

SANDIA REPORT

SAND2012-0640

Unlimited Release

Printed January 2012

What Then Do We Do About Computer Security?

Michael J. Berg, Philip L. Campbell (Editor), Christopher E. Davis, Jackson R. Mayo, Roger A. Suppona, Gregory D. Wyss

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited



Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



SAND2012-0640
Unlimited Release
Printed January 2012

What Then Do We Do About Computer Security?

Michael J. Berg, Philip L. Campbell (Editor)
Network Systems Survivability and Assurance
Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185-0672

Christopher E. Davis
Effects-Based Studies

Jackson R. Mayo
Scalable Modeling & Analysis Systems

Roger A. Suppona
Cyber Enterprise Security

Gregory D. Wyss
Security Systems Analysis

Abstract

This report presents the answers that an informal and unfunded group at SNL provided for questions concerning computer security posed by Jim Gosler, Sandia Fellow (00002). The primary purpose of this report is to record our current answers; hopefully those answers will turn out to be answers indeed. The group was formed in November 2010.

(This page intentionally left blank.)

CONTENTS

1	Introduction.....	7
2	Our Answers	8
2.1	Berg	8
2.2	Campbell	10
2.3	Davis	15
2.4	Mayo	21
2.5	Suppona.....	23
2.6	Wyss.....	25
3	Summary.....	29
4	Distribution.....	32

TABLES

Table 1 Overarching Directives.....	29
Table 2 Single Directives	29

1 INTRODUCTION

In November 2010 Jim Gosler, Sandia Fellow, asked several of us several pointed questions about computer security metrics. Never mind that some of the best minds in the field have been trying to crack this nut without success for decades. Jim asked Campbell to lead an informal and unfunded group to answer the questions. With time Jim invited several more Sandians to join in.

We met a number of times both with Jim and without him. At Jim's direction we contacted a number of people outside Sandia who Jim thought could help. For example, we interacted with IBM's T.J. Watson Research Center and held a one-day, videoconference workshop with them on the questions.

Over the year Jim added more questions to the list upon occasion and upon occasion we provided our then-current answers in the form of short, informal documents, usually about one page each. As we now complete a year on this work we have gathered our now-current answers and present them in this report. The following are Jim's collected questions:

- #1. I have a million dollars; how should I spend it to maximize my computer security?
- #2. I am a program manager for computer security. How do I identify the proposals that will increase my computer security?
- #3. I am a program manager for computer security. When a funded proposal completes how do I determine how much security I got for my money?
- #4. Why is this problem so hard?
- #5. How will our civilization's response to this problem play out?
- #6. How do I address deterrence in this world?

The rest of this report is organized as follows. The next section presents the answers for each of us. The subsequent (and last) section presents a summary.

2 OUR ANSWERS

This section presents our now-current answers for each of the questions.

2.1 Berg

I've re-ordered the questions (but kept the original numbers) to string together an argument.

#4. Why is this problem so hard?

A large part of why the problem is so hard is because of the definitions used and the resulting unsolvable problems. Common definitions of security and insecurity can be summarized as:

- Secure: A system is secure if that system has no vulnerabilities.
- Insecure: A system is insecure if that system has one or more vulnerabilities.

Under these definitions, claiming a system is secure requires proving the non-existence of vulnerabilities, which is an unsolvable problem. Unconditional definitions like these are closely associated with how many people try to treat security as a one-time-investment.

Both [Torgerson] and [Engle] present proofs that systems cannot be proven to be secure. However, each paper points to more constrained definitions of security that may be useful:

- Secure in Practice: A system is secure in practice if that system has no vulnerabilities that are both known to and exploitable by an adversary [Torgerson].
- Real-time Secure: A system is real-time secure if and only if every configuration in its computation history is authorized by the security policy [Engle].

Both of these definitions serve to remind system owners that perceived security can become insecurity at any time and that eternal vigilance is required.

#1. I have a million dollars. How should I spend it to maximize my computer security?

#2. I am a program manager for computer security. How do I identify the proposals that will increase my computer security?

#3. I am a program manager for computer security. When a funded proposal completes, how do I determine how much security I got for my money?

These are all really asking the same question, just at different points in time. Questions 1 and 2 are about how to allocate funds and select projects before work is done. Question 3 is about how to determine success or failure after the work is done. Program managers need to be thinking about question 3 when they are planning and trying to answer questions 1 and 2. The answers to questions 1 and 2 should provide the metrics and the criteria for success or failure needed to answer question 3.

Based on the answer to question 4, [the question] “how much security did I get for my money?” is an unanswerable question. However, “how many known classes of vulnerabilities did I eliminate for my money?” can potentially be answered.

Also try to avoid the conceptual trap of *adding security*. Vulnerabilities can often be eliminated by *removing* some position/role/component instead of by *adding* yet another security product. *Increasing security* is a simple mental translation to *adding a security product*. *Reducing vulnerability* is an easier mental translation to *removing vulnerable parts of the system*.

#5. How will our civilization’s response to this problem play out?

I don’t know. We could each describe how we would like the future to play out. However, I have no idea how things will actually play out given the complex interactions between organizational inertia, external actors, budget constraints, and public opinion.

#6. How do I address deterrence in this world?

Deterrence is all about making the total cost of performing some action greater than the benefit of performing some action.

Retaliation is one way of drastically increasing the cost and is what many people focus on. However, retaliation generally requires attribution, which is another open area of research.

Another way of increasing the cost is to make the process of targeting our systems more expensive due to combinations of human factors, organizational structures, new technologies, diversity and randomization, misinformation, etc. If the costs sufficient to deter are an inherent part of the system architecture, attribution may not be necessary to achieve the desired goal.

References

[Engle] S. Engle, S. Whalen, and M. Bishop, *Modeling Computer Insecurity*, Tech. Rep. CSE-2008-14, UC Davis (2008).

[Torgerson] M. Torgerson, *Security Metrics for Communication Systems*, in Proceedings of the 12th International Command and Control Research and Technology Symposium (DoD CCRP, 2007).

