

Input to the Commission on Enhancing National Cybersecurity

Steven M. Bellovin Adam Shostack
<https://www.cs.columbia.edu/~smb> <http://adam.shostack.org>
Columbia University¹ Independent

Thank you for the opportunity to provide Input to the Commission on Enhancing National Cybersecurity. This is a joint submission by Steven M. Bellovin and Adam Shostack. Steven M. Bellovin, a member of the National Academy of Engineering, is the Percy K. and Vida L.W. Hudson Professor of Computer Science at Columbia University. Adam Shostack is an entrepreneur and the author of *Threat Modeling: Desinging for Security*.

We are writing after 25 years of calls for a “NTSB for Security” have failed to result in action. As early as 1991, a National Research Council report called for “build[ing] a repository of incident data” and said “one possible model for data collection is the incident reporting system administered by the National Transportation Safety Board.” [1] The calls for more data about incidents have continued, including by us [2, 3].

The lack of a repository of incident data impacts our ability to answer or assess many of your questions, and our key recommendation is that the failure to establish such a repository is, in and of itself, worthy of study. There are many factors in the realm of folklore as to why we do not have a repository, but no rigorous answer. Thus, our answer to your question 4 (“What can or should be done now or within the next 1-2 years to better address the challenges?”) is to study what factors have inhibited the creation of a repository of incident data, and our answer to question 5 (“what should be done over a decade?”) is to establish one. Commercial air travel is so incredibly safe today precisely because of decades of accident investigations, investigations that have helped plane manufacturers, airlines, and pilots learn from previous failures.

The problem, in its simplest form, is that we do not have a good idea of what is going wrong in cyber-security. Lacking a repository of incidents, information about the causes of those incidents, or means of discussing controls, we are unable to assess scientifically if our advice is effective. (To your question 1, why is asking the public what are current trends and challenges the core of deciding a research agenda? Why do we not have a more structured way to learn that?)

We lack a repository because we have tacitly agreed that having such a repository is not worth overcoming the barriers to its creation, and we have tacitly agreed to not discuss our mistakes. The reasons for this are worthy of study, and such study should inform efforts to overcome our inability to learn from our mistakes. A better understanding of what goes wrong could better inform discussion of many of the questions put forth in the

¹ Affiliation listed for identification only.

call for evidence, but I submit that being able to address “what has gone wrong recently” and “how is that changing,” represent an important emerging research challenge. Allow us to expand on each point.

We, as professional and research communities, do not have a good idea of what is going wrong in cyber-security. For example, how many computers were compromised last month and how? We know approximately how many people got sick and with what, and that information is be assessed and gathered in a process which begins with “incident reports” at hospitals and morgues. Each member of the advisory group has a view, informed by problems they observe, but we lack a catalog of all information security incidents which occurred or those which met some threshold. We know of some subset as a result of various rules requiring disclosure of breaches which release personal information, but the relation of that set to the larger context is unknown. There is a blind-men and the elephant problem in cyber-security. (In brief, each blind man touches one part of the elephant, and describes it: “it’s like a snake” says the one touching the trunk, “no, it’s like a tree” says the one touching the leg.)

We lack a public repository of incidents, information about the causes of those incidents, and means of discussing controls, even though the National Research Council asked for one in 1991. Perhaps closest to a repository for study are the collections of breach disclosure letters published by some US state governments. After some security breaches, we learn surface information such as “a hacker broke in” or “policy was not followed.” (Sometimes, we learn that we cannot know details because logs were not retained, systems were wiped clean as part of responding to an incident, or other issues, but we cannot even study those issues.) There are some private information sets, for example, held by CERTs. These are generally not available for academic research, peer review, and in conversation with such organizations, the data is not high quality. Information about the hacking technique is rarely forthcoming, even to the level of did the attacker use a vulnerability or convince someone to take action? What controls were in place which might be expected to prevent the attack? Was there a failure to act on information because of too many alarms? A mis-configuration? Did some control simply fail? (Perhaps revealing such information invites additional attacks, but that hypothesis is untested.) Information about the root causes beyond the attack is rarely gathered or shared. (We use the term “root causes” as few complex incidents have a singular cause.) Without such analysis, it is unsurprising that the same mistakes are made over and over again. Compare and contrast to the state of air travel safety information, where a team of trained investigators will show up to investigate incidents of a certain severity, and then publish their reports for all to learn from.

