LA-UR-15-27889
Approved for public release; distribution is unlimited.

Title: What is the current state of the science of Cyber defense?

Author(s): Hurd, Alan J.

Intended for: Report

Issued: 2015-10-09

What is the current state of the science of Cyber defense?

Alan J. Hurd, Director (Acting) of the National Security Education Center
Current employer: Los Alamos National Laboratory
Other affiliations: American Physical Society, Materials Research Society, ScienceCounts!
September 18, 2015

My overall sense of the cyber defense field is one of an adolescent discipline currently bogged down in a cloud of issues, the most iconic of which is the great diversity of approaches that are being aggregated to form a coherent field. Because my own expertise is complex systems and materials physics research, I have limited direct experience in cyber security sciences except as a user of secure networks and computing resources. However, in producing this report, I have found with certainty that there exists no calculus for cyber risk assessment, mitigation, and response, although some hopeful precepts toward this end are emerging.

The area I found to have the most science opportunity is cyber hygiene. This area includes managing and protecting credentials, knowing when credentials have been stolen, and User Access Behavior (UAB) analytics. It is this last subarea where much progress is being made and probably will be made in the future.

The term "user" in UAB analytics can be usefully generalized beyond user to attacker and author (coder). "Behavior" can be such things as time of access (For an attacker, is there a diurnal cycle matching up to work schedules in some part of the world?), device preference (Was this code written to defend a specific type of printer on a network?), browser usage (Is this user likely to be attached inside Internet Explorer), and thousands of other behaviors adopted by a generalized user.

Using the internet as a "Hubble" tool, behavioral sciences are undergoing a revolution in analytics and theoretical understanding driven by huge data streams coming from mobile devices, security cameras, traffic sensors, satellite imagery, and the like. (See for example the research output of the Santa Fe Institute, http://www.santafe.edu/.) The access behavior of users, attackers, and authors can in principle be characterized in useful ways, increasingly so given the growth of information infrastructure.

Practitioners in cyber defense with whom I spoke characterized the current state of UAB as "naïve" and "zeroth order" compared to what is being done in social sciences. Research in complex networks produces over 200,000 papers per year and growing fast, so an accessible body of science exists for cyber science to exploit in complex networks. Cyber defense has not yet fully converged available analytics in a prioritized manner to find interesting correlations that would lead to cause and effect and protecting assets. Current cyber implementations lack built-in intelligence and bandwidth to monitor many channels and correlate big data.

One area of notable progress is the emergence of "scanners" with growing sophistication. A scanner is an automated tool monitoring a series of analytical channels, such as IP addresses through an email gateway. Small companies are developing application program interfaces (APIs)

that communicate with scanners already deployed in the field to correlate their analytical measures with the goal of early warning. An immediate step is to integrate large (~1000) numbers of diverse scanners. A shotgun approach would not be desirable, however; rather, the community should deliberate which analytics provide the best leverage on investment, optimal insight into the complex networks being probed, and are most amenable to machine learning. Aggregation software is being developed and marketed to pull in assessment information from scanners, total up the risk, and prioritize actions by asset. Here is where one needs a calculus of risk assessment that can be automated.

On the cyber offense side, it should be obvious that scanning software and hardware could be used to form a dynamic attack path.

Forensics for incident response may benefit by UAB analytics. Practitioners report that maturity is building but the tools of cyber forensics are scant and unsophisticated. The goals of forensics are to determine how an attack started then how it spread through the sequential steps of incident response and discovery. Currently this agenda is carried out heuristically with a general lack of understanding. Again, applying principles from behavioral sciences of complex networks, forensics should be poised to leap forward.

In a similar vein, the tradition of "white-hatting" to reveal vulnerabilities in a cyber defense is starting to benefit by new UAB tools. This challenge—white hatting, that is—strives to break into a system for friendly and constructive improvement. The standard technique is to sample statistically for weak points in breaking through a defense. It may be more important to think of nontraditional ways to break in by considering, for example, the topology of the target network. If the nodes of a network are segmented logically, its robustness to attack can be improved. Note that segmentation creates the problem of maintaining a large perimeter defense, which implies that nodes with compact depth may payoff in reducing surface-to-volume.

A good example of a brilliant white-hat exercise of a segmented system was described to me as follows. In a hospital environment, an attack began—as it often does—with a printer entry point. (Printers are notorious vulnerabilities set up without thought toward security features.) The white hatter was able to hop from segment to segment, such as printers, EKG machines, EEG machines, and ultimately to security cameras. Next the attacker established behavioral baselines by watching entrees through a cipher-keyed portal. At an opportune time determined from compiled behaviors, the target person of the exercise was seen by camera to key into that portal allowing the white hat to obtain the target's code from the video channel. In addition to revealing weaknesses in segmentation, this white-hat exercise illustrated the power of UAB analytics.

Other aspects of cyber hygiene besides behavior constitute a strong defense and many of the following considerations are training and education issues.

Insecure coding is getting worse with outsourcing. Even in good designs, problems can arise include missing security patches, not invoking built-in security features, and frequent (even nightly) changes. While many of these issues are not in the domain of cyber science or research, they certainly fall into the category of educational shortcomings. In that sense, computer science

departments around the world bear the onus to improve cyber defense through training for secure coding.

As in many new industries, there is a standardization problem in the science of cyber defense. The formats and channels for exchanging security information have not converged well from the days that every company had its own formats. This problem appears to be well known but there is a cost barrier to shifting gears even though MITRE has reportedly developed workable standards. Federal programs have helped push the cyber industry toward more uniformity.

Security policy as opposed to engineered controls is generally a weak link owing to training and compliance shortcomings. Phishing usually gets through a policy defense, in part because compliance is at odds with convenience. Human Performance Indicators—another set of behavior metrics—provide scientific insights as to how policy fails; therefore HPI may inform policy makers of best-design.

As states have begun to sponsor attacks, advanced persistent threats (APTs) organized by large teams with intent have replaced individual "gray hat" attacks with mischief in mind. One might ask why APTs don't succeed more often. One engineer speculated that perhaps the target field is so rich that we have not reached Nash equilibrium, in the game theory sense, with adversaries. This kind of thinking strikes me as the right way to frame cyber defense science to meet today's needs.

References

"The Big Four—What We Did Wrong in Advanced Persistent Threat Detection?", Virvilis, N. and Gritzalis, D., in 2013 Eighth International Conference on Availability, Reliability and Security (ARES), 2-6 Sept. 2013, pages 248 – 254, Regensburg, Germany

Acknowledgments

In developing the material for this report I discussed the state-of-the-art with John Donahue, a software engineer and PhD student at New Mexico Tech, and Dr. Srinivas Mukkamala, CEO of RiskSense, Inc in Albuquerque, NM.