2.2 Campbell

#1. I have a million dollars; how should I spend it to maximize my computer security?

This is the wrong question. Organizations do not allocate resources and then wonder how to apply them. Rather, they allocate resources based on their bottom line. Organizations are under the same evolutionary selection pressure that every biological species is under; the bottom line indicates fitness; so allocating based on the bottom line is a matter of survival.^{1,2} The correct question is, “Which investments in computer security are cost effective for my organization?”

This question calls for making a decision under uncertainty. The tool that organizations use to reduce uncertainty enough to enable a good decision is decision analysis ([Clemen & Reilly], [Hubbard-1]).³

Decision analysis will employ subject matter experts (SMEs) who will probably use guidance such as available from ISACA (www.isaca.org), the Center for Internet Security (www.cisecurity.org), and NIST (for example, SP 800-53 [NIST]), all of which codify hindsight and foresight. These SMEs will probably also use the lists of particular vulnerabilities & particular countermeasures for particular systems as provided by the Defense Information Systems Agency (www.disa.mil) and others. The SMEs will probably also consider technology, such as commercially-available A1 systems that provide “verified protection” [TCSEC], as opposed to the customary C2 systems.^{4,5}

#2. I am a program manager for computer security. How do I identify the proposals that will increase my computer security?

(See answer #1.)

¹ Government agencies and non-profit organizations are under the same pressure but they must rely on a more ambiguous indicator.

² Sowell notes, for example, that “More than once, Intel poured such huge sums of money into the development of improved chips as to risk the financial survival of the company itself. But the alternative was to allow itself to be overtaken by rivals, which would have been an even bigger risk to Intel's survival” ([Sowell], page 111).

³ A good decision does not necessarily imply a good outcome, unfortunately, but that is the best we can do when faced with a decision under uncertainty.

⁴ Bell considers C2 systems as “no stronger than that required to keep order between cooperative colleagues” ([Bell], page 133).

⁵ Bell and Schell & Reed provide summaries of the history of computer security ([Bell-Addendum], [Schell & Reed]).

#3. I am a program manager for computer security. When a funded proposal completes how do I determine how much security I got for my money?

What would you observe if there were to be a change in computer security? The answer to that question defines computer security for your organization and describes how it can be measured. There is a growing set of generally accepted practices that everyone should follow but this set cannot describe all of the practices a given organization should adopt because that is dependent on the objectives, strategies, and tactics of the organization itself (see answer #1).⁶

#4. Why is this problem so hard?

This problem is “so hard” for two reasons. The first reason is that computer security has followed mainstream reductionist science, e.g., closing this hole stops that attack, but reductionism is insufficient for systems, particularly large systems such as the computers of today and our civilization, which now includes computers. As auditors like to say, computer security consists of people, process, and technology.⁷

The second reason that this problem is so hard is that computer security has ignored decision analysis. We are limited in what we can know in the general case about computer programs⁸ and this limitation constrains what we can know about computer security. However, decision analysis enables us to address uncertainty in computer security the way the rest of the world addresses uncertainty in every other problem. At that point this “so hard” problem reduces to the kind of “hard” problem that organizations face every day (see answer #1).

#5. How will our civilization's response to this problem play out?

What we know today as “capitalism,” Hayek calls the “extended order of human cooperation” ([Hayek], page 6)⁹ and it has been evolving for millennia. Even though we live in it and owe our existence to it, it is easily misunderstood. For example, Hayek noted that Aristotle thought that this

⁶ Granted, a negative (e.g., the number of computer incidents did *not* happen) cannot be measured but computer security does not have to be defined as a negative. In fact, if computer security – or anything else for that matter – makes any difference in the world, then there must be some effect that that difference makes. That effect can be observed and this leads to a measurement [Hubbard-2].

⁷ Kline argues that “sociotechnical” systems such as information systems are so complex that we can build and maintain them only by using feedback [Kline], what Braybrooke & Lindblom refer to as “disjointed incrementalism” [Braybrooke & Lindblom].

⁸ Consider, for example, Rice's Theorem which “dashes all hopes of algorithmically testing input-output behavior of arbitrary programs” ([Machtey & Young], page 102).

⁹ Hayek suggests that a better description of Adam Smith's “invisible hand” is an “invisible or unsurveyable pattern” ([Hayek], page 14).

order among men could extend only so far as the voice of a herald could reach (Ethics, IX, x), and that a state numbering a hundred thousand people was thus impossible. Yet what Aristotle thought impossible had already happened by the time he wrote these words. ([Hayek], page 11)

In the past, this order has continued to extend as it successfully addressed problems similar to the problem of computer security that we face today. Consider, for example, the introduction of commercial paper, the adoption of Arabic numerals, the use of credit cards, the rise of the global supply chain, and the countless other changes that have always increased efficiency – that is why they were adopted – and have usually increased risk as well. Or consider that as the order has extended it has enabled the expansion of cities, which in turn has exacerbated the specter of uncontrolled fire.

...let us attempt to take you back to earlier times, fifty, a hundred, two hundred years ago, to a time when the fear of fire dominated men's lives, when small mistakes became conflagrations that destroyed whole cities, when fate seemed to control who lived and who died by fire. Try to imagine living in these times, [facing] a problem that seemed insurmountable. ([Whitley & Yonas], page 6)

Solutions were devised that have gradually, over several hundred years, reduced this threat to a level that allows us to co-exist with the threat of fire by applying constant vigilance and investments in fire protection, but without living in constant fear and dread from fire...Fire protection has become a virtually unnoticed constant in our daily lives. (ibid., Abstract)

Presuming that our civilization survives, its response to the problem of computer security will be similar to its response to innumerable other problems, such as uncontrolled fire, though we undoubtedly do not have "several hundred years" to address this one.

As Machiavelli wrote, "Prudent men are in the habit of saying, neither by chance nor without reason, that anyone wishing to see what is to be must consider what has been; all the problems of this world in every era are found in ancient times" ([Machiavelli], page 351).

