[nominal delivery draft, 6 August 2014]

Cybersecurity as Realpolitik

Dan Geer

Good morning and thank you for the invitation to speak with you today. The plaintext of this talk has been made available to the organizers. While I will not be taking questions today, you are welcome to contact me later and I will do what I can to reply. For simple clarity, let me repeat the abstract for this talk:

> Power exists to be used. Some wish for cyber safety, which they will not get. Others wish for cyber order, which they will not get. Some have the eye to discern cyber policies that are "the least worst thing;" may they fill the vacuum of wishful thinking.

There are three professions that beat their practitioners into a state of humility: farming, weather forecasting, and cyber security. I practice two of those, and, as such, let me assure you that the recommendations which follow are presented in all humility. Humility does not mean timidity. Rather, it means that when a strongly held belief is proven wrong, that the humble person changes their mind. I expect that my proposals will result in considerable push-back, and changing my mind may well follow. Though I will say it again later, this speech is me talking for myself.

As if it needed saying, cyber security is now a riveting concern, a top issue in many venues more important than this one. This is not to insult Black Hat; rather it is to note that every speaker, every writer, every practitioner in the field of cyber security who has wished that its topic, and us with it, were taken seriously has gotten their wish. Cyber security *is* being taken seriously, which, as you well know is not the same as being taken usefully, coherently, or lastingly. Whether we are talking about laws like the Digital Millenium Copyright Act or the Computer Fraud and Abuse Act, or the non-lawmaking but perhaps even more significant actions that the Executive agencies are undertaking, "we" and the cyber security issue have never been more at the forefront of policy. And you ain't seen nothing yet.

I wish that I could tell you that it is still possible for one person to hold the big picture firmly in their mind's eye, to track everything important that is going on in our field, to make few if any sins of omission. It is not possible; that phase passed sometime in the last six years. I have certainly tried to keep up but I would be less than candid if I were not to say that I know that I am not keeping up, not even keeping up with what is going on in my own country much less all countries. Not only has cybersecurity reached the highest levels of attention, it has spread into nearly every corner. If area is the product of height and width, then the footprint of cybersecurity has surpassed the grasp of any one of us.

The rate of technological change is certainly a part of it. When younger people ask my advice on what they should do or study to make a career in cyber security, I can only advise specialization. Those of us who were in the game early enough and who have managed to retain an over-arching generalist knowledge can't be replaced very easily because while absorbing most new information most of the time may have been possible when we began practice, no person starting from scratch can do that now. Serial specialization is now all that can be done in any practical way. Just looking at the Black Hat program will confirm that being really good at any one of the many topics presented here all but requires shutting out the demands of being good at any others.

Why does that matter? Speaking for myself, I am not interested in the advantages or disadvantages of some bit of technology unless I

can grasp how it is that that technology works. Whenever I see marketing material that tells me all the good things that adopting this or that technology makes possible, I remember what George Santayana said, that "Scepticism is the chastity of the intellect; it is shameful to give it up too soon, or to the first comer." I suspect that a majority of you have similar skepticism -- "It's magic!" is not the answer a security person will ever accept. By and large, I can tell *what* something is good for once I know *how* it works. Tell me how it works and then, but only then, tell me why you have chosen to use those particular mechanisms for the things you have chosen to use them for.

Part of my feeling stems from a long-held and well-substantiated belief that all cyber security technology is dual use. Perhaps dual use is a truism for any and all tools from the scalpel to the hammer to the gas can -- they can be used for good or ill -- but I know that dual use is inherent in cyber security tools. If your definition of "tool" is wide enough, I suggest that the cyber security tool-set favors offense these days. Chris Inglis, recently retired NSA Deputy Director, remarked that if we were to score cyber the way we score soccer, the tally would be 462-456 twenty minutes into the game,[CI] i.e., all offense. I will take his comment as confirming at the highest level not only the dual use nature of cybersecurity but also confirming that offense is where the innovations that only States can afford is going on.

Nevertheless, this essay is an outgrowth from, an extension of, that increasing importance of cybersecurity. With the humility of which I spoke, I do not claim that I have the last word. What I do claim is that when we speak about cybersecurity policy we are no longer engaging in some sort of parlor game. I claim that policy matters are now the most important matters, that once a topic area, like cybersecurity, becomes interlaced with nearly every aspect of life for nearly everybody, the outcome differential between good policies and bad policies broadens, and the ease of finding answers falls. As H.L. Mencken so trenchantly put it, "For every complex problem there is a solution that is clear, simple, and wrong."

The four verities of government are these:
. Most important ideas are unappealing
. Most appealing ideas are unimportant
. Not every problem has a good solution
. Every solution has side effects

This quartet of verities certainly applies to the interplay between cybersecurity and the affairs of daily living. Over my lifetime the public expectation of what government can and should do has spectacularly broadened from guaranteeing that you may engage in the "pursuit of happiness" to guaranteeing happiness in and of itself. The central dynamic internal to government is, and always has been, that the only way for either the Executive or the Legislature to control the many sub-units of government is by way of how much money they can hand out. Guaranteeing happiness has the same dynamic -- that the only tool government really has to achieve the outcome of everyone happy or everyone healthy or everyone safe at all times from things that go bump in the night is through the dispensing of money. This is true in foreign policy; one can reasonably argue that the United States' 2007 troop "surge" in Iraq did provide an improvement in safety. One can also argue that the work of those troops, some of whom gave what Abraham Lincoln called "the last full measure of devotion," was materially aided by the less publicized arrival of C-130s full of $100 bills with which to buy off potential combatants. Why should cybersecurity be any different?

Suppose, however, that surveillance becomes too cheap to meter, that is to say too cheap to limit through budgetary processes. Does that lessen the power of the Legislature more, or the power of the Executive more? I think that ever-cheaper surveillance substantially changes the balance of power in favor of the Executive and away

from the Legislature. While President Obama was referring to something else when he said "I've Got A Pen And I've Got A Phone," he was speaking to exactly this idea -- things that need no appropriations are outside the system of checks and balances. Is the ever-wider deployment of sensors in the name of cybersecurity actually contributing to our safety? Or is it destroying our safety in order to save it?

