

A PROACTIVE HOLISTIC APPROACH TO STRATEGIC CYBER DEFENSE

Bradley J. Wood
O. Sami Saydjari
Victoria Stavridou PhD.

SRI International
Cyber Defense Research Center, Systems Development Laboratory

ABSTRACT

We need to change our approach to cyber defense if we are to succeed. We must deeply understand our adversaries, develop effective defensive strategies that will stand the test of time and evolution, and create a new discipline to make this happen. Here is how we start.

INTRODUCTION

The United States and many of its allies are heavily dependent on information technology to insure the safety, security, and prosperity of their people. Military and civilian authorities are seeking “information superiority” to achieve a competitive advantage in both military and economic environments.

The result is the rapid and aggressive introduction of commercial information technology into almost every part of our daily lives. This technology has become pervasive, affecting all economic sectors. The trend toward introducing new technology is also increasing with the passage of time. The phrase “Internet time” is often used to describe the rapid ever-increasing pace of information technology development, simply because traditional economic time scales are no longer relevant.

Information technology has led to many other changes in our society. It is credited with improving quality of life, winning the Cold War, and fostering the longest peacetime economic expansion on record.

Unfortunately, this same technology is inherently vulnerable to a variety of forces, both natural and malevolent. It is arguable that our entire society is becoming dependent on technology that is fragile and easily manipulated by our potential enemies. If this is true (and there is a growing body of data that supports this position), then it is reasonable to conclude that the same technology that is credited for advancing our society does so at great risk.

This potential risk presents the research community with some unique opportunities and challenges:

- How can systems be developed that are inherently robust and survivable?
- What are the economic incentives to improve the survivability of existing critical information systems?
- What are the real vulnerabilities of deployed systems?
- How do we quantify the amount of protection that a given technique might offer?
- What mechanisms provide the best protection?
- What are the inter-dependencies between the nation’s critical information systems?
- What distinguishes a “critical infrastructure” from others?

This potential risk also presents our enemies (both foreign and domestic) with some unique opportunities of their own:

- How can the nation’s critical infrastructures be manipulated?
- What can be accomplished by manipulating these systems?
- What attacks provide the most impact for the least cost?

Therefore, **the race is on** to see who leverages their opportunities first – the US or its enemies.

CURRENT STATE

The authors are prepared to offer data that supports each of these assertions:

- US political, economic, and military systems are already vulnerable due to the pervasive use of commercial information technology.
- Some of the nation’s most important infrastructures (energy, transportation, & communications) are among their most vulnerable.

- These systems are vulnerable to a variety of threats from accidents, naïve adversaries, and sophisticated adversaries.
- These systems can be manipulated to influence decision-makers, the economy, and the will of the general population.
- The tools and techniques needed to disrupt these systems are widely available and can be demonstrated today.
- Decision-makers lack the ability to collect and/or analyze information on threats to the infrastructure.
- Decision-makers lack the ability to detect and monitor the performance of the nation's critical infrastructures and information systems.
- Decision-makers do not have the tools, training, or even the awareness to manage these risks.
- Adversaries enjoy an asymmetric advantage over defenders when attacking an information system.
- The prevalence of commercial information technology is introducing vulnerabilities faster than defenders can discover, evaluate, and mitigate them.

Even if we assume that some of these assertions are false, the situation is still dire.

Currently, defenders are struggling to mitigate the current known threats, rather than redesign systems to protect against known and potentially unknown threats. Some suggest that defenders are falling further and further behind in their ability to protect their systems.

SOME ANALOGIES

What is needed is a new approach to protecting information systems that give defenders the asymmetric advantage over their adversaries.

Several analogies seem appropriate:

- The Japanese Imperial Navy mounted a surprise attack against US naval forces on December 7, 1941 in Pearl Harbor, Hawaii. Analysts predicted this attack. Intelligence reports indicated that the Japanese Imperial Navy was mounting an offensive. Radar warned that enemy aircraft were approaching. Still, Pearl Harbor was arguably one of the most costly Allied failures of the war. Will it take a cyber-equivalent of Pearl Harbor to galvanize support for protecting the nation's critical information systems? What must we do to avoid an electronic Pearl Harbor?

