research and (ii) that levels of funding were dangerously low to solve problems or to sustain a community of researchers. The PITAC report also noted the damage being caused by the lack of continuity in cybersecurity funding and by the inadequate oversight and coordination exerted by Federal government over its cybersecurity research programs. Particularly noteworthy is the observation that short-term

[5] *Toward a Safer and More Secure Cyberspace*. S Goodman and H. Lin (eds), National Academies Press, Washington, DC, 2007. Appendix B.6. http://books.nap.edu/catalog.php?record_id=11925

[6] *Cyber Security: A Crisis of Prioritization*. President's Information Technology Advisory Committee, Feb. 2005. http://www.nitrd.gov/pitac/reports/20050301_cybersecurity/cybersecurity.pdf

goals and a funding culture that targeted easily quantifiable progress have conspired to discourage funding for research efforts that, being more forward-looking, could provide the real pay-offs. PITAC argued, in vain, for a significantly increased investment in "fundamental research in civilian cybersecurity," noting that civilian systems comprise the lion's share of our nation's critical IT infrastructure, and that the government and military rely in large measure on civilian hardware and software components and systems.

The second challenge for cybersecurity research has been programmatic—a disconnect between what research is being solicited and what is really needed. Federal funding programs and commercial priorities have regarded computer security as a game of plugging leaks around a few established technical battlefronts defined by existing cybersecurity markets (e.g. firewalls, anti-virus, intrusion detection, etc). This is a doomed strategy that results from an entrenched mindset a decade or more out of step with the reality of our current adversaries. Thus, we not only need to reinvest in cybersecurity research, but we also need to re-imagine the scope of the cybersecurity problem itself; refocus our attention the same way our adversaries have refocused. We cannot afford simply to develop technologies that plug holes faster, but we need to think of security research more holistically, determining how most efficiently to block, disrupt, or disincentivize opponents.

We must develop a science base for security. We understand that the landscape spans attacks, defense mechanisms, and security properties. But we are only now starting to characterize the lay of the land in terms of how these features relate—answers to questions like: What security properties can be preserved by a given defense mechanism? What attacks are resisted by a given mechanism? Moreover, it is only recently that risk management has begun to supplant an unfortunate historical focus (arising from early military security needs) on absolute security: we are beginning think in terms of deploying defense as appropriate to the threat and to the value of what's being protected; we think of getting some security from the context in which a system is embedded; and we make trade-offs between security and other system properties as a function of the mission. Additionally, the importance of availability and integrity properties, which distinguishes many civilian applications, significantly changes the way one might build a system. In short, we need to articulate and organize a set of abstractions, principles, and trade-offs for building secure systems given the realities of the threats, of our security needs, and of a broad new collection of defense mechanisms and doctrines.

Given the dramatic acceleration in concrete attacks, we must also pursue research that explicitly targets the attacker. We must invest in mechanisms—both operational and forensic—for better attributing cyberattacks to the actors behind them. This is essential for applying virtually all other instruments of policy, from law enforcement to diplomacy. Moreover, we must consider not merely hypothetical opponents, but the real attackers we face today and those we expect to encounter tomorrow. By analogy, our military does not train against a hypothetical adversary with hypothetical resources, strategies and interests, but instead we design our forces and stratagems based on what is known about real and potential adversaries. Thus, in service of undermining the effectiveness of their attacks, we must explicitly research our attackers' motivations and resources, and better understand their

economic value chain, their structural strengths and weaknesses, the nature of their social and organizational networks, and how all these aspects might be best disrupted. Cybersecurity is inherently an asymmetric enterprise, and our best strategy as defenders against an asymmetric threat is to seek and exploit knowledge of the adversary.

Finally, we must prioritize developing better quantitative measures around cybersecurity risk, efficiency, and value. We cannot invest arbitrary amounts in securing our systems without better understanding the return on this investment.

The path forward

This past fall, a CSIS study[7] offered President Obama's administration a broad set of recommendations for improving our nation's cybersecurity posture. Among these is a call for increased investment in longer-term R&D. Asserting that cybersecurity is "now a major national security problem for the United States," the CSIS study suggests both short-term and long-term measures, and makes technical as well as non-technical recommendations. And to address the "Crisis of Prioritization," CSIS recommends placing cybersecurity expertise at the highest levels of the Executive branch.

The recommendations in the CSIS study, if adopted by the Obama administration, would allow the cybersecurity research community to make the progress that will be needed to enable many of the administration's initiatives. There is, of course, some distance to travel from deciding to invest in cybersecurity research to our nation actually making those investments:

- We must decide on an agenda of research to pursue.

