The Future of Security

Nominal Delivery Draft, Rocky Mountain Information Security Conference, 18 May 2012

DANIEL E. GEER, JR.

Dan Geer is the CISO at In-Q-Tel and likes to list the following milestones: the X Window System and Kerberos (1988), the first information security consulting firm on Wall Street (1992), convenor of the first academic conference on electronic commerce (1995), the "Risk Management Is Where the Money Is" speech that changed the focus of security (1998), the Presidency of the USENIX Association (2000), the first call for the eclipse of authentication by accountability (2002), principal author of and spokesman for "Cyberinsecurity: The Cost of Monopoly" (2003), co-founder of SecurityMetrics.Org (2004), convener of MetriCon (2006–present), author of "Economics & Strategies of Data Security" (2008), and author of "Cybersecurity & National Policy" (2010). Creator of the Index of Cyber Security (2011) and the Cyber Security Decision Market (2011). Six times entrepreneur. Five times before Congress.
dan@geer.org

Good afternoon, all. Thank you for the invitation to be here today. I remind you that, as always, I speak for myself.

Were this a formal debate, the title would be the assertion "Resolved: The Internet is no place for critical infrastructure."

I say that, in part, to get your attention and in part to open a line of thought about what is critical and the degree to which that which is critical is defined as a matter of principle and the degree to which that which is critical is defined operationally; that is to say, I am distinguishing what we say and what we do.

Mainstream media and bloggers alike love to turn a spotlight on anything they can plausibly label "hypocrisy," the dictionary meaning of which is

> *the act or practice of pretending to be what one is not or to have principles or beliefs that one does not have, especially the false assumption of an appearance of virtue* [1]

The debate topic I am proposing can therefore be restated as calling "hypocrisy!" on the claim that the Internet is a critical infrastructure either directly or by transitive closure with the applications that run on or over it. If the claim were true, the divergence between our beliefs and our practices would be necessarily narrower.

It is possible that in writing this talk I am merely echoing how a free-range cattleman felt about the coming of barbed wire, roads, and land title to the American West. The great cattle drives of the West lasted twenty years before other kinds of progress made them impossible. Commercial Internet traffic began twenty years ago last summer [2].

Douglas Adams, in the posthumous book *The Salmon of Doubt*, described our reactions to technologies:

1. Anything that is in the world when you're born is normal and ordinary and is just a natural part of the way the world works.
2. Anything that's invented between when you're 15 and 35 is new and exciting and revolutionary and you can probably get a career out of it.
3. Anything invented after you're 35 is against the natural order of things.

I admit all that and more, but recalling Winston Churchill's "the further back I look, the further forward I can see," it seems to me that either the wide-open range that is the freedom of an Internet built on the end-to-end principle must die, or else

we must choose to not allow the critical infrastructure of our lives to depend on that Internet. Freedom and reliability are now at odds.

We have all the evidence we need to show that confidentiality, integrity, and availability for data or systems do not occur by magic, that each must be designed in the first place and each must be renewed often—designed as bolting them on after the fact has been shown to generally fail, renewed often as designing anything to be entirely future-proof is so far from easy as to be unlikely.

A so-called Hobson's Choice is where one is given a "take it or leave it" proposition, which is said to be not much of a choice at all. Although I suspect this applies to no one in today's audience, consider the Internet as a Hobson's Choice. You either get it, warts and all, or you get nothing.

On April 13, the Pew Foundation published a report [3] that talked about so-called "digital differences" in the U.S. As they point out,

> *One in five American adults does not use the Internet. . . . Among adults who do not use the Internet, almost half [said] that the main reason they don't go online is because they don't think the Internet is relevant to them. . . . Though overall Internet adoption rates have leveled off, adults who are already online are doing more.*

For Pew, this is another examination of the so-called "digital divide," but I ask you to consider it in a different light. For those 10% who, presented with a take it or leave it proposition regarding the Internet, choose "leave it," the Internet does not register as a desirable and may, for some of them, be undesirable.

