Challenges of Engineering Cybersecurity: a Government Perspective

Tomas Vagoun

National Coordination Office for Federal Networking and IT R&D Program

The US Government is a principal source of funding for basic research in cybersecurity. In that role, the Government is in a position to direct research against fundamental issues in cybersecurity and toward novel and game-changing solutions. Among the federal strategic cybersecurity research themes, Moving Target Defense and Science of Security are great examples of engineering- and science-based efforts to significantly improve the security of our IT systems.

CALL FOR GAME-CHANGING CYBERSECURITY RESEARCH

The Nation's security, economic progress, and social fabric are now inseparably dependent on cyberspace. Yet, our digital infrastructure and its foundations are not secure. Cybersecurity vulnerabilities are exploited by criminals for illicit financial gains; by state-sponsored mercenaries to compromise our national security interests; and by terrorist groups to potentially cause large-scale disruptions in the national critical infrastructures.

The status quo is unacceptable. Recognizing this problem, the Federal Government has been a champion of high-risk, high-payoff cybersecurity research. This strategy has been refined through a series of steps, most notably with the release of the "Trustworthy Cyberspace: Strategic Plan for the Federal Cybersecurity Research and Development Program" (NSTC 2011). This strategic plan has directed federal agencies and challenged the research community to pursue game-changing advances in cybersecurity.

MOVING TARGET DEFENSE

In the current environment, attackers win by taking advantage of the relatively static nature of our systems. Adversaries can plan at their leisure, relatively safe in the assumption that our key IT assets will look the same for a long time. They can map out our likely responses and stockpile a set of exploits that escalate in sophistication as we deploy better defenses. They can afford to invest significant resources in their attacks because they expect to persist in our systems for a long time and reuse the attacks across many targets. To reverse this asymmetry, we need to decrease the predictability of our systems and reduce the return on investment for developing and executing attacks. By making the cyber terrain appear chaotic to the adversary, we force the attacker to do reconnaissance and launch exploits anew for every desired penetration—ideally, the attacker enjoys no amortization of development costs.

The federal cybersecurity R&D community has proposed the development of such capabilities under the rubric of Moving Target Defense (MTD), calling for the development of technologies such as: non-persistent execution environments; randomized execution of code; randomized network and host identities; randomizing compilers; dynamic address spaces; and automated patch synthesis and installation.

There are many natural systems that are far more complex than our cyber systems but are none-the-less extremely robust, resilient, and effective. The biological immune system functions remarkably well in distributed, complex, and ever-changing environments, even when subject to a continuous barrage of attacks. Immune systems exhibit a wealth of interesting mechanisms that could be the inspiration for new methods relevant to Moving Target Defense objectives, such as distributed processing, pathogenic pattern recognition, multi-layered protection, decentralized control, diversity, and signaling. Designing and developing computing systems that implement such capabilities could bring about game-changing advances in cybersecurity.

DARPA CRASH PROGRAM

Announced in 2010 and ending in 2015, DARPA's Clean-Slate Design of Resilient, Adaptive, Secure Hosts (CRASH) Program is an example of rethinking how computing systems could be better secured, taking inspiration from immune systems. The objective of the program was to design systems that can adapt and continue rendering useful services after a successful attack, learn from previous attacks, and repair themselves after a successful attack.

The program took a multi-pronged approach combining efforts that looked at hardware, operating systems, programming languages, and theorems. Hardware was designed to enforce operating rules by tagging every individual piece of data with its type, size, and ownership to enforce access and use restrictions on data at the hardware level.

New programming languages were developed that are explicit about information flows and access control rights. These languages allow programmers to state exactly what rules apply to each module of code. The operating systems enforce these rules dynamically when the program runs.

Similarly, a new type of operating system has been developed, architected around a large number of cooperative but mutually independent modules. Each module is designed with a specific purpose and the lowest level of access privileges needed. The modules are also designed to be suspicious of each other, checking one another's results to make sure they conform to the rules and policies that govern them. This creates a system where more than one component of the operating system would have to be compromised for an attacker to succeed.

When these self-monitoring systems detect a violation, they invoke built-in system services that attempt to diagnose the problem, using replay and reasoning techniques to isolate and characterize the problem; recover from the immediate problem by having multiple redundant methods to achieve any given goal; synthesize filters that can detect the same type of attack in the future and prevent it from succeeding; and automatically generate a patch to fix the underlying vulnerability.

DARPA CRASH program successfully demonstrated that it is possible to develop significantly more secure computing systems that incorporate game-changing ideas that address core deficiencies of today's cyberspace:

Cybersecurity Problem	Biological Approach	DARPA CRASH
Systems are easily penetrated	Innate immunity • Fast-reacting defenses to known pathogens	New hardware and OS that eliminate common technical vulnerabilities. Examples: CHERI (Capability Hardware Enhanced RISC Instructions): hardware-supported, in-process memory protection and sandboxing (Watson et al. 2015) TESLA (Temporally Enforced Security Logic Assertions): compiler-generated runtime instrumentation for continuous validation of security properties (Anderson et al. 2014).
Cleanup and repair is slow, unpredictable, and costly	Adaptive immunity • Slower reacting defenses to unknown pathogens • Learning and adaptation	Adaptive software that determines causes of vulnerabilities and dynamically repairs flaws. Example: GenProg: genetic programming for automated software repairs (Goues et al. 2012).
Computing homogeneity • Large pool of targets, large ROI for attackers • No enterprise-wide survivability	Diversity • Sustains population survival	Techniques that increase entropy, make systems unique, and raise work factor for attackers: instruction set randomization, address space randomization, functional redundancy. Example: Advanced Adaptive Application (A3) Environment (Pal et al. 2014).

