United States House of Representatives
Committee on Science and Technology
Research and Science Education Subcommittee

Hearing on:
Cyber Security R&D

Dr. Fred B. Schneider
fbs@cs.cornell.edu
(607) 255-9221

Samuel B. Eckert Professor of Computer Science
Cornell University
4115C Upson Hall
Ithaca, New York 14853

June 10, 2009

Testimony of Fred B. Schneider
Samuel B. Eckert Professor of Computer Science, Cornell University

June 10, 2009

Good morning Mr. Chairman and members of the Committee. I appreciate this opportunity to comment on cyber-security research and education. I am Fred B. Schneider, a Computer Science professor at Cornell University and Chief Scientist of the NSF-funded TRUST[1] Science and Technology Center, a collaboration involving researchers at U.C. Berkeley, Carnegie-Mellon University, Cornell University, Stanford University, and Vanderbilt University.

I have been a Computer Science faculty member since 1978, actively involved in research, education, and in various advisory capacities for both the private and public sectors. Besides teaching and doing research at Cornell, I today serve as member of the Dept. of Commerce Information Security and Privacy Advisory Board (ISPAB), as a member of the Computing Research Association's board of directors, and as a council member of the Computing Community Consortium. I also co-chair Microsoft's TCAAB external advisory board on trustworthy computing.

Our nation's increasing dependence on computing systems that are not trustworthy puts individuals, commercial enterprises, the public sector, and our military at risk. If anything, this dependence will accelerate with new initiatives such as the "smart grid" and electronic healthcare records. Increased data, increased networking, and increased processing all mean increased exposure. These systems need to work as we expect—to operate despite failures and despite attacks. They need to be trustworthy.

The growth in attacks we are seeing today should not be surprising. The more we depend on a system, the more attractive a target it becomes to somebody intent on causing disruption; and the more value that is controlled by a system, the more attractive a target it becomes to somebody seeking illicit gain. But more disturbing than the growth in attacks is that our defenses can't keep up. The core of this problem is the asymmetric nature of cyber-security:

- 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.

[1] Team for Research in Ubiquitous Secure Technology.

- New defenses are expensive to develop and deploy; 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.

- The effectiveness of defenses cannot be measured; attacks can. Since we cannot currently quantify how a given security technology or approach reduces risk from attack, there are few strong competitive pressures to develop defenses. So vendors frequently compete on the basis of ancillary factors (e.g., speed, integration, brand development, etc.). Attackers see their return-on-investment and have strong incentives to improve their offerings.

The result has been a cyber-security mentality and industry built around defending against known attacks. Our defenses improve *only* after they have been successfully penetrated. And this is a recipe to ensure some attackers succeed—not a recipe for achieving system trustworthiness. We must move beyond reacting to yesterday's attacks (or what attacks we predict for tomorrow) and instead start building systems whose trustworthiness derives from first principles.

Yet today we lack the understanding to adopt that proactive approach; we lack a "science base" for trustworthiness. We understand that the landscape includes 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? How can we overcome the inevitable imperfections in anything we might build, yet still resist attacks by, for example, forcing attackers to work too hard for their expected pay-off. Having a science base should not be equated with implementing absolute security or even concluding that security requires perfection in design and implementation. Rather, a science base should provide—independent of specific systems— a principled account for techniques that work, including assumptions they require and ways one set of assumptions can be transformed or discharged by another. It would articulate and organize a set of abstractions, principles, and trade-offs for building trustworthy systems, given the realities of the threats, of our security needs, and of a broad new collection of defense mechanisms and doctrines. And it would provide scientific laws, like the laws of physics and mathematics, for trustworthiness.

An analogy with medicine can be instructive here. Some maladies are best dealt with in a reactive manner. We know what to do when somebody breaks a finger, and each year we create a new influenza vaccine. But only after significant investments in basic medical sciences are we starting to understand the mechanisms by which cancers grow, and developing a cure seems to require that kind of deep understanding. Moreover, nobody believes that disease will some day be a "solved problem." We make enormous strides in medical research yet new threats emerge and old defenses (e.g., antibiotics) are seen to lose their effectiveness.

Like medicine and disease, system trustworthiness is never going to be a "solved problem". There will be no "magic bullet" trustworthiness solution, just as there is not

going to be a miracle cure for all that ails you. We must plan to make continuing investments, because the problem will continue evolving:

- The sophistication of attackers is ever growing, so if a system has vulnerabilities then they will find it. Any assumption made when building a system does, in fact, constitute a vulnerability, so every system will have vulnerabilities of one sort of another. And with enough study, attackers will find these vulnerabilities and find ways to exploit them.