I grew up without television and have never myself bought or owned one. I suspect there are a few in this room for whom the Hobson's Choice with respect to television has also been, or has become, "leave it." So far as I know, there is no social opprobrium, no implication that you are a loser, if you opt out of television. It is merely a choice. Such a choice entirely frustrates a whole swath of advertisers, no doubt, and since the majority of the money given to politicians this election season will doubtless be spent on television buys, one might even say that the refusal to participate in television delivers a mildly antisocial side-effect, especially if those television ads are what actually do elect the next President. If your choice to leave television out of your life is so that you can be consistent with an organized set of moral beliefs of which avoiding television is only one, then there is a ready army of sophisticated observers who will immediately suggest that you have been in some way brainwashed. Nevertheless, other than the fraction of the cost of anything you buy that is attributable to the carried-forward advertising budget of its manufacturer, you can be rather independent of television and live a good life.

That 10% of the population that doesn't see any reason to bother with the Internet is surely similar to whatever fraction of the population doesn't see any reason to bother with television. As with those opting out of television, whatever the Internet rejectionists buy will include the cost of Internet advertising bought by the manufacturer, but surely that is all. Surely they can refuse the Internet and have that be just as it sounds, something that they choose not to do anything with, and therefore inconsequential to their life, the way television is inconsequential to mine.

Not so fast. We are at the point where it may no longer be possible to live your life without having a critical dependence on the Internet, even if you live at the end of a dirt road but still occasionally buy nails or gasoline. Unlike television, where, at most, it is choosing the President or deciding what colors will dominate the spring collection, you cannot entirely unplug from the Internet even if you want to. If you are dependent on those who are dependent on television, then so what? If, however, you are dependent on those who are dependent on the Internet, then so are you. Dependence with respect to television is not transitive. Dependence with respect to the Internet is.

The source of risk is dependence, and especially dependence on expectations of system state. My definition of security itself has co-evolved with my understanding of risk and risk's source, to where I today define security as the absence of unmitigatable surprise. It is thus obvious that increasing dependence means ever more difficulty in crafting mitigations, and that increasing complexity embeds dependencies in ways such that while surprises may grow less frequent, they will be all the more unexpected when they do come, and come they will.

Because dependence on the Internet is transitive, those who choose "leave it" with respect to the Internet only get to say that in the first person; they are still dependent on it unless they are living a pre-industrial life. That rejectionists depend on people who are not rejectionist is simply a fact, a fact in the same way that the sun rises in the East is a fact. Everyone has a stake in the game.

At the same time, the rejectionists do have some species of impact on the Internet-happy, something more substantive than not buying geegaws from Internet marketeers. To the extent that we are willing to admit it, the rejectionists are now a kind of fail-safe. If we begin to penalize the rejectionists, that is to say, force them to give up their rejectionism, we will give up a residuum of societal resiliency.

What do I mean? Let me illustrate this at the personal level. I have a 401(k) retirement account with Fidelity Investments, a Fortune 500 firm with which you are all familiar. In the past few months, I have learned that Fidelity no longer accepts client instructions in writing. They only accept instructions over the Internet or, as a fallback for the rejectionist, over the phone. They simply do not accept the canonical wet ink signature on bond paper. I have sent them paper letters. They have responded in email that says what I just said, although I should note that I never gave them my email address and wouldn't have if asked. The main response on Fidelity's side is that their auditors approve of their scheme. The main response on my side is "So what?" which, of course, is my way of saying that Fidelity's auditors work for them, not me. It will doubtless not surprise you that the email letters do not contain a digital signature and, in any case, what is the equivalent of that for a phone call? Mind you, Fidelity still sends paper statements and to the same mailing address from which I have been writing.