- The atom bomb gave the US an asymmetric advantage over the Axis forces in World War II. Cyber-defenders need this type of asymmetric advantage over their potential adversaries.
- The US was surprised by the launch of Sputnik in 1957 (check date). Only then did decision-makers pursue the *Space Race* a national imperative. It is not clear what type of event is needed to galvanize the research and development community toward a focused concerted effort to protect critical information systems.

A DIFFERENT APPROACH

The Cyber Defense Research Center (CDRC) at SRI International is developing a new approach to defending critical strategic information systems. This approach is based these principles:

- Information Assurance and Survivability are best achieved through an integrated multidisciplinary effort across a broad range of technologies and academic disciplines.
- New, innovative processes are needed to stimulate research and to promote development of research information assurance and survivability technologies. Current processes insure that innovators and developers are always playing catch-up to the adversaries.
- Information Assurance is in a trade-off with other critical properties such as system functionality and performance. We need to be able to intelligently adjust this trade-off during system operation to offer up the best defense. Static systems will become ineffective.
- It is vital that the community have a thorough understanding of the potential adversaries, their capabilities, and tactics.
- Any successful solution must be scalable to address the strategic pervasive nature of the nation's modern critical infrastructures. We must learn how to defend in depth as well as in breadth.
- We must think about attack strategy and defensive counter-strategies as an evolution in time and project forward several moves ahead, as in chess playing, to find the most effective next move, whether that move be in system design, operation, or even research itself.

The CDRC will attempt to orchestrate a broad range of activities into a program of integrated information assurance and survivability research services. These services are built around the following concept:

See – Decision-makers need the ability to comprehend what is happening to their systems, especially when they are global.

Act – Timely, appropriate, and coordinated actions are required to mitigate threats to critical systems.

Build – Designers need the tools to develop inherently survivable information systems, especially when they are large and complex.

Share – Operators need the ability to share information as needed among appropriate parties without putting that information at risk.

SEE

To act, you have to first be able to *see* the adversary. Today, systems can detect local known exploits. In the future, we seek to detect sophisticated novel attacks on a national scale. To create a cyber situation understanding capability to see what is happening within systems, the center will adapt and extend advanced tools and techniques currently used in kinetic warfare as a starting point for cyber warfare defense. We expect that many of the design principles and much of the framework will be reusable. We will start with work on a kinetic situation understanding prototype that we will populate with cyber-specific domain knowledge. The goal is to see how far such tools can take us toward the envisioned capability. Once we judge the difference, we intend to plot a research course to achieve the goal capabilities. We expect to use Emerald technology as event input and combine that with technology from SRI's Artificial Intelligence Center to fuse real-world knowledge (such as news stories) to make sense out of the implications of the unfolding situation. In addition to fusion technology, we expect to do research in mission modeling to understand how an organization mission depends on the computing infrastructure services so the effect of attacks can be assessed with respect to the more meaningful mission function. We also expect, under this heading, to create an Indications and Warning capability based on the creation of implicit attack models that are tracked with respect to ongoing events. We intend to use these models to help design and drive a sophisticated sensor grid including a capability to tune and task those sensors for the most relevant cyber events.

ACT

Today, to respond to attack, operators must make on-the-fly judgements about the best action with little context. They have to implement their decisions manually by reconfiguring each individual relevant component (like, for example, blocking specific ports on firewalls, or changing session cryptographic keys on an Secure Sock-

ets Layer (SSL) connection). In the future, we seek to create a decision support system to help quickly develop and evaluate potential courses of actions, a command execution system that allows automated orchestrated response, and a control subsystem that determines if the commands applied had the desired effect. We intend to create this sort of capability at both the tactical and strategic level. The tactical system capability will be based in the application of control theory to cyber defense. Critical elements will include goal-state specification, the creation of "linear" impulse functions, system state projection (requiring a sensor grid within the defended network), and some form of comparison function between the state projection and the goal state which decides on the appropriate impulse function. The strategic capability will be based more in command and control planning techniques using artificial intelligence technology. Under this activity, we expect to sponsor war-games between red (attacker) and blue (defender) forces to develop general-purpose strategy and tactics suited to situations with particular characteristics. The results are what we call the high-value cyber defense play-book.

