BY O. SAMI SAYDJARI

CYBER DEFENSE: ART TO SCIENCE

Seeking the knowledge and means to more methodically detect, defend against, and better understand attacks on networked computer resources.

ILLUSTRATION BY ROBERT NEUBECKER

Imagine that you lead an organization under cyber attack on your critical information systems. What questions are you likely to ask?

Am I under attack; what is its nature and origin?
What are the attackers doing; what might they do next?
How does it affect my mission?
What defenses do I have that will be effective against this attack?
What can I do about it; what are my options?
How do I choose the best option?
How do I prevent such attacks in the future?

Unfortunately, today we must often answer, "We don't know and we have no way of knowing." Informally, it is in being able to answer these basic questions that we find the meaning of the term cyber defense.

More formally, we can define cyber defense from its component words. Cyber, short for cyberspace, refers to both networked infrastructure (computers, routers, hubs, switches, and firewalls) and the information assets (critical data on which an organization depends to carry out its mission). Defense is the act of making safe from attack. Therefore, cyber defense refers to an active process of dependably making critical function safe from attack.

Elements of Cyber Defense

Defense in cyberspace is as complex as traditional warfare—it has the same key elements, corresponding to the basic questions listed at the beginning of this article: sensors and exploitation; situation awareness; defensive mechanism; command and control; strategy and tactics; and science and engineering. Simply put, one needs the knowledge and means to defend oneself, detect and understand attacks, and make good decisions about defense configuration. Each of the six elements are discussed in more detail here.

Cyber sensors and exploitation are the "eyes" of the system; they determine the attack capability, plans, and actions of an adversary—the essential first step to any dynamic defense. A primitive form of determining adversary actions is what we today call intrusion detection. To succeed we must acknowledge that attacks will sometimes succeed and adversaries will get inside the system. To assume otherwise is foolish.

Cyber situation awareness is a process that transforms sensed data into a decision aid by interpreting mission consequences and the context of other activity. For example, situation awareness might tell us that attack A will disable the organization's logistics function for three days and that the attack is pandemic and is thus not targeting our organization specifically.

Cyber defensive mechanism is technology to counter threats. Historically, cyber defense has its roots in this element, with cryptography countering intercepted secret messages, virus scanners countering viruses, and firewalls countering hacker exploitations. Although this element is an important building block, professionals must extend their understanding beyond the cyber defense mechanism to see the bigger picture.

Cyber command and control is the process of making and executing decisions—orchestrating defensive systems, based on input from the situation awareness element. Command decision making requires an understanding of options based on the situation, and the means to evaluate them quickly [12]. Control requires a system to communicate the decisions and execute them reliably throughout the system.

Cyber strategies and tactics is knowledge of what constitutes a good decision in terms of initial defensive policies and configurations as well as changes needed during operations because of attack situations. Ideally, such knowledge is based on a wealth of historical experiences, but we prefer not to sustain the damages required to gain real cyber battlefield experience. As a substitute, we must begin developing strategies and tactics and testing them experimentally in models of real systems with mock adversaries.

Cyber science and engineering is the foundation yielding an understanding of design, composition, building, and maintenance of effective defense systems. Currently, this foundation is dangerously weak to the extent that it exists at all.

Dynamic Defense Is Imperative

Static preventive techniques, while important, are inadequate. In the design of trustworthy cyber defense systems, there is a three-way trade-off among security, performance, and functionality. The security dimension itself has at least three components: confidentiality, data integrity, and availability. One cannot statically optimize all dimensions with respect to all attacks. For example, although spreading many copies of data around a system can hinder denial-of-service attacks, it exacerbates the confidentiality problem by creating more targets of opportunity for the attacker. At the higher level, security functions often degrade both performance and functionality. One would rather not have to incur these costs unless under attack, just as soldiers do not put on chemical suits unless there is a known threat of chemical attack on the battlefield.

We need to create systems that make explicit trade-offs within this space both at design time and at operation time—dynamically moving within the trade-off space depending on the situation. We also, therefore, need systems capable of quickly ascertaining the situation so the correct trade-offs can be made.

The Art of War—Strategy and Tactics

Cyber attacks are becoming sophisticated; attackers routinely use attack design toolkits, apply stealth techniques, and target an increasing spectrum of protocols and applications. Cyber attackers are learning to actively evade countermeasures. Soon they will develop sophisticated tactics and will evolve toward strategic campaigns using multi-pronged attacks against strategic objectives. Moreover, attackers have the advantage because they can carefully plan and choose the best time and the weakest points at which to attack. Creating defenses capable of thwarting such attacks will take years; we cannot afford to wait until we see cyber attack methods evolve to this level.