A second personal example: I choose not to do Internet banking. I use a small, local bank—one that is far from being too big to fail. When they announced the availability of online banking, I sent them a letter stating that as I would not be using that service, that I would appreciate it if they would turn off access to my unused account or, at the least, to raise an alarm if anyone ever tried to use the account waiting in my name. To their eternal credit, they agreed without any argument. That is not the norm—try, as I have done, to tell that to the part of ADP that runs the get-your-W2-online service called iPay. (They refuse, and do so with

the kind of meaningless prose that can only be written by a psychologist retreaded as a lawyer.)

One might well conclude that a company unwilling to turn off your potential access because you ask them to do so is a company that does not, in fact, care about your security. If you will not use the account set up in your name, then you are sure to not notice that someone else has begun to do so, at least while your money or your data are still intact. An ounce of prevention is worth a pound of cure. I don't care if ADP is sure to be outstandingly prompt in sending me a data breach notice and/or buying me three years' worth of credit watch were someone to use the ADP account prepared for me; I care that it is made inoperable. I care that I not have a dependence on ADP's Internet security, however good it may be.

If there are any Estonians in the audience, you are by now sure that I am quite mad. For those of you who are not, Estonia is perhaps the most Internet-dependent country, a fact that has certainly worked well for them on balance. Quoting from an article five Sundays ago in *The Guardian* [4]:

> *42 Estonian services are now managed mainly through the [I]nternet. Last year, 94% of tax returns were made online, usually within five minutes. You can vote on your laptop (at the last election, [the President of Estonia] did it from Macedonia) and sign legal documents on a smartphone. Cabinet meetings have been paperless since 2000. Doctors only issue prescriptions electronically, while in the main cities you can pay by text for bus tickets, parking, and—in some cases—a pint of beer. Not bad for [a] country where, two decades ago, half the population had no phone line. Central to the Estonian project is the ID card, introduced in 2002. Nine in 10 Estonians have one, and—by slotting it into their computer—citizens can use their card to vote online, transfer money and access all the information the state has on them. "There's nothing on the ID card itself, because that could be dangerous if you lost it," says Katrin Pargmae who is in charge of public awareness at RIA, the country's [I]nternet authority. "It only gives you access to the database if you type in the right code." You can also present the card at the pharmacy to pick up a prescription. On public transport, it doubles as a ticket. Many people also have special ID chips on their mobile sim cards that allow them to pay people by text.*

That is entirely impressive and, as the article suggests, a degree of Estonian pride is entirely in order. That degree of dependence happens not to be for me—I want to retain the ability to opt out of most direct dependence on the Internet, viz., to opt out of that dependence which is the root of risk. I mean that as stronger than a preference, but weaker than an ultimatum.

In a free society, that which is not forbidden is permitted. In a non-free society, that which is not permitted is forbidden. The US Supreme Court is presently reviewing whether the Congress can forbid the citizen to not have health insurance, that is to say, whether the government's monopoly on the use of force can be deployed to forcibly collectivize the downside risk of illness. That is not an option I favor, but just as forcibly collectivizing the downside risk of illness has its utopian proponents, so, too, does forcibly collectivizing the downside risk of Internet exposure.

Just as Estonia is well ahead of nearly everybody in productive dependence on the Internet, so too is China well ahead of nearly everybody in forcibly collectivizing

the extent and manner in which the Internet is available to Chinese users. As sovereigns, the former is Estonia's right just as the latter is China's right. I want neither, even though I must acknowledge that as nations decide on their particular mix of dependencies, the Internet will be dramatically balkanized. The Internet will never again be as free as it is this evening.

I spent a decade and a half working in Harvard's teaching hospitals, especially the Beth Israel Hospital. On November 13, 2002, a total computer outage at the Beth Israel began [5]. The initiator was inadvertent high-volume data-sharing amongst researchers; the impact was reverting to paper for four days. The event was severe and unexpected, and recovery was frustrated by complexity. During those four days, doctors and laboratory personnel over 50 years old could effortlessly cope; most of the rest could not. Put differently, that a fallback to manual systems was possible saved the day, and it was those who could comfortably work without network dependence who delivered on that possibility, because they had done so at earlier times.