#6. How do I address deterrence in this world?

Attackers are deterred either by retribution or by the low cost effectiveness of those attacks for those attackers. Retribution requires attribution, which we do not yet have and may never have in the computer world, so we must resort to reducing the cost effectiveness of attacks. Part of this reduction requires continuing to address vulnerabilities (see answer #1). The other part requires resilience. Currently resilience ranges from data backups to hot sites – duplicate systems that can take over on the next clock cycle. This range is good but it is not enough. We need to go beyond electronics: we need graceful degradation all the way to manual. With that extended resilience we can reduce the cost effectiveness of attacks and thereby provide deterrence.

We can learn resilience for computer security by generalizing from organizations that are resilient in other areas, such as under the pressure of lean inventory. For example, consider a solution used by Caterpillar Inc.:

Many companies pool the risks of finished goods inventory by having a single replenishment system with access to their entire inventory in all of their warehouses. Similarly, many retail outlets can direct consumers to a different store of the same retailer, where a wanted item can be found. But what is unique about Caterpillar's system is that it takes such inventory risk pooling a step further by creating a single virtual inventory *involving its customers* – the independent network of dealers – in order to serve the ultimate customer, the equipment owner. The result: “We operate 24 hours a day, every day of the year. We'll ship 99.7 percent of all those items the same day,” Newbanks [the head of inventory research management at Caterpillar] said. ([Sheffi], page 231, emphasis in the original)

Solutions such as Caterpillar's can give us direction on how to address resilience in computer security.

In the meantime, one way to provide a retribution of sorts is to establish private networks with rules of behavior and a large amount of money required from new members up front. If members violate the rules, they lose their membership as well as their money. At some point, as these private networks proliferate, the people who play by the rules will all be in private networks and everyone else will be in each other's empty pockets.

References

[Bell] David Elliott Bell, “Looking Back at the Bell-La Padula Model,” Annual Computer Security Applications Conference (ACSAC), Proceedings, pp. 127-146, December 9-13, 2002.

[Bell-Addendum] David Elliott Bell, “Looking Back: Addendum,” November 27, 2006 (see [Bell]).

[Braybrooke & Lindblom] David Braybrooke, Charles E. Lindblom, A Strategy of Decision: Policy Evaluation as a Social Process (1970), The Free Press.

[Clemen & Reilly] Robert T. Clemen, Terence C. Reilly, Making Hard Decisions: An Introduction to Decision Analysis (1996), Duxbury Press.

[Hakek] F. A. Hayek, The Fatal Conceit: The Errors of Socialism (1988), University of Chicago Press.

[Hubbard-1] Douglas W. Hubbard, The Failure of Risk Management: Why It's Broken and How to Fix It (2009), John Wiley & Sons.

[Hubbard-2] Douglas W. Hubbard, How to Measure Anything: Finding the Value of “Intangibles” in Business, Second Edition (2010), John Wiley & Sons.

[Kline] Stephen Jay Kline, Conceptual Foundations for Multi-Disciplinary Thinking (1995), Stanford University Press.

[Machiavelli] Niccolo Machiavelli, Discourses on Livy, Oxford University Press [1531] 1997. (See also Robert D. Kaplan, Warrior Politics: Why Leadership Demands a Pagan Ethos (2002), Random House.)

[Machtey & Young] Michael Machtey, Paul Young, An Introduction to the General Theory of Algorithms (1978), Elsevier North Holland.

[NIST] Ron Ross, et al., "Recommended Security Controls for Federal Information Systems," NIST Special Publication 800-53. February 2005.

[Schell & Reed] Roger R. Schell, Edwards E. Reed, "Computer Security: a Historical Perspective," pp. 321-341, Encyclopedia of Quantitative Risk Analysis and Assessment (published September 15, 2008).

[Sheffi] Yossi Sheffi, The Resilient Enterprise: Overcoming Vulnerability For Competitive Advantage (2007), The MIT Press.

[Sowell] Thomas Sowell, Basic Economics: A Common Sense Guide to the Economy, Third Edition (2007), Basic Books.

[TCSEC] "Department of Defense Trusted Computer System Evaluation Criteria" (TCSEC), DoD 5200.28-STD, December 26, 1985 (aka Orange Book).

[Whitley & Yonas] John Whitley, Gerold Yonas, "The War on Terrorism and What We Can Learn from our War with Fire," SAND2002-2404, Unlimited Release, Printed July 2002.

2.3 Davis

#1. I have a million dollars; how should I spend it to maximize my computer security?

There is no universal answer to “maximize computer security.” Computer security must bear in mind the assets that are being protected and the impact of loss or compromise of those assets.

Diversify - Defense in Depth

In my opinion, computer security investment must be viewed much like an investment portfolio. You cannot lump all of your investment into a single endeavor and hope that it will maximize your return on investment. Instead, consider the entire cyber system – including the users and administrators, and spread your investment across the system. (Cole, September 08, 2009) (Vacca, May 22, 2009)

Know your system

To be able to best apply an investment, one needs a good sense of what is mission critical within the system. For example, if we consider a banking transactional system like FedWire it is critical that the system be able to complete financial transactions at high speed. I am certain that there are email systems and other secondary systems that are part of such a banking transaction system, but at its core, if the transactions do not clear, the system does not function. As such, protecting the core function of the system is the most important goal. With the ever increasing number of applications and services that a modern enterprise is hosting, it can be tempting to continue to spend resources patching any and every service – however with finite time and resources, one **MUST** focus on the critical mission of your system. (Cole, September 08, 2009)

Understand how your organization defines security

Every organization will have its own definition of “security.” Once you have a solid understanding of your definition of security, you can focus strengthening your implementation of those ideals.