- The technology base used by our systems is rapidly changing. Systems are replaced on a 3-5 year time span, not because computers or software wear out but because newer software and hardware offers improved functionality or better performance (which is then leveraged into new functionality). New systems will work differently, will involve different assumptions, and therefore will require new defenses.

- The settings in which our computing systems are deployed and the functionality they provide is not static. With new settings come new opportunities for attack and disruption, whether it is creating a blackout by attacking the "smart grid" or stalking somebody by planting a virus on a GPS-equipped cell phone.

We can expect to transcend the constant evolution only through the understanding that a science base provides. A science base is also our only hope for developing a suite of sound quantitative trustworthiness measures, which in turn could enable intelligent risk-management decisions, comparisons of different defenses, and incentivize investments in new solutions.

A science base for trustworthiness would not distinguish between classified and unclassified systems, nor would it distinguish between government and private-sector systems. The threats and trade-offs might be different; the principles are going to be the same. But even an understanding of how to build trustworthy systems for the private sector would by itself be useful in military and government settings, simply because so-called COTS (commercial off the shelf) technologies that are developed by the private sector for the private sector are widely used within the government too.

Many equate cyber-security research with investigations solely into technical matters. This oversimplifies. Achieving system trustworthiness is not purely a technology problem. It also involves policy (economic and regulatory). Technological solutions that ignore policy questions risk irrelevance, as do policy initiatives that ignore the limits and capabilities of technology. So besides investing in developing a science base for trustworthiness, we must also invest in research that bridges the technical and the non-technical. We need to understand when we might get more traction for trustworthiness from a policy solution than from a technology one. For example, identifiers—your mother's maiden name, your credit card number, your bank account number, and your social security number—are not a good basis for authentication because they will be known to many. So regulation that prohibits the use of identifiers as authenticators might

more effectively defend against identity theft than new technology could. As another example, there is talk about making the Internet more secure by adding the means to trace packets back to their senders. But the Internet is as much a social construct as a technological one, and we need to understand what effects proposed technological changes could have; forgoing social values like anonymity and privacy (in some sense, analogous to freedom of speech and assembly) in order to make the Internet more-trustworthy might significantly limit the Internet's utility to some, and thus not be seen as progress.

Investments in cyber-security research are best accompanied by investments in cyber-security education, because this provides an efficient path for the research to reach industry where it can be applied. In particular, research undertaken in academia not only engages some of our nation's best and brightest researchers but because these researchers are also teachers, new generations of students can be exposed to the latest thinking from the people who understand it best. And when these students graduate and move into the workplace, they will bring this knowledge and understanding with them. Moreover, faculty in this dual role of researchers and teachers have incentives to write textbooks and prepare other teaching materials that allow dissemination of their work to a very wide audience, including teachers elsewhere.

Question: *Does the current range of federally supported research adequately address existing cyber security threats as well as new and emerging threats? If not, what are the research gaps, and how would you prioritize federal research investments in cyber security?*

Federal expenditures for unclassified cyber-security research do not match the severity of the threat. IT security expenditures are estimated to reach $79 billion annually by 2010[2]. According to the NITRD *Networking and Information Technology Research and Development Program*[3], $342.5M is being requested for FY2010 "Cyber Security & Information Assurance." This means Federal budget requests for unclassified research in system trustworthiness total roughly .4% of the expenditures that might be leveraged by the research. Moreover, anecdotal information about specific funding programs at various key Federal agencies suggests that only a portion of the $342.5M is spent on academic research in cyber-security. It then comes as no surprise to find the recent National Research Council CSTB report *Toward a Safer and More Secure Cyberspace*[4] stating that funding levels for cyber-security research are low, preventing researchers from pursuing their promising research ideas. And this echoes the findings in the President's Information Technology Advisory Committee's independent report *Cyber*

[2] *Information Security Products & Services – Global Strategic Business Report*, Global Industry Analysts, Inc, July 2007.
[3] *The Networking and Information Technology Research and Development Program*. Report by the Subcommittee on Networking and Information Technology Research and Development, May 2009. page 21, http://www.nitrd.gov/Pubs/2010supplement/FY10Supp-FINAL-Preprint-Web.pdf
[4] *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

Security: A Crisis of Prioritization[5] which stated that (i) cyber-security 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 NRC CSTB report also states that, excepting the National Science Foundation (NSF), Federal funding agencies predominantly target short-term problems rather than addressing the harder, longer-term challenges that constitute our only hope to win this war. A culture that targets easily quantifiable progress is particularly dangerous, because it discourages funding research efforts that, being more forward-looking, could provide the real pay-offs.