Let me now state the central thesis of this essay, and it is this: accommodating rejectionists preserves alternate, less complex, more durable means and therefore bounds dependence. Bounding dependence is the core of rational risk management.

Everyone here who has worked in systems administration knows that redundancy enables uptime guarantees. Everyone who has been at the sysadmin game for any significant time also knows that if you don't detect when that redundancy is busy saving your bacon, then you will soon be in bigger trouble. If I only need 4 out of 5 systems to be running, then the failure of any one system will cause no effect. If, however, I don't notice that I've had that failover event, then any subsequent failure is non-recoverable and a surprise. One of the principal arguments for hot standbys is that when the failover has to happen, the equipment to which the failover is directed is known to be working.

Ten years ago, Bill LeFebvre gave a USENIX talk on the operational changes driven by the impact of 9/11 on the Web presence of the Cable News Network, better known as CNN. In it, he described how when demand spikes they shed load, but in the case of CNN, shedding load meant taking, say, the Cartoon Network's servers and re-purposing them on the fly. These days, there are probably lots of VMs and clouds involved, but the idea is the same, that having hot standbys beats having spare, unused capacity any day, since amortizing the cost of the hot standbys through, say, running the Cartoon Network on them is financially sound and, which is more, it guarantees that you know the hot standbys work when the fail-over is necessary, such as when there is an order of magnitude spike in demand for news. Anyone who has ever found that their emergency generator didn't start when it needed to knows what I am talking about. So has anyone who has ever gone to one's backup media only to discover that they are blank.

Summing up so far, 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.

In the language of statistics, common-mode failure comes from under-appreciated mutual dependence. Quoting from NIST's section on redundancy in their High Integrity Software System Assurance documentation [6]:

> *[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.

Which brings us to critical infrastructure and the interconnection between critical infrastructures by way of the Internet. For the purpose of this essay, I will use the definition found in Presidential Decision Directive 63, issued by then-President Clinton [7]: "Critical infrastructures are those physical and cyber-based systems essential to the minimum operations of the economy and government."

Note the wording: "essential to the minimum operations"—that does not read as a requirement that the armor deflect all bullets, only that no bullet be paralyzing. One of the great Allied victories of World War II was getting 338,000 soldiers off the beaches of Dunkirk using 800 "little boats," a paragon of the phrase "essential to minimum operations," as none of those little boats were warships in any formal way.

Defined in its technical sense, the Internet is a network of networks, not a single entity. That the majority of its main protocols were designed precisely for tolerance to random faults and for the absence of common-mode failure has been proven out in practice, perhaps nowhere as spectacularly as when Bill Cheswick and Steven Branigan mapped Yugoslavian networks during the NATO bombardment in the spring of 1999 [8]. Those main Internet protocols worked so well that innovation blossomed simply because the Internet did not depend on the flawless functioning of every one of its moving parts. It was not designed, however, for resistance to targeted faults, which, as Laszlo Barabasi showed, cannot be done at the same time as you are designing for resistance to random faults [9].

In an Internet crowded with important daily-life functions, the chance of common-mode failure is no idle worry. The Obama administration is broadly committed to increasing dependence on the Internet, most notably on two fronts, either of which might be said to be themselves "essential to the minimum operations of the economy and government." First, is the press for electronic health records. Second, is the press for the so-called Smart Grid. As with most garden paths, there is nothing wrong interior to the arguments for electronic health records or for

Smart Grids. Both have eminently useful results for which a desire is rational. Both illustrate my point.

With respect to electronic health records: their utility depends on the smooth functioning of electric power, networks, computers, displays, and a range of security features that we can discuss another day, particularly as it relates to maintaining consistency across multiple practices [10]. With respect to the Smart Grid: its utility depends on almost everything we now know about power, including the absolute necessity of good clocks, a wide range of industrial controls operated flawlessly at a distance and guaranteed not to lie about their state, and another range of security features we can also discuss some other day.

