CDFAI
CANADIAN DEFENCE & FOREIGN AFFAIRS INSTITUTE

Reflections on Re-balancing the Attacker's Asymmetric Advantage

by Michael Locasto

July, 2014

POLICY PAPER

Reflections on Re-Balancing the Attacker's Asymmetric Advantage

By Michael Locasto

Department of Computer Science, University of Calgary

July, 2014

CDFAI
CANADIAN DEFENCE & FOREIGN AFFAIRS INSTITUTE

Prepared for the Canadian Defence & Foreign Affairs Institute
1600, 530 – 8th Avenue S.W., Calgary, AB T2P 3S8
www.cdfai.org

ISBN: 978-1-927573-26-6

Executive Summary

One of the most widely accepted truths in computer security is that the community is playing a game we are destined to lose: attackers will always have quantitatively less work per unit time and more operational leeway than defenders. Computer scientists and policy analysts share a quest to understand and conquer the advantages that accrue to an attacker by dint of this asymmetry. Like all hard truths, it is difficult to accept that we are engaged in an ultimately useless endeavor, and so we persevere in incremental improvements that do nothing to change the underlying structure. Is it really possible to eliminate or significantly reduce the inherent edge that attackers possess? If it is, should we?

This essay considers the nature of asymmetry in cybersecurity. Embedded in this consideration is the need to bridge the language difference between the security community and policy experts. The first reaction to deal with this communication barrier is to link the worlds via analogy. This instinct is vitally wrong. The jargon used in the field of computer security versus that of public policy raises artificial barriers to mutual understanding, diminishing our ability to address the challenges of an asymmetric adversary.

Why is defending computer systems so hard?

Computer system defenders seem to be losing all the time. If the 2011 "summer of Lulz" taught us anything, it is that vast numbers of computer systems require no special expertise to break into and own: the state of modern security defence by most individuals and organizations (including parts of both the Canadian and United State's governments) is abysmal. The recent Heartbleed bug is a catastrophe: a widely used piece of security software is broken in so fundamental a manner as to allow unfettered access to sensitive information stored on all machines that use it.[1] The amount of work to fix this vulnerability is significant as you must patch and redeploy, regenerate keys, regenerate certificates, and have your users change their passwords. In contrast, the amount of effort to take advantage of this vulnerability is simple and amenable to automation.

There are a variety of partial answers to the question above, including market pressures (e.g., make the software work and bolt on security later) and user-interface or user-experience shortcomings, but the underlying reason is the existence of an imbalance between the respective workloads of attackers and defenders. The rational defender must protect against all known (and unknown) vulnerabilities, while the rational attacker need only find one flaw or vulnerability. Accepting this truth leads to the recognition of a fundamental imbalance in effort between attacker and defender. Defenders, if they are not to play a hopeless game, must constantly identify and eliminate flaws in their defence posture. The defender must constantly play the role of all the conceivable possible attackers. But the imaginations of defenders are constrained because they are thinking of how things should work rather than how they can be made to fail.

Some defenders subscribe to a prevailing attitude that attack is always trivial. It is true that some attacks are easy to accomplish. Yet, many attacks worth executing are not trivial; the attacker also has work to do. This is a nuanced position that few defenders understand. Attackers must also identify flaws and then "weaponize" them. The defender must identify and characterize flaws, and we are predisposed to admit that this is hard work. Yet skilled attackers must also undertake this work and maintain a cache of undiscovered vulnerabilities and exploits. This observation leads some authoritative practitioners of offensive security to reject the asymmetric characterization of this game and refer to the attacker's costs as "carrier-class."[2]

So which view is correct? The core of the asymmetric model of cybersecurity is the focus on the distribution of vulnerabilities. Most computers share the same vulnerabilities. This monoculture provides a force multiplier to an attacker.[3] An attacker needs only develop a small set of exploits for a small number of vulnerabilities over a short time scale and then repeatedly reap the rewards of this investment. Defenders, on the other hand, must identify and fix a larger number of vulnerabilities in a smaller timescale. This is the heart of the asymmetric problem: defenders have comparatively more cost in relatively less time. Even if a single vulnerability is eliminated,

