

A Tale of Three Cyber-Defense Workshops

Since 2006, the National Cyber Defense Initiative (NCDI), an ad hoc group of cyber professionals, has been working behind the scenes to help inform the US research agenda for strategic cyber defense. These activities have been in anticipation

of the country and its leadership developing the will to aggressively invest in the area to reduce the significant risk the country faces. An important part of the NCDI's activities has been sponsorship of three workshops so far: the 2006 Safe-Computing Workshop, the 2007 Opening-Moves Workshop, and the 2008 Industry Workshop. Here we look at them and their results. For details and reports, visit <http://ncdi.nps.edu>.

over defenders, and deterrence is so weak, that it seems only a matter of time before a cyber catastrophe occurs. Many studies have identified a significant national threat from these adversaries. Attackers might use automated techniques that take effect at cyber speeds, whereas defenders have few tools to achieve responses at more than human speeds. Our inability to distinguish highly organized, well-funded attacks—for example, by nation states and other powerful groups—from nuisance attacks hampers selection of appropriate responses.

Short-term measures are essential but insufficient. “Patch and pray” efforts are the only large-scale defenses available in the short term. Efforts to solve the problems of the day, although laudable, can't result in incremental progress toward a fundamentally safer status quo. No matter how fast we patch, the spectrum of attacks will outpace our ability to design and deploy patches. System security isn't easily divided into orthogonal subproblems that we can treat separately; it's a holistic property of an overall design, its implementation, and its subsequent operation. Failure to invest in fundamental, long-term solutions will guarantee attackers a continuing advantage.

Market forces won't drive change. Most consumers poorly understand system security. Lack of comprehensive reporting of security incidents impedes quantification of risks and benefits and hides the information essential to

O. SAMI
SAYDJARI
*Cyber Defense
Agency*

CYNTHIA E.
IRVINE
*Naval
Postgraduate
School*

of the country and its leadership developing the will to aggressively invest in the area to reduce the significant risk the country faces.

An important part of the NCDI's activities has been sponsorship of three workshops so far: the 2006 Safe-Computing Workshop, the 2007 Opening-Moves Workshop, and the 2008 Industry Workshop. Here we look at them and their results. For details and reports, visit <http://ncdi.nps.edu>.

The Safe-Computing Workshop

This workshop was a response to concerns that the nation's computers and computer networks are increasingly at risk because of inherent vulnerabilities and escalating threats. The NCDI convened the workshop to explore whether and how present technology trends might be altered to lead us to a future in which each day wouldn't usher in a new set of vulnerabilities critically needing mitigation. More than 60 experts in system security, processor design, OSs, programming languages, networking, and applications from diverse backgrounds in academia, government, and industry attended. They aimed to review past and current strate-

gies for building secure systems and to discuss potential strategies that apply improved or alternative approaches to build more secure, more resilient systems. They also aimed to be as realistic as possible about the future.

The workshop's primary findings are sobering:

- The national information infrastructure is vulnerable, and that vulnerability is increasing.
- There's a compelling and urgent need to dramatically reduce that vulnerability.
- Major, strategic investments could significantly reduce vulnerability over a 5- to 10-year period.

So, without a major change in direction, adversaries would be able to exploit weaknesses in US cybersecurity. They could then deal a critical blow to one or more of the country's major sectors, such as banking, energy, and telecommunications.

Detailed Findings

The workshop's findings fall into the following six categories.

Attackers rule; disasters are likely. Today, attackers have such a strong, asymmetric advantage

make markets work. Even though the rise in attacks of all sorts has made security a more visible concern in the public consciousness, consumers still lack the product security information needed to make wise purchase decisions. Security is a public good in the sense that it's a system property, and frequently one weak or misconfigured component can sabotage an entire system, just as one polluter can ruin a stream.

Until liability for computing is similar to liability for other products, the potentially powerful effect of market forces is unavailable to address this problem. Even if IT product liability were available, it would more likely encourage codification and acceptance of best practices based on current technology than motivate fundamental change.

Usability and manageability must be part of the solution. Many security incidents could be avoided if systems were properly configured and operated, yet assuring proper configuration of an entire system is a daunting task. We must find and apply security metaphors that help people understand how to operate systems safely. To be effective, security must blend into the fabric of the system (and society).