To be entirely clear by way of repetition, this essay is written by someone as his own opinion and not on behalf of anyone else. It is written without the supposed benefits of insider information; I hold no Clearance but am instead informed solely by way of open source intelligence. This path may be poised to grow easier; if the chief benefit of having a Clearance is to be able to see into the future a little further than those without one, then it must follow that as the pace of change accelerates the difference between how far can you see with a Clearance versus how far can you see without one will shrink.

There are, in other words, parallels between cybersecurity and the intelligence functions insofar as predicting the future has a strong role to play in preparing your defenses for probable attacks. As Dave Aitel has repeatedly pointed out, the hardest part of crafting good attack tools is testing them before deployment. Knowing what your tool will find, and how to cope with that, is surely harder than finding an exploitable flaw in and of itself. This, too, may grow in importance if the rigor of testing causes attackers to use some portion of the Internet at large as their test platform rather than whatever rig they can afford to set up in their own shop. If that is the case, then full scale traffic logs become an indispensable intelligence tool insofar as when an attack appears to be de novo those with full scale traffic logs may be in a position to answer the question "How long has this been going on?" The company Net Witness, now part of EMC, is one player who comes to mind in this regard, and there are others. This idea of looking backward for evidence that you didn't previously know enough to look for does certainly have intelligence value both for the Nation State and for the enterprise.

And there is a lot of traffic that we don't have a handle on. John Quarterman of Internet Perils makes a round number guess that 10% of Internet backbone traffic is unidentifiable as to protocol.[JQ] Whether he is off by a factor of two in either direction, that is still a lot of traffic. Arbor Networks estimates that perhaps 2% of all *identifiable* backbone traffic is, to use their term, "raw sewage."[AN] There are plenty of other estimates of this sort, of course. To my way of thinking, all such estimates continue to remind us that the end-to-end design of the Internet[SRC] was not some failure of design intellect but a brilliant avoidance of having to pick between the pitiful toy a completely safe Internet would have to be versus an Internet that was the ultimate tool of State control. In nothing else is it more apt to say that our choices are Freedom, Security, Convenience -- Choose Two.

Let me now turn to some policy proposals on a suite of pressing current topics. None of these proposals are fully formed, but as you know, those who don't play the game don't make the rules. These proposals are not in priority order, though some are more at odds with current practice than others and might, therefore, be said to be more pressing. There are more where these came from, but this talk has a time limit, and there is a meta-analysis at the end.

1. Mandatory reporting -- YES/Tiered

The United States Centers for Disease Control are respected the world around. When you really get down to it, three capabilities describe the CDC and why they are as effective as they are: (1) mandatory reporting of communicable diseases, (2) stored data and

the data analytic skill to distinguish a statistical anomaly from an outbreak, and (3) away teams to take charge of, say, the appearance of Ebola in Miami. Everything else is details. The most fundamental of these is the mandatory reporting of communicable diseases.

At the same time, we have well established rules about medical privacy. Those rules are helpful; when you check into the hospital there is a licensure-enforced, accountability-based, need-to-know regime that governs the handling of your data.[PHI] Most days, that is, but if you check in with Bubonic Plague or Typhus or Anthrax, you will have zero privacy as those are the "mandatory reporting of communicable disease conditions" as variously mandated not just by the CDC but by public health law in all fifty States.

So let me ask you, would it make sense, in a public health of the Internet way, to have a mandatory reporting regime for cybersecurity failures? Do you favor having to report cyber penetrations of your firm or of your household to some branch of government or some non-government entity? Should you face criminal charges if you fail to make such a report? Forty-eight States vigorously penalize failure to report sexual molestation of children.[SMC] The (US) Computer Fraud and Abuse Act[CFAA] defines a number of felonies related to computer penetrations, and the U.S. Code says that it is a crime to fail to report a felony of which you have knowledge.[USC] Is cybersecurity event data the kind of data around which you want to enforce mandatory reporting? Forty-six States require mandatory reporting of one class of cyber failures in the form of their data breach laws,[CSB] while the Verizon Data Breach Investigations Report[VDB] found, and the Index of Cyber Security[ICS] confirmed, that 70-80% of data breaches are discovered by unrelated third parties, not by the victim, meaning that the victim might never know if those who do the discovering were to keep quiet. If you discover a cyber attack, do you have an ethical obligation to report it? Should the law mandate that you fulfill such an obligation?

My answer to this set of questions is to mirror the CDC, that is for the force of law to require reporting of cybersecurity failures that are above some severity threshold that we have yet to negotiate. Below that threshold, I endorse the suggestion made in a piece two weeks ago, "Surviving on a Diet of Poisoned Fruit," by Richard Danzig where he made this policy proposal:[RD]

> Fund a data collection consortium that will illuminate the character and magnitude of cyber attacks against the U.S. private sector, using the model of voluntary reporting of near-miss incidents in aviation. Use this enterprise as well to help develop common terminology and metrics about cybersecurity.
>
> While regulatory requirements for aviation accident reporting are firmly established through the National Transportation Safety Board, there are no requirements for reporting the vastly more numerous and often no less informative near misses. Efforts to establish such requirements inevitably generate resistance: Airlines would not welcome more regulation and fear the reputational and perhaps legal consequences of data visibility; moreover, near accidents are intrinsically more ambiguous than accidents. An alternative path was forged in 2007 when MITRE, a government contractor, established an Aviation Safety Information Analysis and Sharing (ASIAS) system receiving near-miss data and providing anonymized safety, benchmarking and proposed improvement reports to a small number of initially participating airlines and the Federal Aviation Administration (FAA).

Today, 44 airlines participate in that program voluntarily. The combination of a mandatory CDC model for above-threshold cyber events and a voluntary ASIAS model for below-threshold events is what I recommend. This leaves a great deal of thinking still to be done; diseases are treated by professionals, but malware infections are treated by amateurs. Diseases spread within jurisdictions

before they become global, but malware is global from the get-go. Diseases have predictable behaviors, but malware comes from sentient opponents. Don't think this proposal is an easy one or one without side effects.