Because both of these involve new levels of exposure to common-mode risk, some of which are risks that electronic health records and Smart Grids share, both will add new failure modes to the world we live in. On good days, both will deliver far better, more cost-effective benefits than that which we now have. On bad days, the reverse will be true and, as the Beth Israel example proved by demonstration, doing without those benefits will be easier for those who can remember not having had them.

Put differently, each new dependence raises the magnitude of downside risk, the potential for collateral damage, and the exposure of inter-relationships never before contemplated. Forget the banks—it is the Internet that is too big to fail. While there is no entity that can bail out the Internet, there is no meaningful country that is not developing ways to disrupt the Internet use of its potential adversaries. The most a country can hope to do is to preserve the Internet interior to itself, as Estonia demonstrated when under attack from Russia. Of course, at some level of transborder interconnection, the very concept of "interior" loses meaning, as every one of you here knows if you have ever had to explain the limits of perimeter defense to a new client or counterparty.

Now let me hasten to add that where ignoring a risk is negligence, a sin of omission, purposely inflating a risk is fraud, a sin of commission. As Dinei Florencio and Cormac Herley showed in their 2011 paper [11] at the Workshop on the Economics of Information Security, the estimates we have of the impact of cybercrime are all but surely universally inflated, the sin of commentators who may have been able to merge omission and commission. This is an important point, enough so that a shortened form of the WEIS paper appeared as an invited OpEd April 14th in the *New York Times* [12].

I've come to the conclusion that part of what makes a good security person is some sort of intrinsic fascination with failure. I am certain that designing for tolerable failure modes is precisely what security engineering is fundamentally about. If I am right, then the failure mode you did not think of will not be in your design and, therefore, whether it is tolerable will depend on many things, perhaps even the phase of the moon. The question, then, is whether tolerable failure modes can be themselves designed, that is to say, whether a failure mode never before possible can be added to the system such that larger, intolerable failures can be precluded? Is there a critical cyber infrastructure analog to a shear-bolt in the drivetrain of heavy machinery?

No country, no government, no people needs rules against things that are impossible. Obviously, the onrush of things never before possible creates vacuua where, in the fullness of time, there will have to be rules. As the current Mayor of

Chicago (while still in the White House) put it in his characteristically blunt way, the creation of rules is easier in a time of crisis and, therefore, one must "never let a good crisis go to waste." He is right as a matter of observation; he is wrong as a matter of probity. Just as driving under the influence of alcohol is wrong, so is making policy under the influence of adrenaline. The law-making of the last decade is illustrative, but the point is hardly new; eleven years before he became the fourth President of the United States, James Madison said, "Perhaps it is a universal truth that the loss of liberty at home is to be charged to provisions against danger, real or pretended, from abroad."

One wonders how Madison would feel about an interconnected world where "abroad" has so thoroughly lost its meaning, at least with respect to Internet-dependent critical infrastructure if not national frontiers. My guess is that Madison would decide that the Internet is, per se, "abroad."

Two months ago, more or less, I gave a speech asking whether having people in the loop for security is a fail-safe or a liability [13]. I won't recount the arguments here, but I will give my conclusion: that a good security design takes people out of the loop except when it cannot and, when it cannot, it is clear that this is so. I gave two examples:

My previous employer was Verdasys, where I was Chief Scientist for the product design team. For those with a deep background, our product was a distributed, recording version of the Orange Book's "Reference Monitor" implemented as a rootkit. That said, we could do nearly anything to detect and control data handling of any sort at any granularity.

We installed at a major hospital. There, the Chief of Medicine demanded that under no circumstances could our product block access to patient data, since who knows what sort of mortal emergency might be in progress. At the same time, the General Counsel demanded that under no circumstances could our product permit a breach of regulatory controls on data handling. The solution to this standoff was that whenever someone asked for data that was nominally forbidden, a popup window would appear which said "Against policy. Click here to proceed." With that, no data was denied but at the same time no person could deny having intent. This finesse represented the well-placed insertion of a tiny bit of sentience in an otherwise fully automated protection regime.

