

ESSAY

Cybersecurity and National Policy

Daniel E. Geer, Jr., Sc.D.*

What follows is the author's own opinion, *i.e.*, it is not a strategic leak of anything going on in any corridor of power. I am reminded of one of my mentors who said that, if you are lucky, as you age you will compensate for your loss of creativity with burgeoning critical skills. As such, much of what I say here will be critical, but I hope that it is taken as critical in the sense of analytic rather than critical in the sense of harping.

Let me begin with some biases. First, security is a means, not an end. Therefore, a cybersecurity policy discussion must necessarily be about the means to a set of desirable ends and about affecting the future. Accordingly, security is about risk management, and the legitimate purpose of risk management is to improve the future, not to explain the past.

Second, unless and until we devise a scorekeeping mechanism that apprises spectators of the state of play on the field, security will remain the province of "The Few". Sometimes leaving everything to The Few is positive, but not here as, amongst other things, demand for security expertise so outstrips supply that the charlatan fraction is rising.

Third, the problems of cybersecurity are the same as many other problems in certain respects, yet critically different in others. We often misclassify which characteristics are "same" and which are "different", beginning with the sharp differences between the realities of time and space in the physical world versus the digital world. Examples of these differences

* Chief Information Security Officer, In-Q-Tel.

include the original owner continuing to possess stolen data after the thief takes it, and law enforcement lacking the ability to work at the speed of light.

When I think about cybersecurity and national policy, I can only conclude that the problem is the problem statement. At the highest level of abstraction, let me propose that the problem statement for a National Policy is this:

To move from a culture of fear,
to a culture of awareness,
to a culture of measurement.

While this statement is not operationalizable per se, it demonstrates my biases that security is a means and that game play cannot improve without a scorekeeping mechanism.

None of that is new; the worry over fear has been said all but word-for-word for almost five centuries.

The thing I fear most is fear. — Montaigne, *ca.* 1570
There is nothing terrible but fear itself. — Bacon, 1620
The only thing I am afraid of is fear. — Wellington, 1836
Nothing is so much to be feared as fear itself. — Thoreau, 1851
The only thing we have to fear is fear itself. — Roosevelt, 1933

Neither is the idea of a goal state where measurement holds sway, as in Lord Kelvin's iconic 1883 remark:

When you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind; it may be the beginning of knowledge, but you have scarcely, in your thoughts, advanced to the state of *Science*.

But enough of my beliefs, though now that you know them, perhaps you can account for them in my remaining remarks.

To set the rest of what I am going to say on the bedrock of its foundation, the United States's ability to project power depends on information technology, and, as such, cyber insecurity is the paramount national security risk. This point bears repetition: because the United States's ability to project power depends on information technology, cyber insecurity is the paramount national security risk.

Those with either an engineering or management background are aware that one cannot optimize everything at once — that requirements are balanced by constraints. I am not aware of another domain where this is as true as it is in cybersecurity and the question of a policy response to cyber insecurity at the national level. In engineering, this is said as “Fast, Cheap, Reliable: Choose Two”. In the public policy arena, we must first remember the definition of a free country: a place where that which is not forbidden is permitted. As we consider the pursuit of cybersecurity, we will return to that idea time and time again; I believe that we are now faced with “Freedom, Security, Convenience: Choose Two”.

Some time ago, I was asked to categorically state what types of risks rose to such a level that they could legitimately be considered national security concerns. Based on my view, much like the Treasury Department's view on bank failure — that a public loss of confidence is to be avoided at all bearable cost, but that everything short of this amounts to nothing more than some private tragedy — my answer then was that only two kinds of vulnerabilities were that important. The first is any mechanism that, to operate correctly, must be a single point of function, thereby containing a single point of failure. The red telephone on the President's desk is just such a mechanism; having twenty-three red telephones would be far worse than having one red telephone. As such, that single red telephone deserves defense in depth, which is simply a referendum on one's willingness to spend money for layers; it is rarely, if at all, a research-grade problem.

The other national security scale risk is cascade failure,¹ and cascade

¹ A cascade failure is a failure that may begin at any of a large number of equivalent locations but, once begun, proceeds due to the interconnectedness of the components of a larger system. In nature, an avalanche is a cascade failure; in electric power, a cascade

failure is so much easier to detonate in a monoculture — when the attacker has only to write one bit of malware, not ten million. The idea is obvious; believing in it is easy; acting on its implications is, evidently, difficult. Despite what you might think, I am sympathetic to the actual reason we continue to deploy computing monocultures² — making systems almost entirely alike remains our only hope for managing them in a consistent manner. Put differently, when you deploy a computing monoculture you are making a risk management decision: that the downside risk of a “black swan event”³ is more tolerable than the downside risk of perpetual inconsistency.

