

Axelrod *continued*

creating safety-critical systems and those involved with security-critical systems. The goal would be for each group to adopt preferred practices, commonly used by the other group. In this way, failure attributes of security-critical systems will be improved and safety-critical systems will be better protected against attacks. As I have previously tried to demonstrate (Axelrod 2009), there is a great need for resiliency of security-critical software. 🗨️

### Acknowledgement

The author wishes to thank Jennifer Bayuk for her assistance with this article and for providing the pharmaceutical example.

### References

- Axelrod, C. W. 2009. "Investing in Software Resiliency." *CrossTalk*, September–October: 20–25.
- Feron, E. 2010. "From Control Systems to Control Software." *IEEE Control Systems Magazine* 30 (6): 51–71.
- Gross, G. 2010. "Stuxnet Changed Cybersecurity." *NetworkWorld* 22 (Nov.): 10.
- Gutgarts, P. G., and A. Termin. 2010. "Security-Critical versus Safety-Critical Software." Paper presented at the IEEE Homeland Security Technology Conference, Waltham, MA (US), Nov.: 507–511.
- Joint Task Force, Transformation Initiative. 2009. "Recommended Security Controls for Federal Information Systems and Organizations." NIST Special Publication 800-53, Revision 3. Gaithersburg, MD (US): US Department of Commerce.
- Jones, T. 2011. "Japanese Worker inside Stricken Reactor Recalls Quake." Reuters, 26 March 2011. <http://in.reuters.com/article/2011/03/25/idINIndia-55893620110325>.
- National Institute of Standards and Technology (US). 2004. "Standards for Security Categorization of Federal Information and Information Systems." FIPS PUB 199. Gaithersburg, MD (US): US Department of Commerce.
- RTCA. 1992. RTCA DO-178B. "Software Considerations in Airborne Systems and Equipment Certification." Washington, DC: RTCA, Inc.
- RTCA. 2011. SC-216. "Aeronautical Systems Security." <http://www.rtca.org/comm/Committee.cfm?id=76>
- Stoneburner, G. 2001. "Underlying Technical Models for Information Technology Security." NIST Special Publication 800-33. Gaithersburg, MD (US): US Dept. of Commerce.
- Stouffer, K, J. Falco, and K. Scarfone. 2008. "Guide to Industrial Control Systems (ICS) Security." NIST Special Publication 800-82, Final Public Draft. Gaithersburg, MD (US): US Dept. of Commerce.
- Van der Zwan, E. 2010. "Security of Industrial Control Systems: What to Look For." *ISACA Journal Online* 4: 1–9.

## How Do We Measure Security?

Fred Cohen, [dr.cohen@all.net](mailto:dr.cohen@all.net)

While I await an answer to this question that makes sense, I thought I would point out how we, as engineers, measure other things. I am an electrical engineer by education, and as a result, I was introduced to fundamentals of certain physical phenomena early in my life. In electrical engineering, we have these positive and negative charges associated with different physical things, and we can count the things and their charges, measure them through instrumentation that we calibrate against standards, and use those measurements to make mathematical models of electrical components. With the modeled components, we can build systems and use mathematics to model, analyze, and design systems. We have simple symbols and models for simple systems under normal conditions, like  $v = ir$  for a voltage ( $v$ ) across a resistor ( $r$ ) with current ( $i$ ), and more complicated models for more complicated systems under less usual conditions, like the field equations applied to understand propagation of waves through a wave guide.

Just about every electrical engineer in the world understands these same things, uses the same symbols and mathematics, carries out the same calculations, and comes to pretty much the same answers for the same systems. There are similar concepts in other fields of engineering, and they are used to build up mathematical models and understandings, which are then applied to the design and evaluation of components and composites that form systems, under various operating conditions.

As an engineer by education with a generally scientific bent on things, I figure those in the security engineering community might try to look at security in this light. Unfortunately, there really is very little, if any, community of this sort in the security space, and even less of a community in the information-security space. And what community there is, still apparently cannot agree to even the simplest of concepts, like a methodology for evaluating passwords, which have been documented in global use for at least several thousand years (Kahn 1967; Cohen 1997, 2009). In one LinkedIn.com forum of computer-security practitioners, there were almost a thousand responses to a question of how long a password should be, and there seem to be many different schools of thought on the issue.

I maintain, without substantial basis other than the history of science and engineering in other areas, that this problem stems from a lack of a scientific understanding of the very basic components of security in the information arena that we can meaningfully count and the lack of a common language for describing and discussing them. Without this basic set of things, I don't think that meaningful progress is likely to be made, and I don't think that we can reasonably proceed to do engineering in the information-security space.