New technology can catalyze major changes. New hardware and software technologies could bring beneficial disruptive change. Software productivity will be essential to success. More specifically, we must be able to efficiently produce software that also can be rapidly certified as safe to integrate with existing certified system elements. Solving this problem will force us to rethink the foundations of security and develop new human capital—partly by making past experiences in secure-system development more accessible to new researchers. The solution will require contributions from many

disciplines and will necessitate new research-funding strategies to achieve the needed collaboration. This effort can draw from and replenish the deep reservoir of security knowledge that's draining rapidly as today's experts retire, and it can enable progress in the medium term as long-term research progresses.

Only a national initiative will make a real difference. The overall security problem is too big for a single individual, company, institution, or even government sector to address. Moreover, the incentives for security are broken at many levels; fixing them at a single level won't change the landscape. However, a national initiative, properly organized and executed, could. (Authors' note: Of course, to fully address the problem will ultimately require global participation.)

Visions for a National Initiative

Although the workshop didn't develop a consensus vision for a national initiative, attendees developed several vision statements and considered them as possible bases for a national initiative. The group approached the problem as systemic and concluded that an effective national initiative would dramatically improve the US cybersecurity posture over the next 5 to 20 years.

National immunity from cyber attack. The idea of creating wholly new, secure, and adaptive computing systems that are self-checking and self-correcting aroused great interest among the participants. If this technology succeeded, it would transform the whole hardware and software computing fabric into computing resources loyal to individuals' and countries' security policies, even if an adversary penetrated national networks or gained privileged presence on those resources.

Economic growth from personal privacy and identity. Security technology is limiting Internet commerce's growth. If new personal-privacy and identity-management technologies succeeded, new products and services would spark a new round of investment. Such investment would in turn accelerate movement to the next generation of Internet technology: ubiquitous wireless and wire-line networking in which voice, video, and data travel with you anywhere in the world.

Radically reinventing cybersecurity technology to improve the quality of life. Experts at the workshop agreed that security research must start now and continue over the next 5 to 10 years, to fully develop the economic potential of networking technologies that will mature 20 to 30 years from now. A national safe-computing initiative might include incremental improvements to current commercial products. However, it must also extend to a "clean slate" reanalysis of system architectures and development approaches that aims to create fundamentally secure computing architectures and ecosystems. Such an initiative can produce technology sufficiently compelling and diverse for use in components serving high-threat environments. Eventually, this technology can be incorporated into commercial products, just as we've seen spread-spectrum, frequency-hopping, and GPS technologies move from custom military applications to consumer products worldwide.

Vision Statement

After the workshop, the NCDI developed a single vision statement:

Over the next ten years transform the cyber infrastructure to be resistant to attack so that critical national interests are protected from catastrophic

damage and society can confidently adopt new technological advances.

Conclusions

Not long ago, the planning of a national initiative such as the one envisioned in the workshop would have found little backing from a group as distinguished, diverse, and expert as the one assembled at the workshop. But the world has changed significantly since the Internet bubble days, when companies often saw security simply as an impediment to getting a product quickly to market. Now, major software vendors are willing to delay product releases for more than a year to forestall security embarrassments. Spam, phishing, and identity theft have substantially raised public awareness of some aspects of computer and network security. Other vulnerabilities, although less well understood by the public, are no less real.

But even hundreds of millions of dollars invested in trying to secure a single, large, complex system designed with security as an afterthought won't be enough to break the penetrate-and-patch cycle. Putting our computing and networking technologies on a sound footing will require a significant national investment.

The Opening-Moves Workshop

With approximately 35 participants, this workshop aimed to help develop a framework and plan to protect the US cyber infrastructure from strategic damage. The workshop took a bottom-up approach regarding *moves*—strategic actions needed to significantly reduce risk. It took a top-down approach in terms of *end-states*—operational capabilities needed to achieve information sovereignty. The end-states focus on systems critical to US interests. They characterize capabilities to be generated from a proposed 10-year

research program combined with nearer-term actions in computer security. Essential to this program is the creation of systems consistent with privacy expectations in our democracy.

Participants discussed the preliminary results with a small set of industry and government experts to establish the initial conclusions and recommendations. Further work is needed to describe the proposed program's structure. Overall, the workshop's attendees were optimistic that focused research could produce technology setting the stage for industry adoption of significantly more secure and dependable critical systems.

Here we describe a few highlights from the week's activities.

Desired End-States

Participants identified six essential end-states. In many cases, these end states are synergistic; determining their interdependencies will require additional effort.

First, continuity of critical information infrastructure operations is necessary. The US needs a transformed, resilient cyber infrastructure that will sustain critical functions in an attack.