Simplify

Services and applications are promoted and deployed at an alarming rate. While it is critical that IT keep pace with the needs of the workforce, it is also important to balance complexity and the attack surface area that is presented. By reducing the number of services offered, the IT staff can spend more time maintaining those services. (Cole, September 08, 2009)

Roles

Move your enterprise to a role based use model. All users have roles, if they are clearly defined and implemented, you can dramatically increase the security of a system. Initiatives like the Federal Desktop Core Configuration (FDCC) are a move in the right direction, but fall short of

the mark. Finer grained, individual and group based, automated controls are required for role based access to truly succeed.

Multiple and Continuous Forms of Authentication

Current password standards are ridiculous. If one looks at the standard and approaches the problem rationally, it is simple to see that we have created a password requirement that forces us to choose passwords that are very hard for humans to remember (I forgot my password when trying to unlock my laptop TODAY), but trivial for a machine to guess. In passwords, bit entropy is the only metric that matters for security. For usability, easy to remember (without writing down, or keeping a file or email) is what matters. At the intersection of these two constraints are pass-phrases. Long (so hard for a brute force or dictionary attack even if quantum computing becomes feasible), but memorable (so users are less likely to inadvertently compromise them by recording them).




Figure 1 Image from XKCD.com

More widely adopted two or more factor authentication. With multi factor authentication it is harder to share or have stolen authentication tokens. Forms of authentication that are typically considered are "Something you have, something you know, or who you are."

Behavioral biometrics let you continuously authenticate a user. While this proposition may sound rather “big brother-esque,” if a cyber system continuously monitors a user’s behavioral biometrics it can add a layer of resilience to the otherwise static authentication “fence” that users cross. For example, if key-dynamics, acceleration forces, and orientation are measured on a mobile device, we can very accurately and unobtrusively ensure that the authorized person is using the mobile device. If an un-authorized person attempts to use the device, the typing and holding patterns will be very difficult to mimic successfully. Based on the sensitivity of the material the device may hold, automatic actions could include additional challenge questions, temporary lock-out of the device, or even a automated “kill-pill” (a self destruct mechanism built into many portable electronic devices to assist administrators in controlling information leakage) in which the device is securely wiped of all content.

#2. I am a program manager for computer security. How do I identify the proposals that will increase my computer security?

Full Spectrum Support

The proposals that include not only a description of the device/software but also a good description of how that product is configured, deployed, used, maintained, and retired are more likely to increase security.

If the solution addresses the system as a whole it would seem more likely to increase security. Humans are part of the system and cannot be ignored.

#3. I am a program manager for computer security. When a funded proposal completes how do I determine how much security I got for my money?

Measure Mission Specific Quantities

You must have a clear goal before funding any proposal that will “increase your security.” For instance, if the goal is to “Reduce the number of successful phishing attacks,” you should have a clear measure of how many phishing attacks you currently receive and can then measure if there is a decrease in the number of phishing attacks. Some other metrics to consider are included in the table below where “prior” describes the measure that you take before implementing a solution, and “posterior” describes the measure you take after implementation. It is worth noting that while the measures are the same, the prior and posterior have been included to re-enforce consistency in measurement.

Metric	Prior	Posterior
Increase the hour rating of my firewall when under attack from an open source based attacker	Hours the firewall withstood attack by a known red team	Hours the firewall withstood attack by the same red team
Decrease information loss	A statistical measure of the	The same statistical

Metric	Prior	Posterior
	loss size or rate of loss	measure of the loss size or rate of loss
Increase resiliency	Enumeration of failure modes	The same enumeration ¹⁰
Increase simplicity of system	Any number of complexity measures	The same complexity measure
Decrease attack surface	Enumeration of available attack points	Enumeration of available attack points using the same approach

If this question is driving at the **VALUE** returned for the investment, then the task is harder. One must then value the asset (money, time, etc) and determine the value added through some organization approved (or accepted) approach.

#4. Why is this problem so hard?

Measuring cyber security posture is non-trivial. We lack the metrics that cryptography has in its brute force strength approach to security. There are no universal metrics for cyber security at this point. This puts the analyst and engineer at a disadvantage because every situation is a unique problem. Standard analysis and engineering practices benefit from patterns and re-use.

Systems are secure until they are not.

Provable qualities are hard to come by, and lack universality.

#5. How will our civilization's response to this problem play out?

<opens fortune cookie>

"You will face a great challenge"

Joking aside, I think that unless we are clear about our goals, we will spend volumes of money, create even more complex, multifunction devices and systems, and not get much of anywhere. Right now the popular mindset thinks that you can "add security" to systems that are already in use. While this may be possible in some situations, I dare you to ask a bridge engineer to add another lane to a bridge that has been built. For some reason, most people understand the difficulty of adding another lane to a bridge, but think that you can bolt security onto a system.

If however, we are willing to discard large pieces of our cyber-infrastructure and engineer the remaining pieces start to finish with our security goals in mind, I think we will see tremendous improvements in security.

¹⁰ Note this is a very qualitative measure! There are less satisfying quantitative measures.

Some would fault the current innovators for not building security into their systems. We must remember just how young this field is in comparison to other fields. Modern computing has been available (depending on your math) for about 70 years. The parts of computing that were extant before this time all seem to be stronger – cryptography, logic, and math. If we consider “security” in the modern definition, it is a very late addition to the collection of ideas that is computing. We can look back at Marie Curie and think, “How reckless she was to keep radio isotopes in her pocket and desk drawer,” but the challenge she faced was much like the challenge contemporary computer scientists and engineers face – a problem of exploration and of showing that something can be done. In our haste to see utility and financial and social rewards from the advances, we deploy systems that are discovering and inventing new security issues.

Like physical security, this will always be an adversarial “cat and mouse” problem.

#6. How do I address deterrence in this world?