At the same time, defensive mechanisms are proliferating and becoming increasingly complex—exceeding our ability to understand how best to configure

each mechanism and the aggregate of mechanisms.

To effectively manage all the defensive elements, one needs strategy and tactics. Because we have little history in cyberspace, we must look to analogy. We can apply analogies from the battlefield [9]. For example, the battlefield concept of forcing an adversary into disadvantageous terrain has a cyberspace analogue of arranging one's defensive architecture to force adversaries into the "sweet spots" of their intrusion detection algorithms. The battlefield concept of deception has the cyberspace analogue of creating false cyber targets (also known as honey pots) and misleading configurations.

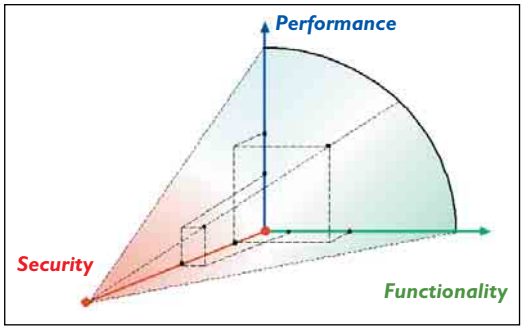

Figure 1. Dynamic design and operating trade-off space.

Similarly, one may borrow from the realm of strategic game playing. The "game" of war is extraordinarily complex because of the great variety of moves, changing rules, and changing capabilities. Yet, some general principles apply, especially as a human decision aid [3]. Determining the right strategic decisions is best performed by creative well-informed humans. This makes cyber defense a matter of art, supported by science, not a matter for total automation. Therefore, we should focus on automating the mundane tasks and providing decision aids to qualified humans for the strategic decision making.

To develop strategy and tactics we accumulate hypotheses based on analogy, and then validate them. We can gain experience through simulation on accurate models of our critical systems interacting with human decision makers. We must engage in many scientific experiments within these models. Adversaries must be accurately modeled using our best red teams. Our strategy and tactics—our cyber defense playbook—need to be validated in such simulations to yield the knowledge to defend our critical cyberspace from sophisticated attack. We must learn how to defend against how real attackers will attack.

Although there is much to be learned from physical war strategy and tactics, there are other areas where the differences are big enough to require a completely new way of thinking about strategy and tactics in cyberspace. As a word of caution, consider some of the key differences. Physical space is three-dimensional; cyberspace is hyper-dimensional, making maneuvering complex. Physical weapon effects are predictable and constrained by physics; cyber weaponry is difficult to predict, often having non-linear damaging effects. Physical attacks occur at human-perceptible speeds; some cyber attacks may aggregate too slowly to be perceived, while many others could occur in milliseconds, making all of them outside the realm of possible human reaction times. Physical attacks often have clear manifestations; cyber attacks can be difficult to detect, making damage assessment problematic.

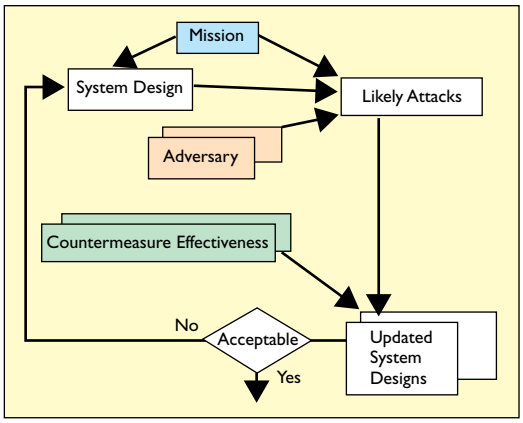

Figure 2. Cyber defense system design models.

Science and Technology Deficits

To achieve a viable cyber defense capability, we need many advances in both science and technology. Here are a few.

We must learn how to create trustworthy systems from untrustworthy components [1]. Trustworthy systems are the building blocks of good cyber defense. The need to create them from untrustworthy components arises from two sources: the vulnerable computer systems that consumers habitually choose and basic system engineering limitations. Trustworthiness, like reliability, is not just in the components, but in the "glue" that holds the components together. Therefore, we need to understand how to achieve trustworthiness through architectures. Without this, we will be building castles in the sand. Some viable approaches have been identified [5] and should be pursued with vigor.