The PITAC report also noted damage being caused by the lack of continuity in cyber-security funding and by the inadequate oversight and coordination exerted by Federal government over its cyber-security research programs. For example, a lack of funding continuity stymies the development of a research community, because younger faculty and graduate students are disinclined to enter fields where future funding is uncertain. This, in turn, leads to a national shortage in cyber-security expertise.

PITAC argued, in vain, for a significantly increased investment in "fundamental research in civilian cyber-security," 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. Moreover, expenditures by the private sector for long-term cyber-security research have historically been quite small, probably because return on such investments is expected to be low. If the Federal government doesn't make these investments then nobody else will, and we all miss the opportunity for the revolutionary advances that are unlikely to result from the current regime of funding evolutionary steps. By the same token, the existence of a healthy IT-security industry suggests that the private sector does make investments in short-term research; so there is a less-compelling reason for Federal investments here.

There is a disconnect between research being funded and what is needed. Federal research funding has been too focused on a few established technical battlefronts (e.g. firewalls, anti-virus, intrusion detection, buffer overflows, etc.). In some cases, this focus reflects views held by researchers; in other cases, the focus comes from program management in the funding agencies. Whichever it is, this mindset is a decade or more out of step with the reality of our current adversaries. We need to re-imagine the scope of the cyber-security problem itself and refocus our attention the same way our adversaries have refocused. We cannot afford simply to develop technologies that plug holes faster; we need to think of security research more holistically, determining how most efficiently to block, disrupt, or disincentivize opponents.

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

- We must establish a goal of developing a science base for trustworthiness, as discussed in detail above. Such a science base is crucial for understanding how to build systems that are trustworthy.

- We must investigate mechanisms—both operational and forensic—for better attributing cyber-attacks to the actors behind them, because this is essential for applying virtually all other instruments of policy, from law enforcement to diplomacy. This approach might well be a last resort, invoked only after defenses to prevent attacks have failed. So it needs to be an option, despite being technically quite challenging as well as raising non-technical questions ranging from privacy all the way to international law.

- We must consider not merely hypothetical opponents, but the real attackers we face today and those we expect to encounter tomorrow. The military does not train against a hypothetical adversary with hypothetical resources, strategies and interests, nor should cyber-security researchers investigate defenses absent that information.

- We must prioritize developing better quantitative measures around cyber-security risk, efficiency, and value. The government and the private sector cannot invest arbitrary amounts in securing our systems without better understanding the return on this investment.

- We must invest in research that bridges policy (regulation and economics) with technology. To do research in technology without knowledge of policy or vice versa risks irrelevance.

- We must better understand the human element in our systems. Too often system security is synonymous with inconveniencing users. And users are inclined to circumvent security controls they find inconvenient, defeating a system's defenses even before it is attacked.

- We must continue to invest in research concerned with building software systems: operating systems, networks, programming languages, formal methods, database systems, etc. Ultimately, the things that undermine a system's trustworthiness will be traced to errors in design, implementation, requirements, or assumptions—subjects that are studied by software researchers. And we must continue making research investments in the relevant theoretical areas, such as logics and cryptography.

While there is certainly both a role and need for undertaking classified research in trustworthy systems, there are significant limitations that come with the secrecy. Classified research does not engage many of the most capable cyber-security researchers, is necessarily less likely to receive broad scrutiny by a diverse community of experts, and does not contribute to educating the next generation of cyber-security researchers and practitioners. Classified research programs are also slow to impact the civilian cyber-

infrastructure and its equipment, on which so much of our nation's critical infrastructure depends.

Having an Ecology of Federal Agencies is Valuable. There once was a diverse ecology of funding sources for the various styles and topics that trustworthiness research spans, but that ecosystem has been eroding as funding agencies have redefined their priorities. Some of these decisions are difficult to defend, given the central role that system trustworthiness plays in the missions these agencies are suppose to support.