2. Net neutrality -- CHOICE

There is considerable irony in the Federal Communications Commission classifying the Internet as an information service and not as a communications service insofar as while that may have been a gambit to relieve ISPs of telephone-era regulation, the value of the Internet is ever more the bits it carries, not the carriage of those bits. The FCC decisions are both several and now old, the FCC classified cable as an information service in 2002, classified DSL as an information service in 2005, classified wireless broadband as an information service in 2007, and classified broadband over power lines as an information service in 2008. A decision by the D.C. Circuit Court of Appeals on this very point appeared earlier this year,[VZF] but settled little. The question remains, is the Internet a telecommunications service or an information service?

I've nothing new to say to you about the facts, the near-facts, nor the lying distortions inherent in the debate regarding network neutrality so far or still to come. What I can say is that network neutrality is no panacea nor is it anathema; peoples' tastes vary and so do corporations'. What I can say is that the varied tastes need to be reflected in constrained choice rather than the idea that the FTC or some other agency can assure happiness if and only if it, rather than corporations or individuals, does the choosing. Channeling for Doctor Seuss, if I ran the zoo I'd call up the ISPs and say this:

> Hello, Uncle Sam here.
>
> You can charge whatever you like based on the contents of what you are carrying, but you are responsible for that content if it is hurtful; inspecting brings with it a responsibility for what you learn.
> -or-
> You can enjoy common carrier protections at all times, but you can neither inspect nor act on the contents of what you are carrying and can only charge for carriage itself. Bits are bits.
>
> Choose wisely. No refunds or exchanges at this window.

In other words, ISPs get the one or the other; they do not get both. The FCC gets some heartache but also a natural experiment in whether those who choose common carrier status turn out differently than those who choose multi-tiered service grades with liability exposure. We already have a lot of precedent and law in this space. The United States Postal Service's term of art, "sealed against inspection," is reserved for items on which the highest postage rates are charged; is that also worth stirring into the mix?

As a side comment, I might add that it was in Seuss' book _If I Ran the Zoo_ that the word "nerd" first appeared in English. If Black Hat doesn't yet have an official book, I'd suggest this one.

3. Source code liability -- CHOICE

Nat Howard said that "Security will always be exactly as bad as it can possibly be while allowing everything to still function,"[NH] but with each passing day, that "and still function" clause requires a higher standard. As Ken Thompson told us in his Turing Award lecture, there is no technical escape;[KT] in strict mathematical terms you neither trust a program nor a house unless you created it 100% yourself, but in reality most of us will trust a house built by a suitably skilled professional, usually we will trust it more

than one we had built ourselves, and this even if we have never met the builder, or even if he is long since dead.

The reason for this trust is that shoddy building work has had that crucial "or else ..." clause for more than 3700 years:

> If a builder builds a house for someone, and does not construct it properly, and the house which he built falls in and kills its owner, then the builder shall be put to death.
> -- Code of Hammurabi, approx 1750 B.C.

Today the relevant legal concept is "product liability" and the fundamental formula is "If you make money selling something, then you better do it well, or you will be held responsible for the trouble it causes." For better or poorer, the only two products not covered by product liability today are religion and software, and software should not escape for much longer. Poul-Henning Kamp and I have a strawman proposal for how software liability regulation could be structured.

```
.......................
0. Consult criminal code to see if damage caused was due to intent
   or willfulness.
.......................
```

We are only trying to assign liability for unintentionally caused damage, whether that's sloppy coding, insufficient testing, cost cutting, incomplete documentation, or just plain incompetence. Clause zero moves any kind of intentionally inflicted damage out of scope. That is for your criminal code to deal with, and most already do.

```
.......................
1. If you deliver your software with complete and buildable source
   code and a license that allows disabling any functionality or
   code the licensee decides, your liability is limited to a refund.
.......................
```

Clause one is how to avoid liability: Make it possible for your users to inspect and chop out any and all bits of your software they do not trust or want to run. That includes a bill of materials ("Library ABC comes from XYZ") so that trust has some basis, paralleling why there are ingredient lists on processed foods.

The word "disabling" is chosen very carefully: You do not need to give permission to change or modify how the program works, only to disable the parts of it that the licensee does not want or trust. Liability is limited even if the licensee never actually looks at the source code; as long has he has received it, you (as maker) are off the hook. All your other copyrights are still yours to control, and your license can contain any language and restriction you care for, leaving the situation unchanged with respect to hardware-locking, confidentiality, secrets, software piracy, magic numbers, etc.

Free and Open Source Software (FOSS) is obviously covered by this clause which leaves its situation unchanged.

```
.......................
2. In any other case, you are liable for whatever damage your
   software causes when it is used normally.
.......................
```

If you do not want to accept the information sharing in Clause 1, you fall under Clause 2, and must live with normal product liability, just like manufactures of cars, blenders, chain-saws and hot coffee.

How dire the consequences, and what constitutes "used normally" is for your legislature and courts to decide, but let us put up a strawman example:

> A sales-person from one of your long time vendors visits and delivers new product documentation on a USB key, you plug the USB key into your computer and copy the files onto the computer.

This is "used normally" and it should never cause your computer to become part of a botnet, transmit your credit card number to Elbonia, or copy all your design documents to the vendor. If it does, your computer's operating system is defective.

The majority of today's commercial software would fall under Clause 2 and software houses need a reasonable chance to clean up their act or to move under Clause 1, so a sunrise period is required. But no longer than five years -- we are trying to solve a dire computer security problem here.

And that is it really: Either software houses deliver quality and back it up with product liability, or they will have to let their users protect themselves. The current situation -- users can't see whether they need to protect themselves and have no recourse to being unprotected -- cannot go on. We prefer self-protection (and fast recovery), but other's mileage may differ.

Would it work? In the long run, absolutely yes. In the short run, it is pretty certain that there will be some nasty surprises as badly constructed source code gets a wider airing. The FOSS community will, in parallel, have to be clear about the level of care they have taken, and their build environments as well as their source code will have to be kept available indefinitely.

The software houses will yell bloody murder the minute legislation like this is introduced, and any pundit and lobbyist they can afford will spew their dire predictions that "This law will mean the end of computing as we know it!"

To which our considered answer will be:

> Yes, please! That was exactly the idea.

4. Strike back -- LIMITED YES

I suspect that a fair number of you have, in fact, struck back at some attacker somewhere or, at least, done targeting research even if you didn't pull the trigger. I'd trust many of you to identify targets carefully enough to minimize collateral damage, but what we are talking about here is the cyber equivalent of the smart bomb. As I implied earlier, cyber smart bombs are what the national laboratories of several countries are furiously working on. In that sense, you do know what is happening behind the curtain, and you know how hard that targeting really is because you know how hard attribution -- real attribution -- really is.