Active Defense

We face a legal dilemma with regards to deterrence. There is currently no international law to explicitly regulate cyber attacks or exploitation. As such, we are left as nations to interpret malicious cyber behavior as either an act of war or as a domestic crime. Both routes have significant flaws as the bodies of international law were written for a different purpose. Interpreting cyber attacks as an act of war places cyber attacks into a category of action considered an armed attack or use of electronic force. While this in of itself may seem benign, the unintended side effect that comes along with this interpretation is that we are limited to passive cyber defense, as an active cyber defense would be considered an act of war¹¹. Interpreting cyber attacks as domestic crime limits prosecution (and as a result deterrence) as many nations are unwilling to prosecute or extradite for cyber actions. (Carr, December 15, 2009)

Politicians, scientists, military, and society must come to an agreement about how to respond to cyber threats and carefully define *jus ad bellum* and *jus in bello* as it is interpreted for cyber actions. *Jus ad bellum* is the body of laws that govern whether entering into war is permissible (i.e., is this a “just war”). *Jus in bello* are the body of laws that govern how war is conducted (i.e., is this war “conducted justly”). (Various)

As a nation and coalition of nations we should pursue the use of active defenses and the laws that govern their use. If a large enough coalition of nations decides to support the use of active defense, it will force sanctuary nations into cooperation with the coalition. No nation wants

¹¹ An active defense is one in which offensive action is taken to stop an ongoing attack. This is in contrast to passive defenses in which the attack is allowed to continue and attempts are made at the defender’s sites to filter or otherwise mitigate the damages.

force used within their borders, and if *jus ad bellum* and *jus in bello* are carefully defined – the proposition of entering war over failure to cooperate can only serve as additional motivation.

References

- Carr, J. (December 15, 2009). *Inside Cyber Warfare*. O'Reilly Media, Inc.
- Cole, E. (September 08, 2009). *Network Security Bible, 2nd Edition*. John Wiley & Sons.
- Vacca, J. (May 22, 2009). *Computer and Information Security Handbook*. Morgan Kaufmann.
- Various. (n.d.). *Jus ad Bellum*. Retrieved 11 30, 2011, from Wikipedia:
http://en.wikipedia.org/wiki/Jus_ad_bellum

2.4 Mayo

#1. I have a million dollars; how should I spend it to maximize my computer security?

The most vital improvement to cybersecurity would be the ability to measure it. Once measurement is possible, improvement is relatively straightforward. This implies that most of the work involved in designing systems with *better* security is the work of designing non-generic systems whose security can be quantified at all. Potential solutions grounded in science involve either (a) reducing and strictly limiting the complexity or, to the contrary, (b) embracing and harnessing it against the attacker. Approach (a) is feasible when the required functionality is simple enough, and is already used for many high-consequence digital systems [1]. Approach (b) is more general but requires additional enabling research [2, 3].

#2. I am a program manager for computer security. How do I identify the proposals that will increase my computer security?

(See answer #1.)

#3. I am a program manager for computer security. When a funded proposal completes how do I determine how much security I got for my money?

Generic cyber systems are inherently unpredictable. Basic theorems (halting problem, etc.) show that such systems are capable of behaviors that are deterministic but that no one can foresee [4]. If the funded project did not provide R&D toward non-generic (analyzable) cyber systems, then the question is unanswerable. We can perhaps collect data on real-world compromises, but this tells us nothing about the compromises that have occurred *without our knowledge* or those that the attacker may invent tomorrow. So-called security technologies that are bolted onto today's *undecidable* systems can only confer *anecdotal* security. On the other hand, a project that did advance toward non-generic (analyzable) cyber systems should be judged by the normal standards of science-based R&D – the theoretical soundness and empirical reproducibility of its security results.

#4. Why is this problem so hard?

As noted, generic cyber systems are inherently unpredictable. This gives an asymmetric advantage to attackers. An implication of cyber systems' complexity is that *no generic metrics* exist for cybersecurity – generic digital systems have an *unknowable* set of vulnerabilities. This explains why more careful programming, more extensive testing, and perimeters like anti-virus and intrusion detection accomplish very little or nothing in staying ahead of adversaries. Further, while dealing with an untrusted supply chain makes the problem all the more challenging, complex systems have plenty of room for unanticipated vulnerabilities even when design and fabrication are done in a trusted setting. This is because the possible behaviors of a complex system are not in general foreseeable even by its designer; moreover, a *composition* of individually predictable components is not predictable in general.

#5. How will our civilization's response to this problem play out?

The biggest question is whether there will be a “cyber Pearl Harbor” of catastrophic (but not civilization-ending) cost that galvanizes attention and resources. This will affect to what degree society is willing to demand quantifiable security as a requirement on cyber infrastructure. Continuing to muddle through with anecdotal security may simply magnify the catastrophes to come.

#6. How do I address deterrence in this world?

Complexity enables unforeseeable obfuscation of cyber attacks. That is, since attacks are carried out with *bits* that act independently of their originator, and since untold numbers of vulnerabilities can be discovered by attackers, the mechanism used to compromise a complex cyber system can be made as convoluted, inscrutable, and untraceable as desired. As a result, attribution (hence deterrence) is generally not feasible. The underlying complexity problem can only be addressed by constraining our system *designs* in appropriate ways, not by any amount of monitoring.

References

- [1] J. Woodcock, P. G. Larsen, J. Bicarregui, and J. Fitzgerald. Formal methods: Practice and experience. ACM Computing Surveys, 41:19, 2009. <http://dl.acm.org/citation.cfm?doid=1592434.1592436>
- [2] Y. Vorobeychik, J. R. Mayo, R. C. Armstrong, and J. R. Ruthruff. Noncooperatively optimized tolerance: Decentralized strategic optimization in complex systems. Phys. Rev. Lett., 107:108702, 2011. <http://link.aps.org/doi/10.1103/PhysRevLett.107.108702>
- [3] R. C. Armstrong and J. R. Mayo. Leveraging complexity in software for cybersecurity. Proc. 5th Cyber Security and Information Intelligence Research Workshop, Apr. 2009. <http://dl.acm.org/citation.cfm?doid=1558607.1558643>
- [4] R. C. Armstrong, J. R. Mayo, and F. Siebenlist. Complexity science challenges in cybersecurity. Sandia Report SAND2009-2007, Mar. 2009. <http://sendsonline.org/wp-content/uploads/2011/02/DOE-Complexity-Whitepaper-2009.pdf>

2.5 Suppona

#1. I have \$1 Million. How should I spend it to maximize my computer security?

