Barriers to Science in Security

Tom Longstaff
The Johns Hopkins University
Applied Physics Laboratory 11100
Johns Hopkins Rd., Laurel, MD
20723
Thomas.Longstaff@jhuapl.edu

David Balenson
The Johns Hopkins University
Applied Physics Laboratory 11100
Johns Hopkins Rd., Laurel, MD
20723
Thomas.Longstaff@jhuapl.edu

Mark Matties
The Johns Hopkins University
Applied Physics Laboratory 11100
Johns Hopkins Rd., Laurel, MD
20723
Thomas.Longstaff@jhuapl.edu

Overview

In the past year, there has been significant interest in promoting the idea of applying scientific principles to information security. The main point made by information security professionals who brief at conferences seems to be that our field of information security is finally mature enough to begin making significant strides towards applying the scientific approach. Audiences everywhere enthusiastically agree and thrash themselves for bypassing science all along, bemoaning the fact that we could be "so much further along" if we only did science. Of course, after the presentation is over, everyone goes back to the methods that have been used throughout our generation to generate prototypes and tools with no regard for the scientific principles involved.

The type of information security[1] projects in scope for this essay are experimental projects that produce a new approach or support/refute a theoretical result. The use of the scientific method in theoretical information security and in computer science more generally is well documented and mature (even if not universally applied). The focus of the "science of security" publications in FY09-10 is in the area of experimentation and applied information security research. Thus our focus here is also in the comparison of experimental information security research that does or does not use a traditional scientific method in the execution of the project and in the publication of the results. The definition of the scientific method we use in this essay is well documented and not further described here.

Finding agreement in the use of the scientific method is practically universal, finding participation in the scientific method is rare. Why? What are the primary barriers to applying the scientific method to information security projects? What are the main differences between the projects that apply the scientific method to experimental information security projects and those that promote software/tool development without applying a traditional scientific approach? In this essay, I explore three main barriers to achieving a more universal application of the scientific method to experimental information security projects. These are:

- Time to publish as a primary driver
- Standard of peer reviews in conferences and journals
- Expectation of a breakthrough in every publication

Although these drivers are evident in many academic publications, it must be noted that much of the work in computer science, and more importantly, information security does not concern the development of the body of scientific knowledge, but in getting a job done. This is closely aligned to computer engineering or software engineering, both of which are often associated with computer science departments. Many practicing computer scientists work in the area of information security by producing innovative tools and techniques to solve specific technical problems in information security. Many of these practitioners have a computer science degree, but have never been formally trained in the application of scientific method, nor do they need to be to have successful careers in information security. The overarching goal in this area of information security is to get the job done in terms of writing a program to accomplish a task, rather than on exploring the and testing the range of possibilities (experimenting) and implementing a better solution. Practical, working systems that can be quickly implemented tend to prevail. This follows the old IETF mantra of "rough consensus and running code." (see http://en.wikipedia.org/wiki/Rough_consensus).

In this sense, the Science in "Computer Science" is a misnomer – many CS graduates are never formally trained in the scientific method and its use in experimental information security. Some CS curricula teach basics in terms of computational logic, programming languages, data structures, database, artificial intelligence, etc., but do not teach scientific experimentation. Many other academic curricula, like Math, and even English, often develop students who ultimately work as programmers, developers, or researchers, but they also lack formal education in scientific method. The curricula that do teach scientific methods, such as Psychology, Biology, Physics, etc. lead to few people who work on information security. However, as more of these professionals enter the field, the call for a scientific approach becomes increasingly urgent.

[1] We use the term *information security* to clarify that the types of projects in scope address the confidentiality, integrity, or availability of information assets. While it is common to use the term *cyber security* to address perhaps a wider set of topics, the definition of *cyber security* is not as well defined or accepted, and thus is more likely to cause confusion over the types of projects included herein.

ACSAC '10 Dec. 6-10, 2010, Austin, Texas USA

Time to Publish as a Primary Driver

The application of the scientific method to experimental information security projects usually takes significantly more time than is available for the development of a demonstration/prototype tool. A carefully conceived experiment requires planning around a well-formed hypothesis, assuring that the tests against the hypothesis are sufficient to potentially refute the hypothesis. In the likely event that the experiment will support the hypothesis, the domain of the test environment must be sufficient to build an argument that the hypothesis holds in a significantly extensive context. This frequently means many runs of the experiment over a wide variety of input variations to assure the relationship between the domain and range of the system under test (SUT) is as predicted by the hypothesis.

In contrast, many experimental tests take a developed prototype or demonstration system and provides a narrow set of performance characteristics. Since there is no hypothesis to test, there is no possibility of refuting a hypothesis. All that is generated is a series of observations of the SUT. The tests can be performed in a narrow set of domain variables since the test is designed to show performance in the environment for which the SUT was designed. Since no failure is possible in this situation, the tests need not be extensive to lead to results that may be published.

Even when a rigorous scientific test is designed, the pressure to publish quickly may lead to an inadequate exploration through extensive and multiple trials. There is a tendency to test a very limited set of functionality or a small number of parameters. This approach supports the hypothesis, but only for a limited environment. These tests answer specific questions such as testing an implementation X in environment Y and it's ability to detect Z. Variations X' in alternative environments Y' may be limited. The full operating range or characteristics of our technology may not be included in the rush to publish.