The other example:

I have good relations with a number of the largest banks. One of them has long since made user-level provisioning a completely automated process. This automated provisioning control includes deprovisioning—what you might describe as removing Dan's access within 120 seconds of the time Dan submits his letter of resignation or, for that matter, slugs a Managing Director on the trading floor. Fast, hands off, one-button deprovisioning makes regulators happy. It makes General Counsels happy. But it's a nightmare if it goes into a loop. The bank I'm thinking of has coded for this explicitly; if 50 resignations have come in within an hour, the deprovisioning system halts and will not proceed until a human gives it authority to do so. Putting a human back into the loop has saved their bacon at least once.

These are examples of where putting a human in the loop, that is to say, falling back from automation, has proven to be a breakthrough finesse. Both were designed that way, neither was an accident, and both required real labor in getting everyone on the same page.

I ask, is such an outcome something that can only be done on a case by case basis, something that cannot become part of a security discipline in the large, something that avoids both sins of omission and sins of commission? One hopes that it can be; there simply is not time to make every security-related architectural decision go down the path that these two examples propose.

The public at large is not and cannot be expert in the way this audience is expert, nor should they have to be. That the public has, shall we say, "volunteered" its unused computing power to botmasters is nothing so much as an historical mirror of how press gangs once filled the rosters of the British Navy. But how is that concretely different from a formal mandate that if they have medical records they shall be electronic, or if they receive electricity that the metering regime be a surveillance tool? How is it different from finding that compliance auditors have certified to the distant regulators that there is no need to accept a signed paper letter detailing the wishes of the financial client, wishes that those self-same regulators demand the financial client formally submit if the financial institution is to be protected from tort claims?

I've come to the conclusion, as recently have others, that security is a proper subset of reliability. The logic is that security is a necessary but insufficient condition for reliability. As such, connecting the insecure (and thus unreliable) to the important and expecting the mélange to be reliable is utter foolishness. As Marcus Ranum says, "A system that can be caused to do undesigned things by outsiders is not 'reliable' in any sense of the word." On that point, I refer you to the work being done by Sergey Bratus, Meredith Patterson, and others whose startling insight deserves full quotation [14]:

> *The Language-theoretic approach (LANGSEC) regards the Internet insecurity epidemic as a consequence of ad hoc programming of input handling at all layers of network stacks, and in other kinds of software stacks. LANGSEC posits that the only path to trustworthy software that takes untrusted inputs is treating all valid or expected inputs as a formal language, and the respective input-handling routines as a recognizer for that language. The recognition must be feasible, and the recognizer must match the language in required computation power.*
>
> *When input handling is done in an ad hoc way, the de facto recognizer, i.e., the input recognition and validation code, ends up scattered throughout the program, does not match the programmers' assumptions about safety and validity of data, and thus provides ample opportunities for exploitation. Moreover, for complex input languages the problem of full recognition of valid or expected inputs may be UNDECIDABLE, in which case no amount of input-checking code or testing will suffice to secure the program. Many popular protocols and formats fell into this trap, the empirical fact with which security practitioners are all too familiar.*
>
> *Viewed from the venerable perspective of Least Privilege, . . . computational power is privilege, and should be given as sparingly as any other kind of privilege to reduce the attack surface. We call this . . . the Minimal Computational Power Principle.*
>
> *We note that recent developments in common protocols run contrary to these principles. In our opinion, this heralds a bumpy road ahead. In particular, HTML5 is Turing-complete, whereas HTML4 was not.*

So far as I can guess, nearly nothing we have in our cyber interfaces to critical infrastructure meets LANGSEC's test. For that reason, if no other, attaching the cyber interface of critical infrastructure to the Internet is a flat-out guarantee of error. As always, such error may be improbable, but probabilistic events do eventually occur. If we are unlucky, those errors will not be prompt.