The issue is shared infrastructure, and that issue is not going away. There are some entities that can operate globally and strike back effectively, Microsoft and the FBI teaming up on the GameOver Zeus trojan for example,[GOZ] but that's an expensive therapy in limited supply that can only be applied to the most damaging malware. Nevertheless, that is the therapy we have. Smaller entities cannot act globally nor can they act in certain ways without pairing with national agencies. That can, and must, go on, but I don't see how the individual or the smaller entity can shoot back. All I see is for the individual or the smaller entity to put all their effort into having fast recovery.

5. Fall backs and resiliency -- TOO COMPLICATED FOR ONE POLICY

There has always been a lot of talk about what to do when failure

is unacceptable and yet failure is inevitable. Heretofore, almost anything that has come to be seen as essential to the public gets some sort of performance standard imposed upon it, electricity and water, say. But let's talk about software.

For one example, a commonly voiced desire for cryptographic protocols is "algorithm agility," the ability to swap from one cryptographic algorithm to another if and when the first one becomes unsafe. The security benefit of such a swap is not what you turn on but what you turn off. For that to be possible, a second algorithm has to already be in place, but that means that the second algorithm had to be designed in at the outset and at both ends, with a way to choose between them such that either end of the proposed connection can force a change-over to the alternate algorithm. One might argue that implementing algorithm agility actually means a single, more complex algorithm. Or maybe what you want is two algorithms where you always use both, such as when you encrypt with one algorithm and super-encrypt with another so that the failure of one has no practical effect on security and nothing has to change.

I say all that just to demonstrate that it is not always simple to have a pre-deployed fallback should something break, that design willpower alone is not enough. So perhaps mandating pre-deployed fallbacks is a bad idea entirely. Perhaps what is needed is a way to reach out and upgrade the endpoints when the time of necessity comes. But today, or real soon now, most of the places needing a remote management interface through which you can remotely upgrade the endpoints are embedded hardware. So let me ask a question, should or should not an embedded system be required to have a remote management interface? If it does not, then a late discovered flaw cannot be fixed without visiting all the embedded systems -- which is likely to be infeasible because some you will be unable to find, some will be where you cannot again go, and there will be too many of them in any case. If it does have a remote management interface, the opponent of skill will focus on that and, once a break is achieved, will use those self-same management functions to ensure that not only does he retain control over the long interval but, as well, you will be unlikely to know that he is there.

Perhaps what is needed is for embedded systems to be more like humans, and I most assuredly do not mean artificially intelligent. By "more like humans" I mean this: Embedded systems, if having no remote management interface and thus out of reach, are a life form and as the purpose of life is to end, an embedded system without a remote management interface must be so designed as to be certain to die no later than some fixed time. Conversely, an embedded system with a remote management interface must be sufficiently self-protecting that it is capable of refusing a command. Inevitable death and purposive resistance are two aspects of the human condition we need to replicate, not somehow imagine that to overcome them is to improve the future.

Lest some of you think this is all so much picayune, tendentious, academic perfectionist posturing, let me inform some of you and remind the others that it is entirely possible to deny the Internet to a large fraction of its users. Home routers have drivers and operating systems that are binary blobs amounting to snapshots of the state of Linux plus the lowest end commodity chips that were extant at the time of the router's design. Linux has moved on. Device drivers have moved on. Samba has moved on. Chipsets have moved on. But what is sold at Best Buy or the like is remarkably cheap and remarkably old. With certainty born of long engineering experience, I assert that those manufacturers can no longer build their deployed software blobs from source. If, as my colleague Jim Gettys has laboriously measured, the average age of the code base on those ubiquitous low-end routers is 4-5 years,[JG] then you can be assured that the CVE catalog lists numerous methods of attacking those operating systems and device drivers remotely.[CV] If I can commandeer them remotely, then I can build a botnet that is on the

outside of the home network. It need not ever put a single packet through the firewall, it need never be detectible by any means whatsoever from the interior of the network it serves, but it is most assuredly a latent weapon, one that can be staged to whatever level of prevalence I desire before I ask it to do more. All I need is to include in my exploit a way to signal that device to do three things: stop processing anything it henceforth receives, start flooding the network with a broadcast signal that causes other peers to do the same, and zero the on-board firmware thus preventing reboot for all time. Now the only way to recover is to unplug all the devices, throw them in the dumpster, and install new ones -- but aren't the new ones likely to have the same kind of vulnerability spectrum in CVE that made this possible in the first place? Of course they do, so this is not a quick trip to the big box store but rather flushing the entire design space and pipeline inventory of every maker of home routers. There appears to be an event at DefCon around this very issue.[SOHO]

Resiliency is an area where no one policy can be sufficient, so I've suggested a trio of baby steps: embedded systems cannot be immortal if they have no remote management interface, embedded systems must have a remote management interface if they are to be immortal, and swap-over is preferable to swap-out when it comes to data protection.

6. Vulnerability finding -- HEGEMONY

Vulnerability finding is a job. It has been a job for something like eight years now, give or take. For a good long while, you could do vulnerability finding as a hobby and get paid in bragging rights, but finding vulnerabilities got to be too hard to do as a hobby in your spare time -- you needed to work it like a job and get paid like a job. This was the result of hard work on the part of the software suppliers including the suppliers of operating systems, but as the last of the four verities of government says, every solution has side effects. In this case, the side effect is that once vulnerability finding became a job and stopped being a bragging-rights hobby, those finding the vulnerabilities stopped sharing. If you are finding vulns for fun and fame, then the minute you find a good one you'll let everybody know just to prevent someone else finding it and beating you to the punch. If you are doing it for profit, then you don't share. That's where the side effect is -- once coin-operated vuln finders won't share, the percentage of all attacks that are zero-day attacks must rise, and it has.

In a May article in The Atlantic,[BS] Bruce Schneier asked a cogent first-principles question: Are vulnerabilities in software dense or sparse? If they are sparse, then every one you find and fix meaningfully lowers the number of avenues of attack that are extant. If they are dense, then finding and fixing one more is essentially irrelevant to security and a waste of the resources spent finding it. Six-take-away-one is a 15% improvement. Six-thousand-take-away-one has no detectable value.