Intrusion detection needs to get a whole lot better. It is inadequate to employ a detect-respond paradigm. Recent attacks such as Slammer and Code-Red are just too fast for today's systems, which are based on detecting signatures of previously detected attacks. Experimental schemes to identify attacks based on detecting anomalies deviating from "normal" activity

have unacceptably high false-alarm rates and relatively poor coverage of the attack space [2]. Further, these schemes often use data from sensors originally designed for auditing security-relevant events, not for detecting attacks. Viable detection requires ground-up analysis of how attacks manifest, custom design of sensors that exploit these manifestations, proper classification algorithms for categorizing relevant events, and detection algorithms that measurably [11] cover the entire attack space.

Intrusion response must be developed so that actions are timely and effective. We need some degree of autonomic response for attacks that are too fast for the human decision cycle. We also must develop decision aids for enumerating situation-dependent courses of action, as well as means to evaluate those courses of action. Using human anatomy as an analogy, we need both the autonomic and the central nervous system, and they must work together to create a systematic defense.

Defending against distributed denial-of-service attacks is needed to ensure availability. Cutting off the many attack paths available to attackers often cuts off the very availability that one is trying to preserve. Further, traditional security and reliability remedies often worsen the problem. Solutions will almost certainly require that quality of service capabilities are added to the Internet.

Countering life-cycle attacks is essential to trustworthiness. If adversaries can infiltrate software development teams and insert malicious code into systems while the software is developed, they can subvert whatever trust that was established. For modern software, the development process and the resulting code are very complex, therefore making preventing and detecting subversion extraordinarily difficult. Nonetheless, we must develop techniques to detect and eradicate malicious code embedded in our code, or find ways to architecturally neutralize it.

Scientific experimental computer science is needed to make real progress. Much of the knowledge in cyber defense today has the status of hypothesis rather than fact. Some have significant evidence in their favor, yet they are still hypotheses. We need experimental methods, based on solid metrics, isolating single variables at a time, to convert these hypotheses into knowledge. Only this will create the firm research foundation needed to enable a sound research track.

Controlled information sharing is needed more than ever. Computer security has its roots in the requirement for Multilevel Security (MLS) processing. The need continues for controlled sharing among groups of differing trust relationships. Further, the need is rapidly growing as collaboration becomes the norm in accomplishing organization goals.

On a final note, even if we make the required scientific and technological advances, we still must find ways to better integrate technology results into mainstream products and systems. Industry's and government's track record of employing useful technology results from research has been poor so far. For example, solutions to defend the vulnerable Domain Naming Service and the Border Gateway Protocol have been available for several years now, but have yet to be incorporated into the network infrastructure.

Creating a Systems Engineering Discipline

Today, the process of designing a well-defended system is matter of black art. One simply hopes designers were adequately knowledgeable about the range of relevant attacks and that they did an adequate job of defending against those attacks. We must evolve toward a systems engineering discipline, which urgently requires several elements.

What gets measured gets done. Without adequate metrics to assess the performance of cyber defense systems, progress is impossible to judge. Some primitive metrics have been proposed [10], but much more work remains to be done.

We need a spectrum of system models and an engineering framework analogous to the CAD/CAM framework used by hardware engineers. The community needs adequate threat models, adversary models [7], mission models, and countermeasure effectiveness models. Each type of model will require tremendous energy to produce, yet little effort is under way in these arenas.

Finally, a methodology to quantitatively trade off design factors and achieve a specified system result is needed. A vision for such a framework should be established and it should be realized with dispatch.

Achieving a National Cyber Defense Capability

So far, I've described cyber defense in the abstract, the principles of which apply at all scales. Here, I examine and discuss cyber defense at the macro scale of defending the national critical information infrastructure. To understand what suffices as a defense, one needs to understand vulnerabilities and the consequences of failure. That the threat is serious has been established [6]. If the reader has any doubts as to the gravity of the problem, consider the major damage done by accidental failures of the telephone system, the power grid,

and banking. Any failures that can happen by accident can likely be induced by an attacker, with significantly more damage potential [4]. The degree of that risk is hotly debated among leading professionals, but, unfortunately, opinions are founded almost entirely on speculation, as scientific study of this question has yet to be seriously undertaken. Such a study is a matter of utmost urgency to our government because the nature and scope of a national cyber defense capability must be grounded in a full comprehension of the susceptibility of our systems.