Technology investments for security will return little or no benefit unless larger investments are made in recruiting and training skilled and motivated security analysts. Given the right set of skilled individuals one could make minimal technology investments and obtain world-class cyber security. Investments not made in technology, including out-year maintenance and upgrade fees, are much better spent on retaining valued analysts. For example, an enterprise could obtain a variety of no cost open source security products (Snort, WireShark, BASE, etc.) and a small number of server-grade platforms on which to run them. The right set of people could develop an infrastructure around those products that is well-matched to that enterprise's computing environment and would likely outperform the most expensive commercial equivalents. It is always important to remember that a vendor's business plan cannot and never will match your security requirements; invest accordingly.

#2. I am a program manager for computer security. How do I identify the proposals that will increase my computer security?

The short answer is none. The longer answer is the proposals that are the most open-ended in terms of specifications, deliverables, and methodologies. Remember, you are the target. The harder you become to hit the more diverse and complex the attack methods must become in order to succeed. Any proposal that cannot adapt to new challenges with change orders or agreement modifications will result in net decreases in security. Rigid agreements do not stand up well to flexible adversaries.

In other words, don't treat security proposals like commodities.¹²

#3. I am a program manager for computer security. When a funded proposal completes how do I determine how much security I got for my money?

One assumes that metrics have been gathered that detail security events and their impact on the enterprise have been measured. If the past predicts the future, over time attacks will continue to increase in severity, complexity and quantity. That will not change regardless of protection methods in place or due to be put in place. You will know if proposals are successful if your state of security is at a minimum unchanged from prior to the proposal's implementation.

¹² The cyber security market is full of opportunistic marketing driven by current/popular threats. It's not uncommon for an existing product – perhaps it's been around for a few years – to suddenly become (in the eyes of the vendor at least) the tool that you need to protect your enterprise from that current threat. Folks reviewing those proposals need to have a very jaundiced eye when it comes to what they promise to deliver. And given the important role that is the outcome of those proposals, they can't be looked at the same way one might consider a proposal for a new copier.

Ideally there will be a positive movement of your state of security but that movement will likely occur in very small increments.

#4. Why is this problem so hard?

It is hard because we tend to treat computer security events as technology problems: a zero-day event occurred, a patch was not applied in time, social engineering occurred, etc. In fact these are people problems, theirs and ours. No amount of technology will cause an adversary to point their skills elsewhere (remember, they likely have access to the same technology or very good analogues). And no amount of technology will cause one of our people to not push the “Don’t Push This Button” button. And it really only takes one person (see various security events of the last 12 months).

#5. How will our civilization’s response to this problem play out?

Civilization will accept this as the status quo for a very long time – most events are one-off [i.e., unique] events that impact such a small percentage of the population that the rest of the population sees it in the terms of “I’m glad that I’m not you.” That will not change until there is an event of sufficient magnitude that a large percentage of the population is impacted in very serious ways including loss of life. Think in terms of the power grid, flight control systems, hospital infrastructure, or other components of our every day lives that most people do not consider to be “cyber.”

We will see significant changes in people’s attitudes once that happens, and we will also observe unintended consequences: We will likely see a move towards “underground computing” where some percentage of the population moves to covert channels, such as TOR,¹³ in order to hide their presence as a protection mechanism. But think TOR on a grand scale and with reasonable bandwidth.

#6. How do I address deterrence in this world?

Just like we think about in the physical world. We are badly in need of a sheriff and there is none to be had. Over time we have made it socially unacceptable and dangerous to various degrees to kick in doors and steal stuff. It still happens and will continue to do so. Nonetheless we respond when the opportunity is presented; homeowners shoot burglars, alarms and security cameras cause others to be caught, and sometimes someone turns someone else over to the authorities. It’s a model that keeps the lawless population to a small enough size that many folks will experience little more than petty crime during their lifetimes. Cyber needs a reasonable analogue.

¹³ See <https://www.torproject.org>.

2.6 Wyss

#1. I have \$1 Million. How should I spend it to maximize my computer security?

(See answer #3.)

#2. I am a program manager for computer security. How do I identify the proposals that will increase my computer security?

(See answer #3.)

#3. I am a program manager for computer security. When a funded proposal completes how do I determine how much security I got for my money?

With all due respect, I fear that we may be asking the wrong question here. I have spent the last couple of years trying to help people understand that risk – both safety risk and security risk – is *not a number*. Risk is something to be understood and managed and, while in some cases, representing those risks as numbers may be helpful, in many cases, it can lead people to a false sense of precision with regard to the actual vulnerabilities and risks that might exist within their system. This is especially true in the area of security risk, and particularly egregious in those situations where the risk analyst asks an expert to state the likelihood of an attack and then multiplies it by an expected consequence (using point estimate values without uncertainty estimates for both numbers). The true uncertainty in such risks is orders of magnitude, and yet investment decisions are often made “confidently” because the risk values differ by a factor of two.

I believe that the question “how much security” can only be answered in a relative sense, and then, only by comparison to some concept of the remaining unresolved (or “residual”) security risks. A qualitative security goal is useful, such as “for those events whose consequences are above X level, are the easiest attacks difficult enough to deter the adversaries I’m concerned about?” If one can assign a number to such a goal, it may be useful to decision makers. However, human decision-making processes are notably qualitative and imprecise and heuristic when it comes to selecting and planning courses of action such as attacks. I believe that a risk management method that does not embrace these features of human decision-making and acknowledge these uncertainties is at best unrealistically precise.