If a couple of Texas brothers could corner the world silver market,[HB] there is no doubt that the U.S. Government could openly corner the world vulnerability market, that is we buy them all and we make them all public. Simply announce "Show us a competing bid, and we'll give you 10x." Sure, there are some who will say "I hate Americans; I sell only to Ukrainians," but because vulnerability finding is increasingly automation-assisted, the seller who won't sell to the Americans knows that his vulns can be rediscovered in due course by someone who *will* sell to the Americans who will tell everybody, thus his need to sell his product before it outdates is irresistible.

This strategy's usefulness comes from two side effects: (1) that by overpaying we enlarge the talent pool of vulnerability finders

and (2) that by making public every single vuln the USG buys we devalue them. Put differently, by overpaying we increase the rate of vuln finding, while by showing everyone what it is that we bought we zero out whatever stockpile of cyber weapons our adversaries have. We don't need intelligence on what weapons our adversaries have if we have something close to a complete inventory of the world's vulns and have shared that with all the affected software suppliers. But this begs Schneier's question: Are vulnerabilities sparse or dense? If they are sparse or even merely numerous, then cornering the market wins in due course. If they are dense, then all we would end up doing is increasing costs both to software suppliers now obligated to repair all the vulns a growing army of vuln researchers can find and to taxpayers. I believe that vulns are scarce enough for this to work and,, therefore I believe that cornering the market is the cheapest win we will ever get.

Let me note, however, that my colleagues in static analysis report that they regularly see web applications greater than 2GB in size and with 20,000 variables. Such web apps can only have been written by machine and, therefore, the vulns found in them were also written by machine. Machine-powered vuln creation might change my analysis though I can't yet say in what direction.

7. Right to be forgotten -- YES

I've spoken elsewhere about how we are all intelligence agents now, collecting on each other on behalf of various overlords.[RSA] There are so many technologies now that power observation and identification of the individual at a distance. They may not yet be in your pocket or on your dashboard or embedded in all your smoke detectors, but that is only a matter of time. Your digital exhaust is unique hence it identifies. Pooling everyone's digital exhaust also characterizes how you differ from normal. Privacy used to be proportional to that which it is impossible to observe or that which can be observed but not identified. No more -- what is today observable and identifiable kills both privacy as impossible-to-observe and privacy as impossible-to-identify, so what might be an alternative? If you are an optimist or an apparatchik, then your answer will tend toward rules of data procedure administered by a government you trust or control. If you are a pessimist or a hacker/maker, then your answer will tend towards the operational, and your definition of a state of privacy will be my definition: the effective capacity to misrepresent yourself.

Misrepresentation is using disinformation to frustrate data fusion on the part of whomever it is that is watching you. Some of it can be low-tech, such as misrepresentation by paying your therapist in cash under an assumed name. Misrepresentation means arming yourself not at Walmart but in living rooms. Misrepresentation means swapping affinity cards at random with like-minded folks. Misrepresentation means keeping an inventory of misconfigured webservers to proxy through. Misrepresentation means putting a motor-generator between you and the Smart Grid. Misrepresentation means using Tor for no reason at all. Misrepresentation means hiding in plain sight when there is nowhere else to hide. Misrepresentation means having not one digital identity that you cherish, burnish, and protect, but having as many as you can. Your fused identity is not a question unless you work to make it be. Lest you think that this is a problem statement for the random paranoid individual alone, let me tell you that in the big-I Intelligence trade, crafting good cover is getting harder and harder and for the exact same reasons: misrepresentation is getting harder and harder. If I was running field operations, I would not try to fabricate a complete digital identity, I'd "borrow" the identity of someone who had the characteristics that I needed for the case at hand.

The Obama administration's issuance of a National Strategy for Trusted Identities in Cyberspace[NS] is case-in-point; it "calls

for the development of interoperable technology standards and policies -- an 'Identity Ecosystem' -- where individuals, organizations, and underlying infrastructure -- such as routers and servers -- can be authoritatively authenticated." If you can trust a digital identity, that is because it can't be faked. Why does the government care about this? It cares because it wants to digitally deliver government services and it wants attribution. Is having a non-fake-able digital identity for government services worth the registration of your remaining secrets with that government? Is there any real difference between a system that permits easy, secure, identity-based services and a surveillance system? Do you trust those who hold surveillance data on you over the long haul by which I mean the indefinite retention of transactional data between government services and you, the individual required to proffer a non-fake-able identity to engage in those transactions? Assuming this spreads well beyond the public sector, which is its designers' intent, do you want this everywhere? If you are building authentication systems today, then you are already playing ball in this league. If you are using authentication systems today, then you are subject to the pending design decisions of people who are themselves playing ball in this league.

After a good amount of waffling, I conclude that a unitary, unfakeable digital identity is no bargain and that I don't want one. I want to choose whether to misrepresent myself. I may rarely use that, but it is my right to do so. If that right vanishes into the panopticon, I have lost something and, in my view, gained next to nothing. In that regard, and acknowledging that it is a baby step, I conclude that the EU's "Right to be Forgotten" is both appropriate and advantageous though it does not go far enough. Being forgotten is consistent with moving to a new town to start over, to changing your name, to a definition of privacy that turns on whether you do or do not retain the effective capacity to misrepresent yourself, a right which I will remind you is routinely granted but to those who have especially helped governmental causes (witness protection, e.g.). A right to be forgotten is the only check on the tidal wave of observability that a ubiquitous sensor fabric is birthing now, observability that changes the very quality of what "in public" means. Entities that block deep-linking to their web resources are neutralizing indexability. Governments of all stripes, irretrievably balkanizing the Internet through the self-same vehicle of indexing controls, are claiming that a right to do so is inherently theirs. The only democratizing brake on this runaway train is for individuals to be able, in their own small way, to do the same as do other entities. I find it notably ironic that The Guardian newspaper's championing of Edward Snowden's revelations about privacy loss is paired with the same paper's editorializing that "No one has a right to be forgotten."[GRF] Au contraire, madames et monsieurs, they most assuredly do.