Engineering Cyber Defense—Manhattan Project. The National Strategy to Defend Cyberspace (released in February 2003) implies that market forces will be adequate to defend at the national scale. This is unlikely to be true [8]. By the same reasoning, market forces would be adequate to defend people, buildings, and towns against kinetic war. We can wish for that to be true. We can observe that we need a certain level of strength to survive normal stresses of life and that this strength provides some measure of defense against trivial attacks. The notion that such protection levels would withstand a nation-state attack is obviously absurd. Why then would this be different in cyberspace? Only a national scientific study will tell us for sure, but intuition tells us that an engineering capability is almost certainly needed.

How might a national cyber defense capability be engineered? What is clear is that this is a very difficult problem. Some of the requisite technology does not yet exist. Yet, all indications are that such a capability is urgent. It therefore seems reasonable to dedicate the finest scientific and engineering minds toward a concerted effort to develop the capability within three years. We need a project in the style of the Manhattan Project with the requisite national priority, resource levels, and structure. Anything less will likely stumble. History has shown us that a distributed lower-priority investment has failed to create the needed capability despite more than three decades of research and development.

Sound National Cyber Defense Policy. Sound national cyber defense policy depends on understanding the problem, its solutions, and the nature of cyber conflicts from economic, behavioral, and political perspectives. Unfortunately, our understanding is currently quite limited. In such a circumstance, a sound policy would be to place high priority and urgency on gaining a deep understanding of all these areas.

Furthermore, the U.S. would benefit by acknowledging that its survival and function now depend on its information infrastructure. Sound policy would reduce the rate at which that dependency is increasing and find ways to minimize the risk as an interim measure.

Finally, I note that achieving a national cyber defense capability has the potential of infringing on citizen privacy in the effort to detect malicious network-based activity. The U.S. national policy must be to avoid abrogating the very rights it is intending to protect.

Conclusion

Cyber defense poses serious technical and policy challenges for the U.S. Much work lies ahead in creating a stronger scientific foundation, the required technology for national-scale cyber defense, and an engineering discipline to provide the means. Policy should follow scientifically based knowledge and understanding; gaining that understanding should now be the primary objective. C

References

1. Committee on Information Systems Trustworthiness, National Research Council. *Trust in Cyberspace*. National Academy Press, Washington, D.C., 1999.
2. Haines, J., Ryder, D., Tinnel, L., and Taylor, S. Validation of sensor alert correlators. *IEEE Security and Privacy 1* (Jan./Feb. 2003), 45–56.
3. Hamilton, S.N., Miller, W.L., Ott, A., and Saydjari, O.S. The role of game theory in information warfare. In *Proceedings of the The Fourth Information Survivability Workshop*, Vancouver, B.C., Canada, March 2002.
4. Letter to President Bush, February 27, 2002; www.uspcd.org/letter.html.
5. Neumann, P. *Principled Assuredly Trustworthy Composable Architectures*. Draft Final Report (Oct. 2003); www.csl.sri.com/users/neumann/chats4.pdf.
6. President's Commission on Critical Infrastructure Protection. *Critical Foundations: Protecting America's Infrastructure*. Washington, D.C., 1997; www.ciao.gov/resource/pccip/PCCIP_Report.pdf.
7. Salter, C., Saydjari, O., Schneier, B., and Wallner, J. Toward a secure system engineering methodology. In *Proceedings of New Security Paradigms Workshop* (Sept. 1998), ACM Press, New York, 1998.
8. Saydjari, O.S. Defending cyberspace. *IEEE Computer 35* (Dec. 2002), 125.
9. Saydjari, O., Tinnel, L., and Farrell, D. Cyberwar strategy and tactics: An analysis of cyber goals, strategies, tactics, and techniques. In *Proceedings of the 2002 IEEE Workshop on Information Assurance*, June 2002, U.S. Military Academy, West Point, NY.
10. Schudel, G. and Wood, B. Adversary work factor as a metric for information assurance. In *Proceedings of New Security Paradigms Workshop* (Sept. 2000), ACM Press, New York, 2001.
11. Tan, K.M. and Maxion, R.A. Determining the operational limits of an anomaly-based intrusion detector. *IEEE Journal on Selected Areas in Communications, Special Issue on Design and Analysis Techniques for Security Assurance 21* (Jan. 2003), 96–110.
12. Tinnel, L., Saydjari, O., and Haines, J. *An Integrated Cyber Panel System*. Supplement to DARPA Information Survivability Conference and Exposition, April 2003, Crystal City, VA.

O. Sami Saydjari (ssaydjari@CyberDefenseAgency.com) is the CEO of Cyber Defense Agency, LLC (www.CyberDefenseAgency.com) and chairman of the Professionals for Cyber Defense (www.uspcd.org).