We are unable to assess if our advice is effective. There are a wide variety of advice for defending a system, ranging from the Australian DSD “top 35” controls, 12 PCI controls, the SANS top 20, and US “Federal Risk and Authorization Management Program” set of 168 controls. Which works better? Which works better for a given level of investment or sort of organization? We don’t know and in many ways, often seem to accept our failure as if the difference is an unsolvable mystery. The standards bodies that put forth advice often lack even rudimentary feedback loops. For example, PCI is mandated for

those accepting credit cards, over a great deal of grumbling. They do not require after-incident reports be filed with the PCI Council, rather asserting that anyone breached was not compliant.

We lack a shared repository of incidents, root causes or analyses because we have agreed to not discuss our mistakes. The steering group is certainly familiar with the reasons for not discussing security incidents, but we can do better than folk knowledge. The reasons for this are worthy of study. Analysis of disclosed breaches (after the passage of American laws requiring such disclosure) shows that stock price barely moves, and that customers do not flee in droves [4]. A more complete list of reasons why organizations prefer not to discuss incidents would enable study of those reasons, and perhaps help us overcome the reticence. (Of course, it's possible that simple shame will win out.)

An incident database will have other advantages. Insurance companies could use it to help set rates for hacking coverage: right now, the lack of actuarial data means that there are not sound incentives for good behavior. This in turn means that the market cannot work its magic to improve security.

Similarly, regulators, such as the Federal Trade Commission, will be able to consult the data to assess whether or not companies have indeed followed best practices, thus putting its common law approach [5] on a sound footing. Even students will benefit; they will have real-world systems to study.

It may be that setting up a proper structure cannot be implemented at this time. In the interim, we suggest that an anonymous reporting system, similar to NASA's Aviation Safety Reporting System (<http://asrs.arc.nasa.gov/>), be deployed with incentives for reporting.² Such schemes won't have the in-depth investigations that are really needed; that said, a good set of problem descriptions will help researchers and conscientious practitioners.

The government could also set an example by establishing a reporting structure and analysis structure for its own incidents. This is similar to the intent of the Privacy Act of 1974:³ Congress declined to create mandatory requirements for the private sector, but imposed (what were for the time) best practices on Federal agencies. Having such a structure in place for security would also end the current uneven record of investigations and publication.

Being able to address “what has gone wrong recently” and “how is that changing,” represent an important emerging research challenge. These questions prompt others: Should there a body chartered and funded to gather information about cyber-security incidents? Would research into what methods for analyzing incident root causes generates the best results (and what metrics should be used for assessing best)? There are

² NASA runs a similar system for railroad incidents at <http://c3rs.arc.nasa.gov/>.

³ 5 U.S.C. § 552a.

a variety of methods for sharing or publishing information. What are the tradeoffs between aggregated, anonymized or other approaches?

As such questions are addressed, and more data becomes available, there will be a need to understand its validity over time. Technologists are fond of claiming that “this changes everything.” How important are those changes to security? Can we take lessons from the rise of personal computers and apply them to the current challenges of “bring your own device?” What lessons will transfer from the internet of general-purpose computers to the internet of things?

References

- [1] System Security Study Committee. *Computers at Risk: Safe Computing in the Information Age*. National Academies Press, 1990.
- [2] Bellovin, Steven M. “The major cyberincident investigations board.” *IEEE Security and Privacy* 10:6, November-December 2012.
- [3] Shostack, Adam, and Andrew Stewart. *The new school of information security*. Pearson Education, 2008.
- [4] Acquisti, Alessandro, Allan Friedman, and Rahul Telang. "Is There a Cost to Privacy Breaches? An Event Study." In *WEIS*. 2006.
- [5] Solove, D.J. and Hartzog, W., 2014. The FTC and the new common law of privacy. *Columbia Law Review*, pp.583-676.