Second, critical assets must be protected and, in some cases, defended from strategic damage. It must be economically infeasible for an adversary to cause strategic damage to the US critical infrastructure.

Third, cyber early-warning systems are necessary. Local through global cyber situation awareness should enable us to know what hardware and software are on critical system platforms, what's connected to the network, who's on the network, and what traffic is flowing over it.

Fourth, data-tight systems must be available to prevent unauthorized leaks or exfiltration (the opposite of infiltration) of critical information and intellectual property. We must ensure accountability for information flows.

Fifth, extensible systems that safely embrace new technology will let system owners confidently add new functions without fear of compromising existing system functionality or assurance claims.

Finally, quantifiable security and dependability will let system owners and defenders make informed decisions. Whenever possible, we should deploy new metrics and measurement technology that let us quantify security and dependability on the basis of realistic assumptions.

Promising Moves

To achieve the end-states, participants suggested actions they deemed particularly promising.

First, we should embrace stratification and partitioning as architectural principles. Networks and systems should be physically and logically organized so that fallback operations as well as rapid recovery and repair from attacks, even unexpected ones, are possible. As a matter of policy, critical security components should be stratified and partitioned. Networks that have moved away from these concepts should be reorganized. Furthermore, as in telephone systems, the control plane's critical data and functions should be separated from the operational plane.

We should determine the value of critical cyber-infrastructure functions and prioritize their criticality. As critical infrastructure functions become automated and integrated, the cost of operating without them because they're vulnerable to cyber attack must be calculated to assess their mission-criticality. We must quantify recovery and rollback for each critical cyber-dependent infrastructure.

We should develop rich test and analysis environments that use the best cyber strategies and tactics to vet theories of cyber defense and offense and new mechanisms and operators. Dif-

ferent test environments, with a range of scales, will be needed. Some might need to be domain-specific. Numerous testbeds are under development but need significant improvement—for example, to be more usable and to provide data and other tools to support experiments. Testbed integration will be needed for large, diverse experiments.

Authentication and attestation mechanisms to establish trust and justify suspicion should be extended and fully utilized. Authentication of individuals to each other and to machines, and of machines to individuals and to other machines, is required to establish trust, especially in new environments where mobility is the norm. Trustworthy identity combined with privacy-protecting mechanisms is a prerequisite for security policy enforcement and for mechanisms such as network admission control.

Human-capital development is urgently needed. National competitions in secure-system engineering should be inaugurated to attract new talent and integrate academic, industry, and government efforts. National security research institutes with academic, private, and government players should be created and operated in an unclassified context. Research-funding processes should be revamped to encourage long-term, focused engagement in crucial areas. In addition, funding should increase for areas that will create a cyber workforce of researchers, system developers, and system administrators for commercial and government-critical systems.

We also need to initiate research in key technology areas. Participants identified six candidate areas. The first is practical techniques and tools for composable architectures, to support safe system design, extension, and evaluation. The second is transparent security mechanisms,

to enable rather than interfere with work. The third is active automated forensics, to hold attackers accountable. The fifth is self-healing and dynamic security, to raise the bar for attackers. The sixth is system security benchmarking and assessment, to develop quantifiable metrics.

Conclusions

Moving beyond the first workshop's rather broad, informative discussions, this workshop determined that we must focus on protecting the most critical information infrastructure elements from strategic cyber-initiated damage. The workshop methodology was an effective way to identify objectives and means to those ends. Connecting a bottom-up "moves" approach with a top-down "end-state" approach broadened thinking, yet quickly brought key strategic moves into focus in a way that let participants identify necessary new technology that could realistically meet critical applications' needs.

Participants deemed that developing and sustaining human capital (smart, well-trained people) is essential to all important cybersecurity endeavors in the near and long terms. In addition, a strategy for understanding and influencing commercial markets is a prerequisite for any move. Industry should be actively developing strategies for securing the most critical infrastructures. Pragmatic solutions and incentives are needed, and

often cite only recent work, was that we should be careful not to undervalue ideas as "old" just because they have been previously identified and discussed. Many of these ideas have never been tried in earnest and need translation to the current context. The participants concluded that many key ideas, some old and some new, are ready for incorporation into design and evaluation methodologies.

Recommendations for Next Steps

To continue the workshop's momentum, participants recommended several follow-up activities. They felt that, although the workshop had made progress, much more work was required in terms of fundamental analysis and in creating stakeholder consensus.