BUILD

Today, trustworthy system design is a black art that is done through exhaustion; one tries to counter as many vulnerabilities as possible until available resources are exhausted (similar to bug testing). In the future, we want to enable the design of systems with engineered assurance properties using tools analogous to Computer Automated Design (CAD) tools used by hardware engineers today. To create an effective Security Engineer's CAD system, we must initiate two critical and deeply related thrusts: analysis and design. In the analysis thrust, we expect to create the world's finest red team. The red team will be focused on research improving: next-generation research systems, the red-teaming process itself with tools and techniques, and experiments to learn effective defensive tactics and strategies. See below for more details. To create a security co-designer workbench in support of the "design" thrust, we expect to quickly initiate work on vulnerability modeling and counter-measure effectiveness modeling. Such models will allow designers to understand the comprehensive set of attacks against a putative system and guide them toward the countermeasures that are most effective against the most significant attacks.

SHARE

Today, there is tremendous pressure to share information between inter-company systems for the sake of speed and efficiency. Still, because of a lack of trust in technology, the amount of such sharing is limited to well

below what it would be if we could share with higher confidence. Today, we have all-or-nothing sharing. There is no good way to specify the domains of sharing and keep the transactions to those domains. In the future, we seek to create tailored on-the-fly private collaborative cyberspaces. To do this we must create powerful specification languages for policies, a means to negotiate sharing policies on-the-fly, and a means to verifiably (to all connected parties) demonstrate that the constraints of all parties involved in the sharing are satisfied.

RED TEAMS

As part of the “Design” initiative, the CDRC will make extensive use of Red Teams or model adversaries. The purpose of a red team is to provide a credible model of a realistic threat or adversary so that systems can be evaluated against that threat. Red teams are one of the only qualitative metrics in today’s system technology discipline, thus it plays an essential role. The overall goal of any red team is to improve system defenses. Some of the activities that are envisioned include:

- Evolve “red teaming” from a haphazard activity to a professional practice (much like engineering) and then eventually into a precise science of modeling cyber adversaries.
- Discover effective defensive strategies and tactics against increasingly sophisticated offensive strategies and tactics.
- Develop, exercise, and refine tools to support the execution of these effective strategy and tactics.
- Create a synthetic wargaming environment to support the discovery of these strategy and tactics.
- Develop a process for performing *active assessments* – a process for evaluating a systems strengths and vulnerabilities against a specific threat or adversary.
- Develop refined processes for experimenting with and engaging red teams to evaluate relevant survivability metrics.
- Develop processes for integrating red teams into the design process so that more vulnerabilities can be mitigated through a system’s design process.
- Develop the tools, techniques, and processes for delivering the benefits of a full-blown red team at a fraction of their current cost.
- Develop tools, techniques, and metrics for evaluating the performance of red teams and other model adversaries.

SUMMARY

If the community stays on its present course, the situation will not improve. With these key technology initiatives, we expect the Cyber Defense Research Center to help lead the way into activities fully engaged in rapidly addressing the urgent and critical cyber defense problems of today against the most sophisticated class of adversary. We urge the community to follow.

ABOUT THE AUTHORS

Bradley J. Wood is the leader of SRI’s Research Red Team under the Cyber Defense Research Center. Brad was one of the pioneers of applying the scientific method to the discovery of effective defense principles using red teaming as an aid to high-consequence system design. Brad came to SRI from Sandia National Laboratories where he headed the Information Design Assurance Red Team.

O. Sami Saydjari has been an information assurance researcher since 1984. He has held research positions at the National Security Agency, the Defense Advanced Research Projects Agency, and is the creator and leader of the newly formed Cyber Defense Research Center at SRI.

Victoria Stavridou specializes in research in software engineering, system design and information security. She is the director of SRI’s System Design Lab and the Cambridge Computer Science Research Centre. She earned her Ph.D. from the University of Manchester, UK. Prior to joining SRI, she was a Reader (Associate Professor) in Computer Science at the University of London, UK.