8. Internet voting -- NO

Motivated & expert opponents are very nearly undefendable against. People like us here know that, which is why it is natural for people like us here to oppose voting over the Internet. The National Center for Policy Analysis thinks online voting is a bad idea. NIST thinks online voting is a bad idea. With Pamela Smith, Bruce McConnell editorialized in the pages of the Wall Street Journal[BMC] that online voting is a bad idea. The fact that we here have near universal disdain for the idea has not seemed to change much policy.

Now it is always true that a thorough security analysis will get much less attention than a juicy conspiracy theory even if both lead to the same conclusion. How do we explain this? If I knew that, then I would commence to explaining, but we may not need to explain it if the integrity of some election is put at question by events. I'd like to think that we don't need carnage to motivate a re-think, but perhaps we do. If we do need carnage, then may its

coming be sooner rather than later.

9. Abandonment -- CERTAINTY OF CONSEQUENCES

If I abandon a car on the street, then eventually someone will be able to claim title. If I abandon a bank account, then the State will eventually seize it. If I abandon real estate by failing to remedy a trespass, then in the fullness of time adverse possession takes over. If I don't use my trademark, then my rights go over to those who use what was and could have remained mine. If I abandon my spouse and/or children, then everyone is taxed to remedy my actions. If I abandon a patent application, then after a date certain the teaching that it proposes passes over to the rest of you. If I abandon my hold on the confidentiality of data such as by publishing it, then that data passes over to the commonweal not to return. If I abandon my storage locker, then it will be lost to me and may end up on reality TV. The list goes on.

Apple computers running 10.5 or less get no updates (comprising a significant fraction of the installed base). Any Microsoft computer running XP gets no updates (likewise comprising a significant fraction of the installed base). The end of security updates follows abandonment. It is certainly ironic that freshly pirated copies of Windows get security updates when older versions bought legitimately do not.

Stating what to me is the obvious policy stance, if Company X abandons a code base, then that code base must be open sourced. Irrespective of security issues, many is the time that a bit of software I use has gone missing because its maker killed it. But with respect to security, some constellation of {I,we,you,they} are willing and able to provide security patches or workarounds as time and evil require.

Would the public interest not be served by a conversion to open source for abandoned code bases? I believe it would. But wait, you say, isn't purchased software on a general purpose computer a thing of the past? Isn't the future all about auto-updated smartphone clients transacting over armored private (carrier) networks to auto-updated cloud services? Maybe; maybe not. If the two major desktop suppliers update only half of today's desktops, then what percentage will they update tomorrow?

If you say "Make them try harder!," then the legalistic, regulatory position is your position, and the ACLU is already trying that route. If smartphone auto-update becomes a condition of merchantability and your smartphone holds the keying material that undeniably says that its user is you, then how long before a FISA court orders a special auto-update to *your* phone for evidence gathering?

If you say "But we already know what they're going to do, don't we?," then the question is what about the abandoned code bases. Open-sourcing abandoned code bases is the worst option, except for all the others. But if seizing an abandoned code base is too big a stretch for you before breakfast, then start with a Public Key Infrastructure Certifying Authority that goes bankrupt and ask "Who gets the keys?"

10. Convergence -- DEFAULT DENY

Let me ask you a question: Are the physical and digital worlds one world or two? Are cyberspace and meatspace converging or diverging over time? I conclude that they are converging, but if they are converging, then is cyberspace looking more and more like meatspace or is meatspace looking more and more like cyberspace? That is not so clear.

Possibility #1 is that cyberspace becomes more and more like meatspace, ergo the re-creation of borders and jurisdictional boundaries is what happens next. Possibility #2 is that meatspace becomes more and more like cyberspace, ergo jurisdictional boundaries grow increasingly irrelevant and something akin to one-world technocratic government more or less follows. The former is heterogeneous, the latter is the monoculture of a single nation-state. As we all know, resiliency and freedom obtain solely from heterogeneity, so converging meatspace to cyberspace is the unfavorable outcome, but what can be done about it?

At the end of last year, the Pew Research Center invited 12,000 "experts" to answer a single Yes/No question:

> By 2025 will there be significant changes for the worse and hindrances to the ways in which people get and share content online compared with the way globally networked people can operate online today?[PEW]

Of the 12,000 invited, some 1,400 did answer. Putting aside whatever selection bias may be reflected in who chose to answer and who did not, Pew found four themes dominated respondent comments:

> 1) Actions by nation-states to maintain security and political control will lead to more blocking, filtering, segmentation, and balkanization of the Internet.
>
> 2) Trust will evaporate in the wake of revelations about government and corporate surveillance and likely greater surveillance in the future.
>
> 3) Commercial pressures affecting everything from Internet architecture to the flow of information will endanger the open structure of online life.
>
> 4) Efforts to fix the "too much information" problem might over-compensate and actually thwart content sharing.

My colleague Rob Lemos mapped Pew's themes to the two alternative futures I mentioned above,[RL] saying that "If cyberspace converges to our physical reality, then we will have balkanization and commercial efforts to artificially create information monopolies, while if the physical world goes toward digital space, then we have greater surveillance, the erosion of trust, much information leakage, and the reaction to that leakage." More crucially, Lemos also observed that the growth of technology has greatly increased personal power:

> The impact that a single person can have on society has significantly increased over time to where a single individual can have a devastating effect. The natural reaction for government is to become more invasive {possibility #2 above} to better defend its monoculture, or more separate {possibility #1 above} to firewall threats from one another. Because threats and kinetic impacts can increasingly travel through the digital realm, they necessitate that the policy and legal frameworks of the digital and physical world converge.

In other words, Lemos argues that convergence is an inevitable consequence of the very power of cyberspace in and of itself. I don't argue with Lemos' idea that increasingly powerful, location independent technology in the hands of the many will tend to force changes in the distribution of power. In fact, that is the central theme of this essay -- that the power that is growing in the net, per se, will soon surpass the ability of our existing institutions to modify it in any meaningful way, so either the net must be broken up into governable chunks or the net becomes government.

It seems to me that the leverage here favors cyberspace whenever

and wherever we give cyberspace a monopoly position, which we are doing that blindly and often. In the last couple of years, I've found that institutions that I more or less must use -- my 401(k) custodian, the Government Accounting Office's accounts payable department, the payroll service my employer outsources to, etc. -- no longer accept paper letter instructions, they each only accept digital delivery of such instructions. This means that each of them has created a critical dependence on an Internet swarming with men in the middle and, which is more, they have doubtlessly given up their own ability to fall back to what worked for a century before.