A significant recommendation was to continue the analysis started in the workshop by creating a small team of highly experienced people with security engineering, research, and operational backgrounds. These experts would try to identify the most critical moves and plan actions for them over the near, medium, and long term.

The participants identified several moves that had no corresponding end-states. These moves' importance and relevance must be determined. In addition, the possible quantification of end-states must be explored to prioritize R&D investment.

On the workshop's last day, a group of industry leaders at-

Over the next ten years transform the cyber infrastructure to be resistant to attack so that critical national interests are protected from catastrophic damage and society can confidently adopt new technological advances.

key vendors can ensure those solutions' practicality.

An interesting conclusion, in these days when research papers

tended briefings on the workshop results and critiqued its progress and recommendations. The ensuing discussion made it clear that

a national cybersecurity effort will much more likely succeed if a cadre of key technical leaders from industry is engaged early and often in the process to determine ways to effect change consistent with commercial-market mechanisms and behavior.

Ultimately, the attendees recommended a series of follow-up workshops over the next year with domain (for example, power, banking, and telecommunications), technology, and industry experts. The workshops would aim to extend the cybersecurity plan. The participants also recommended that government agencies responsible for security R&D and the critical cyber-dependent sectors become more involved, initially through the US government's overarching cybersecurity coordination offices.

The Industry Workshop

This workshop included representatives from the leading US cybersecurity commissions and working groups chartered by various government organizations, and a balance of industry leaders and academic researchers. It provided a forum for members of the computing and communication industry to contribute input to multiple planning and strategic efforts and to define actionable plans to relay back to their organizations.

The workshop took a focused, pragmatic view toward the previous workshop's recommendations and aimed to leverage partnership opportunities between government, academia, and industry. This workshop's results and recommendations are consistent in many ways with those of reports such as *Securing Cyberspace for the 44th Presidency*¹ and other commissions and panels. This is no surprise; key participants in these groups overlap. However, the workshop's recommendations reflect the participation

of key industry leaders and the metalevel recommendations of an economics-driven strategy.

The workshop's key recommendations fell into two broad areas. The first was strategic policy. The US must create a national strategy, national response and recovery plans, and command authorities analogous to US nuclear-defense strategies during the height of the Cold War. Such an effort will require national-level cyber-infrastructure measurement, economics-based consequence analysis, and vigorous programs for the nation's academic leaders and researchers. The Cold War analogy also applies to the proposed academic and research initiatives.

The second area was technical mechanisms: researching, developing, and implementing technological approaches for protecting information systems. These protection approaches must stay ahead of attack technology that will attempt to exploit and corrupt those systems. Specifically, we must develop techniques for modeling security, mechanisms to establish trust relationships, attribution methods, techniques for supply chain assurance, and techniques for quantifying problems and impacts. Support for such development should include proactive security management, an operational environment based on mutual suspicion, and security technology as a key component of commercial products.

The NCDI mapped this workshop's recommendations to the end-states defined at the previous workshop. So, the Opening-Moves Workshop and Industry Workshop resulted in a desired vision (for end-states in a secure world) and actionable recommendations that are under review by key government, industry, and academic participants. These combined results provide an initial road map toward the NCDI vision.

Much work remains to develop a sound national investment plan to avert a catastrophic attack on critical infrastructures. The NCDI continues to organize workshops in areas needing new thinking. We're pleased that cybersecurity has reached the radar screen of national leadership in the US Congress and the White House. At the same time, a public mandate for a major investment in radically improving the US cyber-defense posture still appears missing. Cybersecurity professionals need to make a compelling case for such investment and continue to develop the plans for that investment. □

Reference

1. *Securing Cyberspace for the 44th Presidency*, Center for Strategic and Int'l Studies, 2008; http://csis.org/files/media/csis/pubs/081208_securingcyberspace_44.pdf.

O. Sami Saydjari is the founder and president of the Cyber Defense Agency. He's also the head of Professionals for Cyber Defense, a nonprofit strategic-policy advisory group. Contact him at ssaydjari@cyberdefenseagency.com.

Cynthia E. Irvine is a professor in the Naval Postgraduate School's Department of Computer Science and the director of the school's Center for Information Systems Security Studies and Research. Contact her at irvine@nps.edu.

cn Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.

Interested in writing for this department? Please contact editors O. Sami Saydjari (ssaydjari@cyberdefenseagency.com) and Vijay Varadharajan, (vijay@ics.mq.edu.au).