[1] CBC. "Heartbleed web security bug: What you need to know" http://www.cbc.ca/news/technology/heartbleed-web-security-bug-what-you-need-to-know-1.2603988

[2] Dave Aitel. "The Three Cyber-War Fallacies" USENIX Security 2011 Invited Talk." Video at https://www.usenix.org/conference/usenix-security-11/three-cyber-war-fallacies

[3] Dan Geer. "Monoculture on the back of the envelope" USENIX ;login magazine. December 2005, Volume 30, Number 6. https://www.usenix.org/publications/login/december-2005-volume-30-number-6/monoculture-back-envelope

the attacker can still rely on the potency of the rest of their collection to attack the rest of the computing population.

IMPLICATIONS

The information security community largely accepts this version of the asymmetric game, yet ignores its implications and does not question the fundamental assumption on which it is based: that computers should retain their current form and function.

We have some opportunities to try and tip the balance. We could increase the cost to the attacker, decrease the defender's cost, decrease the attacker's benefit, or try to do all three at once. If we could control the amount of overlap of vulnerabilities, we might be able to decrease the benefit enough to change the slope of this relationship. While we have not increased the cost of vulnerability identification or changed how difficult it is to weaponize an exploit (which most security mechanisms today seem to focus on), the "rebalancing" approach can make these activities less profitable. In essence, this approach scales the workload necessary to gain the same benefit. It is by doing this, rather than removing constant factors of accumulated workload, that we can potentially even out the balance. For example, Neti et al. examine the issue of how we might create computer systems that are still vulnerable, but are largely vulnerable in disjointed or incompatible ways, thus minimizing the impact of any one attack.[4] They suggest the radical redesign of computers along the lines of evolved organisms.

But that is not what we are doing now, and communicating its nuances to policy makers is a near impossibility given the current disconnect between technical language and policy language.

JARGON AS AN EMBLEM OF DISCONNECTIONS: WHAT DO WE MEAN BY WHAT WE SAY?

The fields of information security and security policy share a common vocabulary, and they both study the art of how to *be* and *defeat* a capable and malicious adversary. Yet, this shared vocabulary comes with different semantics, and this divergence in meaning is a monumental obstacle when discussing the security of modern computing systems and critical information infrastructure.

Modern information systems are large, complex, and fast. Scale and speed kill immature notions. The first notion it kills is that the game of cybersecurity is anything like warfare. That is not to say that it is harsher or gentler, better or worse, but only *different*. Unfortunately, jargon causes a short-circuit in critical thinking. As Dijkstra, the noted Computer Scientist, tells us, "By means of metaphors and analogies we try to link the new to the old, the novel to the familiar. Under sufficiently slow and gradual change, it works reasonably well; in the case of a sharp discontinuity, however, the method breaks down."[5]

Discussing cybersecurity in terms of analogies fails us in many places. One example is the term "cyberweapon." Shown a listing of the code for a computer exploit, many people would simply see a list of strange numbers. In fact, most people cannot easily tell the difference between the

[4] Saran Neti, Anil Somayaji, and Michael Locasto. "Software Diversity: Security, Entropy, and Game Theory". USENIX Hot Topics in Security 2012. https://www.usenix.org/system/files/conference/hotsec12/hotsec12-final25.pdf

[5] Edsger W. Dijkstra "On the Cruelty of Really Teaching Computing Science" http://www.cs.utexas.edu/~EWD/transcriptions/EWD10xx/EWD1036.html

piece of code that issues a calendar reminder and a piece of code that copies your password. This lack of knowledge is not unexpected, even very few Computer Science students really ever gain any facility with understanding information as a single bit of information that can represent many things at once.

This pile of encoded numbers speaks to a fundamental disconnect between the fields of public policy and information security. Such information might form the basis of a software exploit, but who is to say whether this information is a cyberweapon or a cybermunition? Is it like a bullet or a gun? Perhaps neither. Need it be licensed and controlled? Must we busy ourselves with international conventions and treaties that the strong will ignore and the weak subvert?