Since first considering that question, I have decided there is now another national security scale issue, arising as a side effect of global supply chains and device complexity. It is simply not possible to provide product or supply chain assurance without a surveillance state. This matters not just philosophically, but practically.

The root source of risk is dependence — dependence on system state, including dependence on expectations of system state reliability. Indeed, my definition of security has co-evolved with my understanding of risk and risk’s source, to where I currently define security as the absence of unmitigatable surprise. Thus, increasing dependence results in heightened difficulty in crafting mitigations. This increasing complexity embeds

failure occurs when one power station goes offline and increases the load on others, leading to the failure of the weakest remaining operating unit, and so forth.

² Creating a computing monoculture contributes to managerial efficiency but risks having existing vulnerabilities being exploited with similar industrial efficiency. In the agricultural setting, the spread of a blight across a susceptible species is made worse if that species occupies entire counties or larger. In the computing setting, an exploitable vulnerability invites attack in proportion to its widespread identicality, as the attacker has a much better return on his investment in crafting the exploit when the breadth of vulnerable targets is all but universal.

³ “Black swan event” is a term of art attributable to Nicholas Taleb, who authored a book of the same name. The idea is that the things we (should) fear most are too rare to develop experience with handling and, frequently, represent complex failure modes for which no one can be faulted for not having recognized *a priori*. Taleb also stresses that our reaction to an unprecedented beneficial event will be mild and relaxed whereas our reaction to an unprecedented malicious event will be panic-driven and anything but relaxed, and this asymmetry of effect is present even when the beneficial and malicious alternatives have equal probability. *See generally* NASSIM NICHOLAS TALEB, *THE BLACK SWAN* (2007).

dependencies in a manner that may diminish the frequency of surprises; however, the surprises will be all the more unexpected when they inevitably occur.

And that is the crux of the matter: our dependence on all things cyber as a society is now inestimably irreversible and irreversibly inestimable. That sounds more apocalyptic than I intend, but the competent risk manager always asks, “How bad could it be?” or, in the altogether American tortious style, “Who will have to pay?”

This leads to my first conclusion about cybersecurity and national policy: our paramount aim cannot be risk avoidance but rather risk absorption — the ability to operate in degraded states, in both micro and macro spheres, to take as an axiom that our opponents have and will penetrate our systems at all levels, and to be prepared to adjust accordingly.⁴ To this extent, security becomes a subset of reliability in that an insecure system will not be reliable, but a reliable system is not necessarily secure.

That tenet of a free society, *viz.*, anything not forbidden is permitted, interacts strongly with the rate of change in the digital world. No society needs rules against impossibilities. The rate at which we are turning the impossible into the possible is accelerating and will continue to do so because technologic change is now in a positive feedback loop. This leads to my second conclusion: free society rulemaking will trail modalities of risk by increasing margins, even if that rulemaking comes (God forbid) from some one-world government. This second conclusion evokes the Second Amendment, that an armed citizenry is a *sine qua non* of freedom.

Ed Giorgio, then chief cryptanalyst for the NSA, famously remarked that, “In our line of work, security and privacy are a zero sum game.” I do not intend to argue his point for him. I do not have to; every proposal for reinventing the Internet stresses the law-and-order essentialness of strong authentication in service of attribution, based on the finding that without attribution there is no deterrence to cybercrime because forensics will never close the widening evidentiary gap. Looking backwards, we see ready examples: the general public was only too happy to yield to cell-phone location tracking so that they could call 911 without having to know where

⁴ I note, for the record, that the United States can absorb substantially more risk than most small countries, which, on a relative basis, are suffering or will suffer the most harm.

they were. Looking forward, without universal strong authentication, tomorrow's cybercriminal will not need the fuss and bother of maintaining a botnet when, with a few hundred stolen credit cards, he will be able to buy all the virtual machines he needs from cloud computing operators. In short, my third conclusion is that if the tariff of security is paid, it will be paid in the coin of privacy.

As much as I would wish, the market is unlikely to come to the rescue here, since a market only exists where there is demand. I do not see that demand; I do not see it in the general population (which is far more dependent on the digital world than it wants to realize, inured as it is to convenience *über alles*), and I do not see it in government. It has been said over and over for twenty years, "If only we could make government's procurement engine drive the market toward secure products." This, ladies and gentlemen, is a pleasant fiction. First, the United States government's buying power is staggering but, on a world scale, it is less than market-driving, and growing less so. Second, we have long since demonstrated that the single greatest barrier to introducing innovation into government is that procurement rules obliterate any hope of real entrepreneurs approaching the city gate. Third, 90-plus percent of the installable base is not in government, but in the private sector. Sadly, then, my fourth conclusion is that market demand is not going to provide, in and of itself, a solution.