This is not to say that it is a fool’s errand for program managers to understand the benefits of their security investments. A security investment is beneficial to the extent that it addresses attack scenarios that would be particularly advantageous or attractive to an adversary. One criterion for an attractive scenario is one that embodies an opportunity for an adversary to achieve especially high consequences to the defender (or value to the adversary’s goals) with a relatively low level of investment or difficulty *by comparison to other available attack scenarios*. Thus, a security investment is of greater benefit if it significantly “raises the bar” of difficulty for an adversary to achieve particular consequences. It is of lesser value (but still possibly of some value) if it addresses one of several comparably attractive attack scenarios while leaving the

others still available to the adversary. It is of little value if the only attack scenarios that it addresses are those that are relatively unattractive to an adversary. Thus, one examines an adversary's opportunities before and after implementation of a security upgrade and to evaluate how security has improved, using criteria derived from the degree to which the adversary's job of attacking the system has been made more difficult by the security upgrade. This is not necessarily quantitative, but it represents risk management, in which the system owner understands the risks inherent in the system and prioritizes responses to those risks in a rational and systematic way.

A 2010 report by the National Academy of Sciences for DOE embodied similar types of recommendations, albeit at a higher level. The report said that DOE should focus more on *risk management* and less on "how much or how little risk exists." The report indicated that, while some probabilistic risk assessment (PRA) tools may be useful in understanding the system, PRA should not be the main basis for risk management, and that qualitative risk analysis and management methods may be very useful in this endeavor. A lot of these same thoughts are also applicable to questions 1 and 2.

Reference: National Academy of Science, "Understanding and Managing Risk in Security Systems for the Doe Nuclear Weapons Complex," Official Use Only report from the Committee on Risk-Based Approaches for Securing the DOE Nuclear Weapons Complex, Nuclear and Radiation Studies Board, Division on Earth and Life Studies, National Research Council, National Academy Press, Washington, D.C., 2010.

#5. How will our civilization's response to this problem play out?

To date, our society has often subconsciously engaged in the type of risk management described above, but with a twist. Our risk management decisions sometimes involve addressing the most attractive (generally measured as the most common) attack scenarios directly (e.g., firewalls, patches, virtual private networks); sometimes involve reducing the potential for consequences (e.g., not storing credit card information); and sometimes involve transferring the risk to others (e.g., limitations of liability on credit card accounts, identity theft insurance, business indemnity agreements). We have seen evolution of Internet-based commerce in the wake of changes in attack tactics by adversaries. And as long as the money to be made by enduring some level of such attacks dramatically exceeds the amount of money lost to such attacks, this model will continue to be exercised. This is reactionary, and may be adequate when dealing with ongoing criminal behavior.

A reactionary approach, however, is likely inadequate when one thinks about a potential cyber "Pearl Harbor," and I am unable to predict how we will respond to attacks that result in national and/or existential threats, which may require a major paradigm shift in our entire cyber infrastructure. We may be able to increase the resilience of systems, decrease consequence potential, and so forth, and in so doing, manage our vulnerabilities to these game-changing attacks. However, current investment philosophies for infrastructure and public utilities focuses on extracting ever more capacity from existing antiquated systems through increasingly complex and highly optimized system controls. We seek higher efficiencies and

eliminate underused resources in the name of cost savings. In so doing, we are working directly against resiliency and setting ourselves up for even greater consequences in the event of such an attack.

#6. How do I address deterrence in this world?

The most oft-cited and highly-studied modern example of deterrence is the Cold War between the United States and the Soviet Union. It is tempting to view the issue of cyber deterrence through this lens, but there are several critical dissimilarities between these situations. The Cold War was a bi-polar conflict in which two easily identifiable superpowers had comparable objectives (avoiding nuclear conflict while expanding their sphere of influence in the world) and comparable ability to wage devastating attacks. Cyber deterrence is a multi-polar conflict with unidentified and sometimes small non-state actors who engage in asymmetric conflict toward differing objectives, often based on fundamentally different social values (and, hence, definitions of success). Cold War deterrence relied on each nation's ability to invoke awesome retaliatory damage that may not be possible in many cyber-attack scenarios. Cyber deterrence is severely complicated by many adversaries with many different attack objectives or consequences possible, with great difficulty attributing attacks to specific adversaries and far less damaging retaliatory options. Thus, cyber space does not lend itself to the straightforward Cold War deterrence model.

To develop a workable cyber deterrence model, one must understand that deterrence is primarily related to the adversary's state of mind, as influenced by the defender and the surrounding world. That is, deterrence is fundamentally a decision by the adversary *not* to pursue a particular course of action because they believe it to be in some way against their interests. Evidence clearly shows that terrorists like Al Qaeda perform a cost-benefit analysis in evaluating attack opportunities. The defender's actions affect this decision mainly through three "C's":

- Capability - the defender has the ability to defend against or act to counter the attacker's move.
- Credibility - the defender's capability is perceived to be credible by the attacker.
- Communication - the defender communicates its credible capability to the attacker.

Clear communication is essential to deterrence of physical attacks, but all communication is susceptible to misinterpretation and deception. Uncertainty is also a key component of deterrence as an adversary may decide not to pursue a course of action if they are uncertain whether it will lead to an unacceptable end state. Both of these factors of deterrence are very adversary-specific. Thus, effective deterrence in cyberspace may mean that we will only stop some attacks and some attackers, and it seems fairly certain that we will not achieve a one-size-fits-all cyber deterrence solution. This confounds our ability to make sense of the problem at a national level.

The current lack of consequences to an attacker has had an additional effect that actually encourages rather than deters cyber-attacks. Many young people use technology to play competitive games. Computers and networks have historically been seen by many smart young

people as puzzles: things to explore, challenges to overcome, things to learn about, etc. Our best-defended systems (or those systems incorporating the latest security) are magnets for their attention. Hacking can be done from the convenience of an attacker's home, school, or smart phone. The technology they use to mount attacks is cheap because it is commercially available (i.e., they can buy it if necessary to support their "learning" (i.e., hacking)). Both free and inexpensive training and tools are available. In the end, the very nature of cyber/IT security creates inspiration in the hacker community to get better and better so that they can obtain bragging rights, career advancement, and so forth. Thus, our cyber systems and defenses have become the trophies that hackers collect. Then, nation states, organized crime, and hacktivist gangs recruit these people and give them even more training, a mission, and/or a job. This is very different from physical security world: kids don't run around testing our security at weapons sites because they will get arrested, go to jail, get hurt, or die.