Funding from a single agency (NSF) now dominates unclassified Federal cyber-security research. In the past, DARPA had been a significant source of funding for university researchers doing work in systems and security, but for the last eight years DARPA has not been making those investments. DHS has funded work in cyber-security, but at significantly lower levels and focusing on problems with a short-term horizon. DoD, through AFOSR, ARO, and ONR, does fund some fundamental research in security, but the number of projects supported is relatively small and some of the funding is for special one-time initiatives (i.e., the MURI program). IARPA inherited from its predecessor organizations a small but strong trustworthiness research program. That, however, is being terminated, and new programs to take its place have been slow to get started. Also, the funding philosophy at IARPA appears to be oriented more toward production of quantifiable results than toward open-ended curiosity-driven explorations.

This ecology of different government agencies with their different needs, goals, and cultures, could 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 reduce their cyber-security research investments and/or to focus those expenditures on short-term work, which they saw as better suited for their missions.

Today, NSF is the only natural home for fundamental research in civilian cyber-security. They not only fund single-investigators doing more-theoretical work, but they also fund larger-scale multi-investigator efforts that involve prototyping non-trivial systems. NSF's Trustworthy Computing (formerly Cyber Trust) 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 complementary, and both ought to be supported. Another point of contrast between the different styles concerns the manner they review and select proposals for funding. External peer-review by the research community leads to funding work having a different character from internal review (where programmatic goals play a role in project selection).

There is a tension between maintaining a diverse ecology of federal agencies to fund trustworthiness research and allowing each individual funding agency the autonomy to alter its priorities. So we must be mindful: seemingly local decisions within an agency

actually can have a broader impact by changing the Federal portfolio of trustworthiness research (as well as changing the total amount of Federal expenditures for trustworthiness research). This tension would be resolved if a coordinating body were to monitor such decisions and offset their impact on the Federal portfolio by allocating additional resources and recreating the now-absent styles at agencies electing to continue funding trustworthiness research.

Finally, it is worth noting that new initiatives in energy (e.g., a "smart grid"), transportation, and electronic medical records will almost certainly require solving new trustworthiness research questions. A failure to engage the community early in such initiatives is a mistake. This kind of trustworthiness research is not done well in a vacuum from applications; there is no substitute for direct experience with the application area. Thus, part of these new initiatives should be to involve the trustworthiness research community, so they can help ensure that the inter-networked systems required will be ones we can depend on.

Question: *What is the state of cyber security education? Are future cyber security professionals being adequately trained by colleges and universities to meet anticipated demands of the private sector? If not, what kind of cyber security training is appropriate and necessary for institutions to develop, and for what kinds of students?*

The University Landscape. Cyber-security professionals are today not being adequately trained to meet the needs of either the private sector or the public sector.

- **Part of the problem is resources**. University Computer Science (CS) departments lack the faculty to offer the relevant courses. Few faculty members have the necessary expertise to offer courses in this area. And even if a CS department has managed to hire a few cyber-security specialists, they will likely also be involved in teaching the large complement of other classes that need to be covered by a department giving undergraduate and graduate CS degrees.

- **Part of the problem is content**. The field is relatively young and fast moving. There is not yet widespread agreement about what technical content must be covered, which makes this an exciting time to be teaching cyber-security at the university level. But it also means that textbooks and other teaching materials have short lives unless they are frequently revised, which is a disincentive to some authors. So there are fewer good textbooks than would be found in a more mature subject. Yet, creating agreement on content by legislating a curriculum would be a serious mistake at this point, because it would retard the dissemination of new ideas to students and it would discourage faculty from writing texts that reflect improvements in our understanding of the field.

A Cyber-Security Professional Degree. I believe that a well trained cyber-security professional needs to have exposure to a broad variety of topics. One would expect to see courses that cover technical topics, such as computer security principles, distributed

systems and networking, systems reliability, software engineering, cryptography, and user interfaces and human factors. But I also strongly advocate exposure to non-technical topics, including cyber-law (intellectual property law, communications law, privacy law), ethics, economics of computing and networking, business strategy, and human relations (i.e., management of people). This broad education would enable a cyber-security professional to use all conceivable technical and policy tools for achieving trustworthiness. It would also ensure that solutions could be evaluated in a broader societal context, so that risk-management and trade-offs between different social values (such as privacy versus accountability) can be contemplated.