It would be rude and facile to repeat that when you do not know where you are going, any direction will do. But there are so many directions we could go now that I risk sounding rude and facile. Putting aside bread and circuses, I wish we had some sort of consensus on what goal state we wanted, but I am not even sure of that, and, in any case, the rate of change may not only make means temporary, it may also do so for ends. You can think of this as evolution in that evolution's winners are a random selection.

Government's usual triad of tools — regulation, taxation, and insurance pricing — are all potential avenues; however, we must remember that all are subject to what I refer to as the Four Verities of Government:

Most meaningful ideas are not appealing.

Most appealing ideas are not meaningful.
Not every problem has a good solution.
Every solution comes with side effects.

Nevertheless, let us consider what we might do.

Government regulation is easiest to apply when the regulatees are few and those few are already well regulated; it is far harder to regulate a fragmented vastness and/or to introduce regulation where there was none. For this reason, and whether just or unjust, the major telecoms will continue to be compelled to play the role of government's outsourced private police force. This took one form under the Bush administration and it is taking an all but identical form under the Obama administration. The demand for "safe pipes" inexorably leads to deputizing those who own the most pipes. Even so, that beats nationalization.

Accretive sequestration of social policy in the Tax Code is precisely why the Tax Code is complex. Using the Tax Code to encourage investment in cybersecurity is just as possible as using the Tax Code to encourage investment in research and development. One must note, however, that using the Tax Code to do anything has perhaps the richest source of side effects, also known as unintended consequences. Inserting cybersecurity concerns into the Tax Code would be no different.

Government can drive insurance pricing primarily through the codification of liability, which, in turn, forces collectivization of liability risk into insurance pools. There are many options for liability here, and, as any software buyer knows, the high-order bit of every page of every end-user license agreement (EULA) reads, "It is not our fault." EULAs are an outrage and require fixing, yet it is all but clear that attempting to regulate software quality, even given the poor quality of monopolistic providers, just exports the software business to China in perpetuity.

Consistent with the paradigm of risk tolerance, it might be possible to say, as an engineer would, that no system may fail silently. In a sense, that is what the "data breach" laws assume — that breaches are inevitable and the proper response is notification of affected parties. Interestingly, the first of these rules, California SB 1386, was drafted by taking a toxic waste spill rule and substituting data-on-the-information-superhighway for poison-

on-the-public-street.⁵ If you consider notification an adequate mitigation, then this law creates a form of security as the surprise is followed by mitigation. The regulation here would be performance standards for latency of cleanup steps, like notification.

Verizon's 2009 Data Breach Investigations Report⁶ included two findings that are more important than all the others: one, that 75 percent of all data losses are discovered by unrelated third parties; and two, that whether data breaches are preponderantly insider attacks or outsider attacks depends on your definition of insider. If "insider" means "on the payroll", then insider attacks are not the most important issue. If, however, you define insider to include folks on your payroll plus employees of your partners who have as-of-right access to your data, then the majority of data losses are insider attacks.

Those two findings must be true of government too, and even more so as is implied by the mention of procurement. Government could lead by example here, not as a market-directing procurement model, but in terms of requiring government suppliers to be sufficiently instrumented such that data loss events are discovered by the second party, not some third party, and that, as a condition of contract, the contracting firm must attest to its recent data-handling performance at regular intervals. In other words, treat data as if it were money. As data represents an increasing fraction of total corporate wealth, this is less cathartic than other options.

It is important to understand that cyber insecurity is driven by sentient opponents, not by bad luck or stray alpha particles. At the same time, that the opponents think and calculate and are not inanimate random processes only changes the clock, not the azimuth of drift. This, perhaps, strengthens the argument for latency-based performance standards.

That 90-plus percent of the critical infrastructure is in private hands means that state-level opponents unremarkably spend over 90 percent of their effort on the private sector. No component of the commercial world is without compromise, but the primary targets are essential industries, the secondary are the counterparties to those primary targets, and the tertiary

⁵ See ANN. CAL. CIV. CODE §§ 1798.29, 1798.82, 1798.84 (effective July 1, 2003).