I don't want to be misunderstood in this. I am not saying that people who work in information security don't measure things. To the contrary, they measure lots of things. For example, there is a fad of counting the number of known "vulnerabilities" (whatever those are) in computers, and the number of those vulnerabilities "mitigated" (whatever that is) over a period of time. This is measured by "scanning" for known vulnerabilities using a tool that nobody actually knows how to calibrate (or what it would mean to calibrate it), sold by a commercial vendor who makes money by having different vulner-

ability lists than other vendors and scanning for them better, faster, and cheaper, of which only two can be measured, and only in the simplest of ways.

Indeed there is a whole information-security metrics community that has emerged, and as a community, they seem to have largely agreed that they do not want to address any fundamental challenges like developing a theoretical basis for measurement, identifying basic underlying things that might be worth measuring, or providing any sort of mathematical basis for doing anything with the measurements they provide. (See <http://www.securitymetrics.org/> and their mailing list for details.) They seem to be stuck in the challenge of all mysticism before science emerges. They look at things and make up notions about what might be interesting, use the similarity of appearance or descriptive statistics to claim or imply causality when no mechanism has been proposed and scientifically evaluated, and they introduce casual theories that are not evaluated but become part of the urban legends of the field. Attempts to burst these bubbles are met with a level of hostility and frustration, and any attempt to introducing the notion of starting to look at the underlying principles or theoretical issues is met with scorn. Indeed it is easy to become a pariah in that community by bringing up the underlying fallacies of the approaches too often, and if you mention anything serious about developing a physics for the area, don't count on getting any serious attention from this "research" community.

Of course physical security does not have the same set of problems that information security has in this regard. There are some pretty substantial notions underlying physical security and they are widely accepted in the communities that really care about such things. Information security has largely ignored these things, even though they could be quite helpful. For example, physical security measures progress in an attack graph with time as a metric for a physical protective system. They take measurements of different defense mechanisms (e.g., fences or gravel) relative to specific attack mechanisms (tanks, people with tools) by doing experiments, and form a mathematical model of the overall system by enumerating all of the topological paths from source to target (and back out depending on the issues involved) (Garcia 2001). These are done for different threats identified as part of the threat-assessment activity, which also has a scientific basis in facts and measurements, along with an intelligence component, also based on facts.

A similar approach has been taken in scientific experiments with information protection, producing a measurement of progress in an attack graph with time (Cohen 1999, 2001, 2003). Results of these approaches can be used to make design selections and to assess response methodologies (e.g., determine mean time to repair for a desired availability) and thus approach systems engineering. Unfortunately, such approaches have not found their way into the larger information-security community. Funding to support such scientific development has largely

been lacking, and for some reason, there does not appear to be a lot of support for pursuing such lines of enquiry in the larger community. I don't claim that this is the best, or even a good approach, only that it is an approach that has some of the trappings of a scientific basis for an engineering discipline. Notionally, it is at least a place to start thinking.

Returning for a moment to electrical engineering, I recall the first time when working with digital circuits that I encountered the notion of fan-out and fan-in. Here is an example of a design rule that states, in essence, that any output of a digital circuit within the same family can be connected to  $N$  inputs within the same family, where  $N$  is the fan-out. Instead of having to perform calculations at a detailed level for each time I interconnect two logic gates, I could simply make sure that I stayed within the fan-out (and fan-in) limits, and I was assured that the ultimate implementation would work reliably and within the overall specifications.

The ability to connect one digital circuit to another to form an overall circuit of arbitrary size means that I can design reliably without having to worry about complex interactions ranging across the entire system. Indeed I can connect systems to other systems to form ever-larger systems and, assuming I continue to follow appropriate design rules, build up a system of virtually any size, knowing that it will continue to work properly. I may also have some parameters like clock speed and propagation times to concern myself about as a designer, but those too are known parameters of the overall system within the defined operating range of the technology.

In the information-security arena I do more or less the same thing when, for example, I design a set of web servers to support a domain. I know the performance characteristics of the computers, and can measure them reliably in terms of the number of web pages they can serve over time, the total number of bytes they can serve over time, and so forth. As long as each is scaled so that it can handle the maximum bandwidth at the external interface without service failing, I can guarantee against denial of services from excessive requests. If I need more bandwidth, I can set up multiple web servers, and I can limit bandwidth by limiting the external services rates. I have done this since the mid 1990s, and as a result, despite a wide array of attempts at distributed denial of services attacks against some of my servers, they have never failed from excessive load.