SCIENCE OF SECURITY

Prioritized by the federal cybersecurity R&D strategy and supported by research funding from a number of federal agencies, Moving Target Defense has become an active area of research and development. At least 40 moving target techniques have been proposed, at all levels of a computing system, from hardware, operating system, applications, network, to system of systems (for examples see MIT Lincoln Laboratory report, Okhravi et al. 2013).

While the techniques propose innovative approaches to increasing agility, diversity, and redundancy of computing systems, and hence increase attackers' workload and decrease attacks' ROI, MTD techniques are subject to the same limitation as other techniques: we don't know how to systematically assess the efficacy of security techniques, how to measure security benefits, how to compare different techniques, or how to provably determine the security characteristics of the techniques. MTD techniques can make the systems appear chaotic and unpredictable to attackers, however, they do so at the cost of increased complexity. How do we assess whether the benefits outweigh the costs? Some approaches have been proposed, for example, in incorporating MTD into formal security models such as the Hierarchical Attack Representation Model (HARM) (Hong and Kim 2015). Nevertheless, it remains to be seen whether such approaches provide the foundations necessary for reasoning about MTD.

Our inability to assess the strengths and weaknesses of security measures, MTD or otherwise, in a systematic, measurable, and repeatable manner, points to a fundamental weakness. We do not have the foundation to ground the development of secure systems in a rigorous and scientific approach that would facilitate the discovery of laws, hypothesis testing, repeatable experiments, standardized metrics, and common terminology.

The lack of scientific foundations is one of the critical problems and barriers to achieving effective and sustained improvements in cyber security. As a result, nurturing the development of a science of security is also one of the key objectives of the federal cybersecurity R&D strategy.

The most focused science-of-security research initiative funded by the Federal Government is the set of Science of Security Lablets, funded by the NSA. Initiated in 2012, four universities (Carnegie Mellon University, University of Illinois at Urbana-Champaign, NC State University, University of Maryland) were selected to lead research and education projects specifically aimed at investigating scientific foundations of cyber security. Initially, the projects target five areas of interest: resilient architectures, scalability and composability, secure collaboration, metrics, and human behavior.

The growing emphasis on the science of security is strengthening foundations of security across many areas, including MTD. Efforts to reason about MTD techniques from a theoretical basis are growing, including, for example, assessing MTD techniques as changes in a system's entropy (Zhuang et al. 2014).

SUMMARY

We are far too dependent on cyber infrastructure to hope that incremental enhancements will bring about substantial security improvements. In the absence of market-driven solutions, the Federal Government has initiated a series of high-risk/high-payoff R&D programs with the focus on game-changing advances in security. The Government's strategy includes the development of Moving Target Defense techniques and the development of the field of science of security, both showing promising results.

REFERENCES

Anderson J, Watson R, Chisnall D, Gudka K, Davis B, Marinos I. 2014. TESLA: Temporally Enhanced System Logic Assertions. Proceedings of the 2014 European Conference on Computer Systems (EuroSys 2014), Amsterdam, Netherlands, Article No. 19.

Le Goues C, Dewey-Vogt M, Forrest S, Weimer W. 2012. A systematic study of automated program repair: Fixing 55 out of 105 bugs for $8 each. Proceedings of the 2012 International Conference on Software Engineering (ICSE), Zurich, Switzerland, pp. 3-13.

Hong JB, Kim DS. 2015. Assessing the effectiveness of moving target defenses using security models. IEEE Transactions on Dependable and Secure Computing, vol. pp, issue 99.

National Science and Technology Council (NSTC). 2011. Trustworthy Cyberspace: Strategic Plan for the Federal Cybersecurity Research and Development Program. www.nitrd.gov/subcommittee/csia/fed_cybersecurity_rd_strategic_plan_2011.pdf

Okhravi H, Rabe MA, Mayberry TJ, Leonard WG, Hobson TR, Bigelow D, Streilein WW. 2013. Survey of cyber moving targets. Technical Report 1166, ESC-EN-HA-TR-2012-109, MIT Lincoln Laboratory, www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA591804.

Pal P, Schantz R, Paulos A, Benyo B. 2014. Managed execution environment as a Moving-Target Defense infrastructure. IEEE Security & Privacy, vol. 12, no. 2, pp. 51–59.

Watson R, Woodruff J, Neumann PG, Moore SW, Anderson J, Chisnall D, Dave N, Davis B, Gudka K, Laurie B, Murdoch SJ, Norton R, Roe M, Son S, Vadera M. 2015. CHERI: A hybrid capability-system architecture for scalable software compartmentalization. Proceedings of the 2015 IEEE Symposium on Security and Privacy, San Jose, California, pp. 20-37.

Zhuang R, DeLoach SA, Ou X. 2014. Towards a theory of moving target defense. Proceedings of the First ACM Workshop on Moving Target Defense, Scottsdale, Arizona, pp. 31-40.