⁶ VERIZON BUSINESS, 2009 DATA BREACH INVESTIGATIONS REPORT (2009), available at www.verizonbusiness.com/resources/security/reports/2009_databreach_rp.pdf.

are sites that can be prepped for future attacks. It would be wise to structure our defenses accordingly. This leads directly to whether government's cooperation with the private sector should not focus on the Defense Industrial Base and if the Defense Industrial Base should be expanded to include cybersecurity firms and technology within its remit.⁷

Some degree of international engagement is essential for no other reason than that our opponents are location-less. Much work has been done on this, but the path to any treaty is steep and the clock of upward progress ticks in years, not minutes. The Council of Europe's Convention on Cybercrime⁸ is a case in point. At the same time, the recent decision of the Internet Corporation for Assigned Names and Numbers (ICANN) to wildly proliferate the number of top-level domains and the character sets in which domains can be enumerated is the single most criminogenic act ever taken in or around the digital world. United Nations treaties are all but useless — unenforceable and thus popular with the worst state offenders — so “coalitions of the willing” are the best we can hope for, taking the G8's Financial Action Task Force⁹ as an example. Put differently, international engagement is likely necessary but certainly insufficient.¹⁰

It is straightforward to see the value of information sharing: as a matter of logic you cannot, for example, know whether you are a target of choice or a target of chance unless you compare your attack pressure to that

⁷ See U.S. DEP'T OF HOMELAND SECURITY, NATIONAL INFRASTRUCTURE PROTECTION PLAN: PARTNERING TO ENHANCE PROTECTION AND RESILIENCY (2009), *available at* http://www.dhs.gov/xlibrary/assets/NIPP_Plan.pdf.

⁸ Convention on Cybercrime, *opened for signature* Nov. 23, 2001, T.I.A.S. 13174, E.T.S. 185.

⁹ The Financial Action Task Force (FATF) is an inter-governmental body committed to the development and promotion of international policies to combat money laundering and terrorist financing. FATF Home Page, <http://www.fatf-gafi.org>.

¹⁰ Note that I am not covering the question of what some call cyberwarfare. I refer you to two further readings if that is your interest: the analysis of the Russian-Georgian conflict by the U.S. Cyber Consequences Unit, and the National Research Council's consideration of the state of cyberwarfare policy. See U.S. CYBER CONSEQUENCES UNIT, OVERVIEW BY THE US-CCU OF THE CYBER CAMPAIGN AGAINST GEORGIA IN AUGUST OF 2008 (2009), *available at* <http://www.registan.net/wp-content/uploads/2009/08/US-CCU-Georgia-Cyber-Campaign-Overview.pdf>; NATIONAL RESEARCH COUNCIL, TECHNOLOGY, POLICY, LAW, AND ETHICS REGARDING U.S. ACQUISITION AND USE OF CYBERATTACK CAPABILITIES (William A. Owens et al. eds., 2009), *available at* <http://www.anagram.com/berson/nrcoiw.pdf>. These two collectively make the case that non-governmental actors can play a major role in cyberwarfare and that cyberwarfare is freighted with unanticipatable collateral damage.

of others. Sharing information between the private sector and the government has been worked on in many ways, but, to my mind, little has come of all that good intent and effort. To begin with, no General Counsel in any industry believes that the protections against Freedom of Information Act (FOIA) access to security data shared with the government will actually work when the time comes, nor do they believe that they would have effective recourse if proven correct in their skepticism. Now repeat that same sentence substituting “antitrust” for “Freedom of Information Act”. General Counsels are not being unduly risk averse here — the codified protections against FOIA and antitrust, as applied to private security data, have never been tested in court. And as our newest Supreme Court Justice has candidly said, “Real policy is made at the Circuit Court of Appeals.”

Yet for all of that, we clearly need information sharing. My own hope has long been for a technical guarantee of non-reversibility of shared data, something you can call de-identification or even anonymization of log data, thereby removing the General Counsel from the equation without giving congressional committees new weapons. Others I trust and admire have repeatedly proven by demonstration that it is all but impossible to craft such a technical guarantee and thus it is the technologists who argue for a procedural guarantee even as the sadder-but-wiser policy people pine for a technical guarantee. There does not seem to be a simple solution to this problem, though, in the private sector, some sharing does take place; for example, banks quickly share suspicions about stocks actively involved in pump-and-dump schemes. Providing security clearances to the management committees of every U.S. business with a claim to criticality does not seem workable either.

Above I opined that the ability to operate in a degraded state is an essential capability for government systems and private sector systems. A second essential capability is a means to assure broad awareness of the gravity of the situation. To the extent that awareness is a trained response, we have to ask whether we can get awareness without the “training” that a thoroughgoing crisis provides. Yes, decision-making under crisis conditions is especially fraught with unintended side effects, but memory of a crisis usefully serves to keep certain issues clearly in our mind. The unanswered question is whether we can proactively keep awareness of cybersecurity clearly in our mind. There are many who have proposed that the process we went through as a nation to understand, set policy for, and contain

nuclear weapons has lessons for us here, because the gravity of the nuclear situation and our awareness of it did not flag despite decades of relative quietude.