The publication of a well-designed experiment must follow a rigorous structure that will allow readers of the publication to fully repeat the experiment. This includes the domain (data and input settings), full description of the SUT (including any implementations), and the architecture of the test environment. This implies that this data was carefully captured during the experiment, which again takes a carefully planned experimental methodology. When simply executing performance tests of a prototype/demonstration system, the standard is not to capture the experiment in full detail, but to instead describe the performance of the prototype/demonstration. The publication is not designed to allow others to re-create the experiment but instead to motivate the use of the prototype in their environment. Since there is not carefully described domain description for the test, the result (range) of the prototype in a new environment cannot be accurately predicted.

A well-defined experiment has a much more powerful predictive value, but given that it takes a much longer time to achieve, there is significant pressure on researchers to publish a higher volume of results more quickly than running a series of experiments. Since the metric for most academics in the area of information security is number of publications rather than quality of experimental results, rewards are gained by minimizing a scientific approach and putting out as many publications on new prototypes/demonstrations as possible. Since we get what we incentivize, time to publish becomes a primary driver for choosing prototyping over science.

Standard of Peer Reviews in Conferences and Journals

Of course, rapidly producing many publication submissions based on prototypes and demonstrations would be irrelevant if the selection criteria in conferences and journals favored science over demonstration.

In many natural and social science journals and conferences, a submission must demonstrate the use of good science principles in order to be considered for publication. In scientific areas such as Physics, Chemistry, Psychology, and many others, the entire culture is focused on the critical evaluation of scientific evidence. A reviewer in these disciplines has an enormous responsibility to represent the critical review of the entire readership. Her primary responsibility is to discredit the potential publication before it can be discredited by the readership. A well respected journal or conference gains a reputation for the inclusion of only a small subset of submissions that cannot be discredited, so thus must be published to allow another researcher to reproduce the results or possibly refute the hypothesis while hopefully proposing an alternative.

In cultures such as the natural and social sciences described above, critical reviewers are trained throughout their career to evaluate submissions for scientific rigor. New ideas are not simply given credence for being clever, but must be supported with scientific evidence. Only then can the new idea be incorporated into the scientific body of knowledge and used to make further predictions.

In sharp contrast to this culture, the majority of the information security reviewers consider the technological implementation of new ideas to be of high worth. A description of a new tool that implements a feature that has not yet been conceived is of great interest to most of the reviewing community. A critical review of this type of submission usually focuses on the quality of the description itself, and of any duplication the tool might have with previous tools that have been created (often to assure there is a reference to this prior work). In this case, experimental design is neither desired nor appreciated in the submission, and may be excluded for a reduced page count.

Expectation of a Breakthrough in Every Publication

If you accept the previous two points (time to publish and standard of peer review) as driving the culture of scientific discourse in information security, a natural expectation for short-term R&D is to create a novel new system and publish the result. These new systems are designed to solve particular problems (such as intrusion detection or secure computing), but the approach to solving the problem is to use insight to create a novel solution that attempts to solve the problem at large. The "breakthrough" solutions are shown to be effective in a lab environment or small set of enterprise environments and described as a prototype demonstration of the novel concept.

While there is absolutely nothing wrong with the generation of technology based on novel concepts (this is how many companies succeed), this is not a scientific approach to solving problems in information security. Using a scientific approach would create reusable knowledge or explore causal relationships rather than

focus on the apparatus used to gain these results. By equating the process of "scientific discovery" with technology innovation, we create an expectation that scientific publications should always contain a breakthrough technology as a core benefit. This expectation leads to a reduced number of accepted publications that show incremental progress in the understanding of how information security actually works, and instead promotes publications that fully describe a technology breakthrough.

Conclusions and Way Forward

It is certainly possible that in this field, the traditional scientific approach is not commercially viable from a product development standpoint. It can easily be argued that given the rapid pace of technological advance, we should be promoting innovative technological solutions over scientific investigation. We do have mature and rigorous scientific investigation in computer science more generally and in information security from a theoretical and cryptographic perspective. While we don't often use these results to drive innovation, there are specific instances where we have used results from theoretical computer security to drive a security product.

If this is the case, why the clamor for scientific method in experimental information security? Given the advances in other experimental sciences, the hope is that we can begin to develop lines of information security products that are incrementally better as time goes on, not just by adding features to an implementation, but by understanding the underlying causality of information security and addressing the problem at its most fundamental level. Applying the scientific method to our experiments will enable a more purposeful approach to discovering the exact conditions under which our innovations can be expected to operate, providing much greater utility in our future products.

If this is a goal to be at least partially achieved, the three barriers to adoption described in this article must be addressed. Each of these poses a significant challenge to the field as they address the culture of our process, which one can argue has successfully produced commercially successful products. Yet the basic problem of information security remains. Could we begin to eliminate these problems through the application of experimental science in information security? If we do not create at least a small sub-culture that applies scientific method to experimental information security, we may never know. If we do create such a sub-culture that embraces experimental science in information security, it might be best to treat this delicate new community as a "skunk-works" from the main body of information security R&D. This would involve creating a series of publication venues that use reviewers from this new community, create expectations that will appeal largely to this community (and not to the information security community at large), and which creates a body of knowledge that is formed outside of the mainstream of information security R&D. The success or failure of this community will pivot on its ability to solve fundamental questions in information security in a way that cannot be ignored by the mainstream.

It possible that the current climate of our funding agencies in the US and EU are disposed to fund the creation of this community given a clear definition and leadership in its formation. For members of this conference that both have a deep understanding and appreciation of experimental science and for future program managers that might fund such an approach, it is time to come together to produce the "grand experiment" of the creation of a sub-community of information security that rejects ad-hoc solutions in favor of scientific evidence that increase our understanding of information security.