- We must decide on the agents and manner for carrying out this research, mindful not only of the agenda but of having efficient ways to transition research results into practice.

Research Agenda. The question of defining a cybersecurity research agenda is not new, and the research community has devoted a good deal of time and energy to formulating answers. National Research Council CSTB reports in 1998[8] and 2007[5] offer remarkably consistent visions of what research needs to be done. The vision articulated not only includes technical topics in cybersecurity but also notes ways in which the cybersecurity problem should be seen as considerably broader.

The NRC CSTB research agendas suggest that beyond building a science base for cybersecurity and exploring various specific mechanisms and doctrines, investments also be made in two less obvious areas. First, attacks are possible because the systems we build have flaws ("bugs"). Thus, investing in research for building "correct" software pays dividends in improved

---

[7] *Securing Cyberspace for the 44th Presidency*. Center for Strategic and International Studies, Dec 2008. http://www.csis.org/media/csis/pubs/081208_securingcyberspace_44.pdf

[8] *Trust in Cyberspace*. F.B. Schneider (ed), National Academy Press, Washington, DC, 1998. http://books.nap.edu/catalog.php?record_id=6161

cybersecurity. Second, solutions to cybersecurity problems will almost certainly involve non-technical elements, and investments should be made there. Economic and regulatory incentives, for instance, are required to foster the needed substantial investments in defense by users and producers of systems, although such incentives can only be designed with knowledge of cybersecurity technologies; as another example, identity theft can often be traced to conflating an "identifier" (which necessarily is not secret) with an "authenticator" (which could be a secret) when trying to validate an identity claim, so changing how the law defines authentication could be a high leverage, but non-technical, method to reduce identity theft. Cybersecurity research that ignores these non-technical dimensions risks irrelevance.

Ecology of Federal Agencies. Fundamental research in cybersecurity today is supported primarily by the National Science Foundation. In the past, DARPA had been a significant source of funding for university researchers doing work in systems and security, but DARPA is no longer making those investments, instead preferring short-term and classified programs. DHS and the new I-ARPA have funded work in cybersecurity, but at significantly lower levels, and these agencies have also limited their investments to problems with a short-term horizon. DoD, through AFOSR and ONR, does fund some fundamental research in security, but the number of projects supported is relatively small and the funding is often through special one-time initiatives (i.e., the MURI program).

Various government agencies themselves perform research and development in cybersecurity. NIST's Security Division is probably the most visible player, through its role in defining Federal civilian cryptography standards and in advising civilian agencies on security best practices. With increased dependence on computing, additional investments in this NIST group would be highly leveraged both for government agencies and the nation. There are also significant research efforts within the services and the intelligence agencies. While in some cases this research is widely reported (NRL is extremely visible in the academic cybersecurity research community), frequently the work is classified and thus, by definition, has a far more circumscribed impact. Overall, the numbers of cybersecurity researchers in government employ are dwarfed by the numbers of researchers based in universities and the private sector.

This ecology of different government agencies with their different needs, goals, and cultures, in theory should yield a robust and diverse research climate. However, many of the potential benefits have not materialized, both because the inter-agency coordination has been voluntary and because tight budgets led some of the participants to focus cybersecurity research expenditures on short-term work, which they saw as better suited for their missions. The CSIS recommendation for cybersecurity expertise at the highest levels of the Executive branch could go far in addressing the coordination problems across the different agencies that fund research in cybersecurity.

Today, NSF is the only natural home for fundamental research in civilian cybersecurity. NSF's Cybertrust program, the likely agent for funding investigations that will have high payoff, is woefully under-resourced. In the past, what had been DARPA's style complemented NSF's style by supporting larger groups (3-5 investigators) to work for relatively longer periods (5-10 years)

in order to take a game-changing idea to a demonstrable embodiment. The NSF and former DARPA styles are indeed complementary, and both ought to be supported. Another point of contrast between the different styles concerns the manner they use for review and selection of proposals for funding. External peer-review by the research community leads to funding work with a different character from internal review, which enables programmatic goals to play a role in project selection.

New initiatives in energy, transportation, and electronic medical records will almost certainly require solving new cybersecurity research questions. Past experience is that failing to engage the community early in such initiatives is a mistake. Cybersecurity research is not done well in a vacuum from applications, and there is no substitute for direct experience with these applications. Thus, part of these new initiatives should be to involve the cybersecurity research community, so they can help ensure inter-networked systems required will be sufficiently trustworthy.