Thus, this break in meaning leads to the lesson Dijkstra wishes to impart: the purposeful elimination of analogies as tools of explanatory power in this interdisciplinary problem. Sometimes "new" really is radically new and cannot be described by our old thinking. We are quite poor at classifying an array of data values as being either like a bullet or a blunderbuss. We must first seek technical understanding before dictating policy and law.

TAKEAWAY MESSAGES

Cybersecurity is a complex game. Understanding who the players are and the purity of their motivation is a subject of great debate. Even if we agree on the players, knowing their disposition is a costly impossibility because we are moving in an atmosphere of uncertainty and partial observation. There is no simplified version of this game, and attempting to play one leads to incorrect conclusions.

Overcoming this temptation toward simplification takes significant effort. Early on in a recent graduate seminar on "Cyberwar, Cyberterror, and Cyberprotest", the class sought at length to clarify its thinking about the differences between traditional kinetic war and "cyberwar." This discussion also considered what the terms "cyberweapon", "cybermunition", and "cyberspace" mean. There appeared to be the risk that imposing a primarily military perspective on cybersecurity issues could cloud the conversation by obscuring the actual nature of the technology artifacts involved. It seemed to the class participants that the practice of offensive and defensive information security is distinguished from "war" between nation-states by a few characteristics, but chief among them is asymmetry.

In the name of improving security, we can jointly work toward rebalancing the nature of this asymmetric game, but this redistribution merely implies a leveling-out of the workload, *not* an elimination of the workload or an asymmetry in favor of defenders. One possible way to obtain this rebalancing is to abandon what we think computer systems should look like. Computer Scientists must have the courage to embrace this paradigm because it means dealing with computer systems that are radically different than current designs.

One definition of “security” that I offer my students is: “the imposition of one principal's will on another." While admittedly an incomplete expression of the complexities of security, the main value of this definition is its moral neutrality: the principals (i.e., actors) in this drama are both attacker and defender: they each strive to impose their will on the other. In short, there are no good guys and no bad guys. We are *all* a composition of the varied cast of characters routinely used by Computer Scientists to model and define secure communication protocols.

Whether you like or dislike that definition, here is a subversive thought arising from it: maybe we should not seek to eliminate asymmetry after all. Perhaps the status quo is a desirable equilibrium. Perhaps you do not want to eliminate the attacker's advantage because one day, the attacker might be you.

About the Author

Dr. Michael Locasto studies the security of computer systems. More precisely, he tries to understand why it seems difficult to build secure systems and how we can get better at it. He has particular interests in making software defense mechanisms automatic, correct, and efficient. He works on intrusion defense, debugging & software trustworthiness, and innovative approaches to information security education.

Canadian Defence & Foreign Affairs Institute

CDFAI is the only think tank focused on Canada's international engagement in all its forms - diplomacy, the military, aid and trade security. Established in 2001, CDFAI's vision is for Canada to have a respected, influential voice in the international arena based on a comprehensive foreign policy, which expresses our national interests, political and social values, military capabilities, economic strength and willingness to be engaged with action that is timely and credible.

CDFAI was created to address the ongoing discrepancy between what Canadians need to know about Canadian international activities and what they do know. Historically, Canadians tend to think of foreign policy – if they think of it at all – as a matter of trade and markets. They are unaware of the importance of Canada engaging diplomatically, militarily, and with international aid in the ongoing struggle to maintain a world that is friendly to the free flow of goods, services, people and ideas across borders and the spread of human rights. They are largely unaware of the connection between a prosperous and free Canada and a world of globalization and liberal internationalism.

In all its activities CDFAI is a charitable, nonpartisan organization, supported financially by the contributions of foundations, corporations and individuals. Conclusions or opinions expressed in CDFAI publications and programs are those of the authors and speakers and do not necessarily reflect the views of the Institute staff, fellows, directors, advisors, or any individuals or organizations that provide financial support to CDFAI.