There is a third essential, one that flows from recognizing the limits of central action in a decentralized world, and that is some measure of personal responsibility and involvement. We all know that patching behavior leaves much to be desired — Verizon’s report showed that data loss events frequently involve open vulnerabilities for which patches had been available in excess of a year at the time of breach, and Qualys demonstrated that actual in-the-field patching follows a half-life curve, and thus never completes.¹¹ We all have seen the scanning results that show a majority of home machines are compromised. We all have heard that the price of stolen personal data is falling as the supply side grows ever more efficient and automated. We are all aware that at the present levels of infection, peer-to-peer pairings almost always involve a transmission opportunity. And so I ask, whose responsibility is this?

You may view an infected machine as a weapon. If I do not lock up my guns and they are used for the commission of a crime, then I will, at the very least, have some explaining to do. You may simply not want to drive through an intersection if you know that a majority of opposing traffic lacks brakes. I do not believe we will find the political will to make personal culpability a serious enough matter to effect widespread change, but I am at a loss to argue in any other direction. I ask this: if it is not the responsibility of the end user to avoid being an unwitting accomplice to constant crime, then whose responsibility is it? If you say that it is the responsibility of Internet Service Providers (ISPs) — that “clean pipes” argument¹² — then you are flatly giving up on not having your traffic inspected at a fine level of detail. If you say that it is the software manufacturer’s responsibility, we will soon need the digital equivalent of the Food and Drug Administration to set

¹¹ See QUALYS, INC., THE LAWS OF VULNERABILITIES: SIX AXIOMS FOR UNDERSTANDING RISK (2006), available at <http://www.qualys.com/docs/Laws-Report.pdf>.

¹² “Clean pipes” is a term of art describing one possible mechanism for general Internet safety where ISPs are responsible for the data they carry and, as such, are obliged to conduct inspection before cartage begins. This is in complete and sharp contrast to the “common carrier” assumption that underlies not only Internet service provision but also commercial freight handling and many other aspects of modern life — *i.e.*, it is not the responsibility of the carrier to inspect what it carries beyond the obvious ideas of not accepting, say, a package which is emitting smoke at the time of acceptance.

standards for efficacy and safety. If you say that it is the government's responsibility, then the mythical Information Superhighway Driver's License must soon follow. To my mind, personal responsibility for Internet safety is the worst choice, except for all the others.

At the risk of repetition, let me be clear that contemplative reaction to compromised counterparties is the core of awareness. They must be dealt with, and they demonstrate why there must be a personal responsibility component, or else we'll be left with Big Brother. Ever since my team created the Kerberos network authentication system,¹³ the idea has been "I'm OK and you're OK, but the big bad network in between us cannot be trusted for a second." Authentication, authorization, and accountability all begin with authentication — and that, in turn, begins by asking the Operating System the name of the user. What has really changed is that it is not true that "I'm OK and you're OK", since it is entirely likely that the counterparty to whom you are connecting is already compromised. It is more like "I think I'm OK, I have to assume you are Owned and the network may make this worse." A secure network connection? Who cares if the other end is hosed? Purdue's Gene Spafford was correct, but early, when he likened network security in the absence of host security to hiring an armored car to deliver gold bars from someone living in a cardboard box to someone sleeping on a park bench.

These are heady problems. They go to the heart of sovereignty. They go to the heart of culture. They go to the heart of "Land of the Free and Home of the Brave". They will not be solved centrally, yet neither will they be solved without central assistance. We have before us a set of bargains, bargains between the Devil and the Deep Blue Sea. And not to decide is to decide.

For me, I will take freedom over security and I will take security over

¹³ The Kerberos network authentication system is a product of MIT's Project Athena, which was the first effective and freely available mechanism for two parties who share a common administrative authority to establish mutual proof of identity. It is now so commonplace that it has become part of the woodwork, so to speak, and is very likely the single most commonly used program in the world. The author was privileged to be the manager in charge of all technical development for Athena, including Kerberos.

convenience, and I will do so because I know that a world without failure is a world without freedom. A world without the possibility of sin is a world without the possibility of righteousness. A world without the possibility of crime is a world where you cannot prove you are not a criminal. A technology that can give you everything you want is a technology that can take away everything that you have. At some point, in the near future, one of us security geeks will have to say that there comes a point at which safety is not safe.

I know full well that my views are neither pleasant nor fashionable, nor even attention getting enough to be dismissed. While time will tell if I am right, it would give no man pleasure then to say “I told you so.”