It is that giving up of alternative means that really defines what convergence is and does. It is said that all civil wars are about on whose terms re-unification will occur. I would argue that we are in, to coin a phrase, a Cold Civil War to determine on whose terms convergence occurs. Everything in meatspace we give over to cyberspace replaces dependencies that are local and manageable with dependencies that are certainly not local and I would argue much less manageable because they are much less secure. I say that because the root cause of risk is dependence, and most especially dependence on expectations of system state. I say "much less secure" because one is secure, that is to say that one is in a state of security, if and only if there can be no unmitigatable surprises. The more we put on the Internet, the broader and unmitigatable any surprises become.

This line of thought is beginning to sink in. Let me quote from a Bloomberg article a month ago:[CWC]

> Wall Street's biggest trade group has proposed a government-industry cyber war council to stave off terrorist attacks that could trigger financial panic by temporarily wiping out account balances, according to an internal document.
>
> The proposal by the Securities Industry and Financial Markets Association calls for a committee of executives and deputy-level representatives from at least eight U.S. agencies including the Treasury Department, the National Security Agency and the Department of Homeland Security, all led by a senior White House official.
>
> The document sketches an unusually frank and pessimistic view by the industry of its readiness for attacks wielded by nation-states or terrorist groups that aim to "destroy data and machines." It says the concerns are "compounded by the dependence of financial institutions on the electric grid," which is also vulnerable to physical and cyber attack.

So here you have the biggest financial firms saying that their dependencies are no longer manageable, and that the State's monopoly on the use of force must be brought to bear. What they are talking about is that they have no way to mitigate the risk of common mode failure.

To repeat, risk is a consequence of dependence. Because of shared dependence, aggregate societal dependence on the Internet is not estimable. If dependencies are not estimable, they will be underestimated. If they are underestimated, they will not be made secure over the long run, only over the short. As the risks become increasingly unlikely to appear, the interval between events will grow longer. As the latency between events grows, the assumption that safety has been achieved will also grow, thus fueling increased dependence in what is now a positive feedback loop. Accommodating old methods and Internet rejectionists preserves alternate, less complex, more durable means and therefore bounds dependence. Bounding dependence is *the* core of rational risk management.

If we don't bound dependence, we invite common mode failure. In the language of statistics, common mode failure comes exactly from

under-appreciated mutual dependence. Quoting [NIST]:

> [R]edundancy is the provision of functional capabilities that would be unnecessary in a fault-free environment. Redundancy is necessary, but not sufficient for fault tolerance... System failures occur when faults propagate to the outer boundary of the system. The goal of fault tolerance is to intercept the propagation of faults so that failure does not occur, usually by substituting redundant functions for functions affected by a particular fault. Occasionally, a fault may affect enough redundant functions that it is not possible to reliably select a non-faulty result, and the system will sustain a common-mode failure. A common-mode failure results from a single fault (or fault set). Computer systems are vulnerable to common-mode resource failures if they rely on a single source of power, cooling, or I/O. A more insidious source of common-mode failures is a design fault that causes redundant copies of the same software process to fail under identical conditions.

That last part -- that "A more insidious source of common-mode failures is a design fault that causes redundant copies of the same software process to fail under identical conditions" -- is exactly that which can be masked by complexity precisely because complexity ensures under-appreciated mutual dependence.

In sum, as a matter of policy everything that is officially categorized as a critical infrastructure must conclusively show how it can operate in the absence of the Internet. The 2008 financial crisis proved that we can build systems more complex than we can operate, the best policy counter to which has been the system of "stress tests" thereafter administered to the banks. We need other kinds of stress tests even more.

Conclusion

I titled this talk "Cybersecurity as Realpolitik." Realpolitik means, in the words of British historian E. H. Carr, that what is successful is right and what is unsuccessful is wrong, that there is no moral dimension in how the world is, and that attempting to govern based on principles cannot succeed. Realpolitik is at once atheistic and anti-utopian.

I find that distasteful and, it seems, that in governing my own life I daily give up power advantage for principle. At the same time, having principles such as "Might does not make right" may well be a failing on my part and, by extension, a failing on the part of those who govern according to principle. Cybersecurity as we describe it in our mailing lists, on our blogs, at our cons, and so forth is rich in principles and utopian desiderata, all the while we have opponents at all levels and probably always will for whom principle matters little but power matters a lot. As Thomas Ray said, "Every successful system accumulates parasites" and the Internet plus every widely popular application on it has parasites. For some observers, parasites and worse are just a cost of doing business. For other observers, design which encourages bad outcomes is an affront that must be fixed. It is realism and realism alone that remains when all else fails.

Political realism of the sort I am talking about is based on four premises:

. The international system is anarchic
. States are the most important actors
. All states within the system are unitary, rational actors
. The primary concern of all states is survival

This is likewise the realism of the cybersecurity situation in a global Internet. It is anarchic, and states have become the most

important actors. States' investment in offensive cyber is entirely about survival in such a world. States are driven to this by the dual, simultaneous expansion of what is possible and what their citizens choose to depend on.

The late Peter Bernstein, perhaps the world's foremost thinker on the topic, defined "risk" as "more things can happen than will."[PB] With technologic advance accelerating, "more things can happen than will" takes on a particularly ominous quality if your job is to ensure your citizens' survival in an anarchy where, daily, ever more things can happen than will. Realpolitik would say that under such circumstances, defense becomes irrelevant. What is relevant is either (1) offense or (2) getting out of the line of fire altogether. States that are investing in offense are being entirely rational and are likely to survive. Those of us who are backing out our remaining dependencies on digital goods and services are being entirely rational and are likely to survive. The masses who quickly depend on every new thing are effectively risk seeking, and even if they do not themselves know it, the States which own them know, which explains why every State now does to its own citizens what once States only did to officials in competing regimes.