There is likely more than 1 year's worth of content past today's CS BS degree, but there is probably less than 3 years of course material. This would argue for creating some sort of graduate, professional degree program. It would be designed so that its students would learn both the technical and the non-technical topics needed to define and develop trustworthy computing systems, manage them, and oversee their deployment, use, and evolution.

Undergraduate Education. Computer Science departments today educate students to pursue a rather diverse set of careers. And, in particular, not all undergraduate Computer Science majors are headed for system-building careers. Thus, it would be inappropriate to impose a cyber-security requirement on all graduates from a Computer Science department. The more sensible model would be for universities to offer a *programme of study* for system trustworthiness, analogous to pre-law or pre-med. Such a programme is typically not associated with a single university department but rather offered in conjunction with a various majors; it prescribes a set of courses for the electives available in that department's major. The courses would cover the subjects outlined above in connection with the cyber-security professional degree. And it should be open to students in the various relevant majors.

Finally, it certainly seems reasonable that students destined to build systems—no matter what their major—should have exposure to the basic ideas needed for making those systems trustworthy. This means that they need exposure to basic cyber-security, software engineering, and various systems topics (operating systems, networking, etc.). Such students will be found enrolled in various majors. So while the CS department is the obvious place to offer these courses, the courses will not be populated only by CS majors. And this has implications concerning what pre-requisites can be assumed.

Biographical Sketch

Fred B. Schneider is Samuel B. Eckert Professor of Computer Science at Cornell University. He joined the Cornell faculty in Fall 1978, having completing a Ph.D. at Stony Brook University, preceded by a B.S. in Engineering from Cornell in 1975. Schneider currently also serves as the Chief Scientist for the NSF-funded TRUST Science and Technology Center, which brings together researchers at U.C. Berkeley, Carnegie-Mellon University, Cornell University, Stanford University, and Vanderbilt University.

Schneider's research has focused on various aspects of trustworthy systems—systems that perform as expected, despite failures and attacks. His early work concerned formal methods to aid in the design and implementation of concurrent and distributed systems that satisfy their specifications; he is author of two texts on that subject: *On Concurrent Programming* and *A Logical Approach to Discrete Mathematics* (co-authored with D. Gries). He has also known for his research in theory and algorithms for building fault-tolerant distributed systems. For example, his paper on the "state machine approach" for managing replication brought an SOSP "Hall of Fame" award for seminal research. More recently, his interests have turned to system security. His work characterizing what policies can be enforced with various classes of defenses is widely cited, and it is seen as advancing the nascent science base for security. He is also engaged in research concerning legal and economic measures for improving system trustworthiness.

Schneider was elected Fellow of the American Association for the Advancement of Science in 1992, the Association of Computing Machinery in 1995, and the Institute of Electrical and Electronics Engineers in 2008. He was named Professor-at-Large at the University of Tromso (Norway) in 1996, and was awarded a Doctor of Science *honoris causa* by the University of NewCastle-upon-Tyne in 2003 for his work in computer dependability and security.

Schneider has served since Sept 2006 as a member of the Information Security and Privacy Advisory Board (ISPAB), which advises NIST, the Secretary of Commerce, and the Director of OMB on information security and privacy issues pertaining to Federal Government Information Systems. He chaired the National Academies CSTB study on information systems trustworthiness that produced the 1999 volume *Trust in Cyberspace*. He also served as a member of CSTB from 2002-2008 and from 2004-2007 on the CSTB study committee for improving cyber-security research. Schneider was a member of the NSF CISE advisory committee 2002-2006. And in Fall 2001, he chaired the United Kingdom's pentennial external review of research funding for academic Computer Science.

In 2007, Schneider was elected to the board of directors of the Computing Research Association (CRA) and appointed to the steering committee of CRA's Computing Community Consortium. CRA is an association of more than 200 North American academic departments of computer science, computer engineering, and related fields; part of it's mission is to strength research and advanced education in the computing fields and

to improve public and policymaker understanding of the importance of computing and computing research in our society.

Schneider is a frequent consultant to industry, believing this to be an efficient means of implementing technology transfer as well as learning about the real problems. He is co-chair of Microsoft's Trustworthy Computing Academic Advisory Board, which comprises outside technology and policy experts who meet periodically to advise Microsoft about products and strategy. He also provides technical expertise in computer security as well as more broadly to a variety of firms, including: BAE Systems, Fortify Software, Lockheed Martin, and Microsoft.