For cyber deterrence to be effective, one must affect the mindset of the myriad potential cyber adversaries, from recreational puzzle-solvers to high-end criminals and nation states. Different deterrence models are probably required for the different types of adversaries because they have different value sets and decision criteria. Reducing the effect achieved by a successful cyber attack may be sufficient to deter adversaries like criminals or terrorists, while attribution and effective retaliation may be required for others. Communication of credible capabilities, which is so necessary in physical deterrence, actually encourages attacks in the cyber domain because of the absence of effective retribution. Whatever the ultimate solutions for cyber deterrence turn out to be, it is likely that they will be widely varied and will follow a very different model from that of physical deterrence. And starting from the mindset of Cold War deterrence may limit our creativity so much that it does more harm than good when seeking a workable model for cyber deterrence.¹⁴

¹⁴ *Acknowledgment:* Most of the thoughts on deterrence expressed in this section originated with John Clem and Mark Mateski of the Security Systems Analysis Department of Sandia National Laboratories. This section summarizes some of their internal communications and working papers on the subject of deterrence. The author is deeply grateful for the opportunity to echo and excerpt their thoughts for this forum

3 SUMMARY

The answers we provide are diverse and resist consolidation. This reflects the nature of the problem. To ease the work of the reader an overarching directive is provided, as shown in the next table. This directive is the one that we each consider to be the most important for improving computer security. We have also provided a single directive for each of the questions, shown in the subsequent table.

Table 1 Overarching Directives

Respondent	Overarching Directive
Berg	Eliminate known classes of vulnerabilities.
Campbell	Start with Decision Analysis.
Davis	Focus on the core.
Mayo	What's hurting us is what we <i>don't know</i> . Demand and create systems with quantifiable security.
Suppona	It's a people problem. Get over it.
Wyss	Increase the job of the attacker.

For ease of reference with the next table, the list of questions shown above is reproduced here:

- #1. I have a million dollars; how should I spend it to maximize my computer security?
- #2. I am a program manager for computer security. How do I identify the proposals that will increase my computer security?
- #3. I am a program manager for computer security. When a funded proposal completes how do I determine how much security I got for my money?
- #4. Why is this problem so hard?
- #5. How will our civilization's response to this problem play out?
- #6. How do I address deterrence in this world?

Table 2 Single Directives

Respondent	Question					
	1	2	3	4	5	6
Berg	You maximize your computer security if you eliminate known classes of vulnerabilities.	(See answer #1.)	(See answer #1.)	This problem is so hard because our definitions of "secure" and "insecure" are wrong.	(There are too many unknowns to provide an answer.)	You address deterrence by increasing the cost to the adversary of targeting your systems.
Campbell	You maximize	(See answer #1.)	You determine	This problem is	Our civilization	You address deterrence

Respondent	Question					
	1	2	3	4	5	6
	your computer security if you start with Decision Analysis.		how much security you got by using your organization's definition of security.	so hard because we use reductionism on systems and ignore Decision Analysis.	will respond by using the same approach we have used for other changes that have increased both efficiency and risk.	by increasing resilience.
Davis	You maximize your computer security if you diversify, focus on the core, define security, simplify, use roles, use better authentication.	The best proposals are the most holistic, detailed proposals.	You determine how much security you got by comparing with a clear goal chosen before the proposal began.	This problem is so hard because each situation here is unique.	Our civilization will respond by building security in. (Be patient: computing is still young.)	You address deterrence by supporting better national and international legal definitions.
Mayo	You maximize your computer security by if you design systems whose security is <i>measurable</i> .	(See answer #1.)	You determine how much security you got by achieving analyzability ; anecdotal security is not an answer.	This problem is so hard because generic cyber systems are inherently unpredictable.	Our civilization will respond, in the <i>best</i> case, by being awoken by a "cyber Pearl Harbor."	(Deterrence that relies on attack attribution is infeasible.)
Suppona	You maximize your computer security if you invest in people.	The best proposals are the ones that maximize your system's flexibility.	You determine how much security you got by measuring the impact of attacks: if they are not	This problem is so hard because we mistakenly think this people problem is a technology	Our civilization will respond by continuing as we are now until a big event occurs,	You address deterrence by treating cyber criminals like criminals.

Respondent	Question					
	1	2	3	4	5	6
			getting worse you are doing well.	problem.	whereupon we will compute "underground."	
Wyss	(See answer #3.)	(See answer #3.)	You determine how much security you got by measuring the increase in difficulty for the attacker.		Our civilization will respond by making incremental responses to evolving attacks, because of the profits to be made in cyber commerce, unless a cyber "Pearl Harbor" forces a major paradigm shift.	You address deterrence by convincing each adversary type not to attack. We are limited to partial and piecemeal deterrence: the asymmetric nature of cyber attacks and varied adversaries means cyber space does not lend itself well to Cold War deterrence models.

4 DISTRIBUTION

1	MS 0899	RIM-Reports Management	9532 (electronic copy)
1	MS 0260	D.R. White	5620
1	MS 0359	A.N. Campbell	1950
1	MS 0423	R.L. Craft	00247
1	MS 0672	P.L. Campbell	5629
1	MS 0672	H.W. Lin	5629
1	MS 0673	M.J. Berg	5629
1	MS 0757	J.F. Clem	6612
1	MS 0757	M.E. Mateski	6612
1	MS 0757	G.D. Wyss	6612
1	MS 0795	R.A. Suppona	9317
1	MS 1205	J.R. Gosler	00002
1	MS 1248	C.E. Davis	5623
1	MS 9151	R.L. Hutchinson	8960
1	MS 9159	J.R. Mayo	8953