Unfortunately, while some individuals may have such design rules and may have built up limited metrics that they use for internal purposes, publishing results in this arena and making them part of the science and art of information protection seems to be out of vogue for now. If as a community we took this sort of approach, we could build up metrics that could be measured reliably with known instrumentation, create designs with known characteristics at their interfaces, and create design rules and the necessary underlying knowledge to build up an engineering discipline that could create systems of systems with known characteristics that could with measurable

Cohen *continued*

certainly provide defined protective functions, be calibrated, measured for performance and changes in performance, applied repeatedly to build up larger structures, and produce known and defined operating characteristics over a defined range of operating conditions.

Systems engineers who want effective system security must understand that if the designers don't treat security like any other design criteria, they will get designs that fail for security reasons. If a systems engineer fails to take into account the operating temperature of space systems, the pressure requirements of underwater systems, or the radiation environment of nuclear-safety systems, they get systems that fail when put into use in those environments. The same is true of the security environment. If you plan to put a computer on the Internet, it had better be measurably specified and designed for that environment or it will fail. If the operating conditions will change over a range, we had better be able to measure and design for them.

So here we sit. As systems engineers, it would be nice to be able to use the same sorts of notions of design for information security as we use for other sorts of design. It would be nice to be able to have standard units of measurement against which we could test things. It would be nice to be able to develop tools for measurement that could be calibrated against the standards, to have a theoretical basis for developing a mathematics and testing it, and then to be able to build up a systems engineering approach to information security like we do in other engineering fields. But first, we need to be able to make meaningful measurements. 🗨️

## References

- Cohen, F. 1997. "Managing Network Security—Part 10: Change Your Password" *Network Security* 9: 8–11.
- \_\_\_\_\_. 1999. "Simulating Cyber Attacks, Defenses, and Consequences" *Computers and Security* 18 (6): 479–518
- \_\_\_\_\_. 2009. "Passwords Again: Why We Can't Leave Well Enough Alone." *Fred Cohen & Associates Analyst Report and Newsletter*, 1 Nov. 2009. <http://all.net/Analyst/2009-11.pdf>.
- Cohen, F., and D. Koike. 2003. "Leading Attackers through Attack Graphs with Deceptions." *Computers and Security* 22 (5): 402–411.
- Cohen, F., I. Marin, J. Sappington, C. Stewart, and E. Thomas. 2003. "Red Teaming Experiments with Deception Technologies." 2003. <http://all.net/journal/deception/RedTeamingExperiments.pdf>.
- Garcia, M. L. 2008. *The Design and Analysis of Physical Protection Systems*. Boston, MA (US): Butterworth-Heinemann.
- Kahn, D. 1967. *The Codebreakers: The Story of Secret Writing*. New York, NY (US): Macmillan.

# Making People the Center of Systems Security

Jennifer McGovern Narkevicius, [jennifer.narkevicius@incose.org](mailto:jennifer.narkevicius@incose.org)

Every system is developed to enhance the capabilities of people. There is not a single system developed in the world that does not exist solely for the humans it supports. The most successful systems intentionally incorporate the people that comprise them as well as integrating the engineered elements. That is, the inclusion of the considerations of people position the system to actually *work*. Systems security is essential—it protects organizations, individuals, material, and concept. Further, systems security is critically dependent on the users to perform as an essential component. Nevertheless, many ignore the user as part of the design space. If the users of systems are considered at all, it is often as a source of problems rather than as the boundary of system behavior. This absence of humans in system-security design has led to stunning instances of failure because of the unwillingness of designers to identify the contributions of self-organizing security within a system. People are adaptive: they actively work within the constraints of equipment and environment to pursue their activities. However, just because collections of sensors can be integrated into networks generating characteristics (known and learned), this does not generate security—unless the users understand the sensors, the networks, their behavior, and what those mean.

Failure to account for the people in and around a system has led to many mishaps, some quite spectacular. Systems are developed to extend or enhance the capabilities of humans, and to achieve this, systems engineering must represent humans at the same level of detail and accuracy as hardware and software. All systems—manned and unmanned—include people as operators, maintainers, and support personnel, and therefore the designers must take these people into consideration.

The human systems integration (HSI) technical process integrates human considerations within and across all system elements and is an essential enabler to systems engineering practice. HSI merges the human-based technical work with the other engineering domains, allowing engineers to incorporate human capabilities and limitations into their assessment of all the design issues; HSI also ensures that risks are addressed in an integrated manner. To accurately define system requirements, it is essential to understand the inherent capabilities and limitations of users, hardware, and software, defining boundaries for trade spaces and achieving the desired system performance.

The HSI analysis data must be included in global systems engineering processes to define a total requirements set and to establish inclusive trade-space boundaries, ensuring that a system will perform as envisioned in the operational environment when the system comes together. To achieve this precision, HSI tools must be plugged into the systems engineering process throughout the development cycle to deliver the technical information needed for accurate trade-offs and improved design. Primary HSI tools focus on modeling,