You have politely listened to a series of "get off the dime" policy proposals around mandatory reporting, net neutrality, source code liability, strike back, fall backs, resiliency, vulnerability finding, the right to be forgotten, Internet voting, abandonment, and convergence, all by one guy that no one ever elected. I thank you, friends and countrymen, for lending me your ears. But I shall be happier still if some one or several of you find the articulateness that overcomes the dynamic which we now inhabit, namely that if what is successful is right and what is unsuccessful is wrong, the observable allocation of success and of failure is utterly disconnected from the technical facts of cybersecurity as we know them here. In the end, reality always wins, and the reality of technical facts has more staying power than the reality of market share or utopian enthusiasm.

Nevertheless, cybersecurity is all about power and only power. Realpolitik says that what cybersecurity works is right and what cybersecurity does not work is wrong and Realpolitik thus resonates with Howard's "Security will always be exactly as bad as it can possibly be while allowing everything to still function." Realpolitik says that offense routinely beating defense is right, and imagining otherwise is wrong, that those whose offense wins are right while those whose defense loses are wrong. Realpolitik says that offense's superiority means that it a utopian fantasy to believe that information can be protected from leakage, and so the counter-offense of disinformation is what we must deploy in return. Realpolitik says that sentient opponents have always been a fact of life, but never before have they been location independent and never before have they been able to recruit mercenaries who will work for free. Realpolitik says that attribution is impossible unless we deploy a unitary surveillance state.

I have long preferred to hire security people who are, more than anything else, sadder but wiser. They, and only they, know that most of what commercially succeeds succeeds only so long as attackers do not give it their attention while what commercially fails fails not because it didn't work but because it wasn't cheap or easy or sexy enough to try. Their glasses are not rose-colored; they are spattered with Realpolitik. Sadder but wiser hires, however, come only from people who have experienced private tragedies, not global ones. There are no people sadder but wiser about the scale and scope of the attack surface you get when you connect everything to everything and give up your prior ability to do without. Until such people are available, I will busy myself with reducing my dependence on, and thus my risk exposure to, the digital world even though that will be mistaken for curmudgeonly nostalgia. Call that misrepresentation, if you like.

There is never enough time. Thank you for yours.

= = = = = = = = = = = = = = =

To the reader, see also: "algorithmic regulation"

==========

[CI] Chris Inglis, confirmed by personal communication

[JQ] John Quarterman, personal communication

[AN] "2% of Internet Traffic Raw Sewage"
www.arbornetworks.com/asert/2008/03/2-of-internet-traffic-raw-sewage

[SRC] "End-to-End Arguments in System Design"
web.mit.edu/Saltzer/www/publications/endtoend/endtoend.pdf

[PHI] Protected Health Information, abbreviated PHI, as defined by
Section 1171 of Part C of Subtitle F of Public Law 104-191, "The
Health Insurance Portability and Accountability Act of 1996," also
known as HIPAA

[SMC] "Penalties for failure to report and false reporting of child
abuse and neglect," US Dept of Health and Human Services, Children's
Bureau, Child Welfare Information Gateway

[CFAA] U.S. Code, Title 18, Part I, Chapter 47, Section 1030
www.law.cornell.edu/uscode/text/18/1030

[USC] U.S. Code, Title 18, Part I, Chapter 1, Section 4
www.law.cornell.edu/uscode/text/18/4

[CSB] Security Breach Information Act
www.leginfo.ca.gov/pub/01-02/bill/sen/sb_1351-1400/sb_1386_bill_20020926_chaptered.pdf

[VDB] Verizon Data Breach Investigations Report
www.verizonenterprise.com/DBIR

[ICS] Index of Cyber Security
www.cybersecurityindex.org

[RD] "Surviving on a Diet of Poisoned Fruit; Reducing the National
Security Risks of America's Cyber Dependencies"
www.cnas.org/surviving-diet-poisoned-fruit

[VZF] Verizon v. FCC, 740 F.3d 623 (D.C. Cir. 2014)
www.cadc.uscourts.gov/internet/opinions.nsf/3AF8B4D938CDEEA685257C6000532062/$file/11-1355-
1474943.pdf

[NH] Nat Howard at USENIX 2000, per Marcus Ranum

[KT] Ken Thompson, "Reflections on Trusting Trust," 1984

[GOZ] "Microsoft and FBI team up to take down GameOver Zeus botnet"
www.techradar.com/us/news/internet/web/microsoft-and-fbi-team-up-to-take-down-gameover-zeus-
botnet-1251609

[JG] Gettys J, former VP Software, One Laptop Per Child, personal
communication

[CV] Common Vulnerabilities and Exposures, cve.mitre.org/cve

[SOHO] SOHOpelessly Broken, www.sohopelesslybroken.com

[BS] "Should U.S. Hackers Fix Cybersecurity Holes or Exploit Them?"
www.theatlantic.com/technology/archive/2014/05/should-hackers-fix-cybersecurity-holes-or-exploit-

them/371197

[HB] "Hunt Brothers Corner Silver Market"
web.archive.org/web/20060118031501/http://www.wallstraits.com/main/viewarticle.php?id=1298

[RSA] "We Are All Intelligence Agents Now"
geer.tinho.net/geer.rsa.28ii14.txt

[NS] National Strategy for Trusted Identities in Cyberspace,
www.nist.gov/nstic

[GRF] "The Right to Be Forgotten Will Turn the Internet into a Work
of Fiction,"
www.theguardian.com/commentisfree/2014/jul/06/right-to-be-forgotten-internet-work-of-fiction-david-mitchell-eu-google

[BMC] "Hack the Vote: The Perils of the Online Ballot Box"
online.wsj.com/articles/pamela-smith-and-bruce-mcconnell-hack-the-vote-the-perils-of-the-online-ballot-box-1401317230

[PEW] www.pewinternet.org/2014/07/03/net-threat

[RL] Rob Lemos, personal communication

[CWC] "Banks Dreading Computer Hacks Call for Cyber War Council"
www.bloomberg.com/news/print/2014-07-08/banks-dreading-computer-hacks-call-for-cyber-war-council.html

[NIST] High Integrity Software System Assurance, section 4.2,
hissa.nist.gov/chissa/SEI_Framework/framework_16.html, but you'll
have to look in the Internet Archive for it

[PB] _Against the Gods_ and this 13:22 video at
www.mckinsey.com/insights/risk_management/peter_l_bernstein_on_risk

This and other material on file at http://geer.tinho.net/pubs