In a conference panel, I was once asked what malware I would write if it was not a question of labor or difficulty. My answer remains the same: I'd find a way to make the occasional odd modification to your Excel formulae and I would embed this malware in the spreadsheet itself so that any sharing of the spreadsheet would propagate my malware. As Excel formulae are probably the world's most prevalent programming language, in a period of time I would de-synchronize all copies of what are ostensibly the same document. This wouldn't end the world, but think about, say, how derivative pricing is done. If I could do this, for which I have neither skill nor desire, then I might even be properly called a terrorist insofar as terrorism is a means of coercion by way of spreading fear. I would not have to destroy your data—you would.

At this point, I am at serious risk of being exactly the kind of fear mongerer that quickly becomes a fraud. That is, of course, not my point. My point is that the working definition of critical infrastructure is broad and, which is more, it is indistinct.

There has been much talk about whether to grant the President a so-called kill-switch for the Internet. There is a considerable logic to that if you accept what I have been saying, namely that in the presence of interdependence that is inestimable there may be times when it is not possible to disambiguate friend from foe. Were someone on an inbound airplane found to have smallpox, the passengers and crew would be quarantined as a matter of public health until such time as each of them could be separately certified as disease free. Many important enterprises, public and private, quarantine inbound email with nearly as much vigor as they quarantine inbound DHL packages. The logic is sound. The time scale is human.

In a kind of living history, we have residing amongst ourselves cloistered communities such as the Amish. We accommodate them. I expect that if a food crisis of some sort were to materialize, it is the Amish who would be least affected. We have amongst ourselves so-called Neo-Luddites. In some sense, the Luddites of old had a more principled analysis—they knew where the machines would lead and on the basis of their analysis they acted. The Amish merely wish to be left alone, such as to remove their children from compulsory education at the close of the eighth grade. So far as I know, their case, Wisconsin v. Yoder, is the only such case to ever reach the US Supreme Court, which found in their favor. I ask, is there room in our increasingly wired world for those who choose merely to be left alone, in this case to choose to not participate in the Internet society? Do those who do not participate deserve to not have their transactions of all sorts be exposed to a critical infrastructure dependent on the reliability of Internet applications as a class?

Paraphrasing Melissa Hathaway from her 60-day review of US cyber policy [15] for President Obama, the United States' ability to project power depends on information technology, and, as such, cyber insecurity is the paramount national security risk. Putting aside an Internet kill-switch, might it be wise for the national authorities to forbid, say, Internet Service Providers from propagating Telnet or SSH v1 or other protocols known to be insecurable? If not that, should

cyber components of the critical infrastructure be forbidden to accept such connections? There is certainly a freedom versus reliability debate topic in that—if not a natural policy. As with most things, there is a direct historical echo here as well: in 1932, the foremost political commentator of the age, Walter Lippmann, told President Roosevelt, "The situation is critical, Franklin. You may have no alternative but to assume dictatorial powers."

Again, when 10% of the population sees nothing in the Internet for them, should we respect and ensure that, as with the Amish, there is a way for them to opt out without choosing to live in a cave? Should we preserve manual means?

I say "YES," and I say so because the preservation of manual means is a guarantee of a fallback that does not have a common-mode failure with the rest of the interconnected, mutually vulnerable Internet world. That this is not an easy choice is the understatement of the day, if not year. I do not (yet) claim to have a fully working model here, but neither do our physicist friends (yet) have a unified field theory.

My colleague Mukul Pareek and I run the "Index of Cyber Security" [16]. It is a survey-based index of sentiment, modeled on the Consumer Confidence and Purchasing Managers Indices. It has been in operation for a year. The respondents are all CISOs, individuals whose view of cyber security is based on direct operational responsibility for their firm's piece of the networked world. It is a risk index, which is to say that when perceived risk rises, so does the Index. Over the course of the year, the Index has risen inexorably, reflecting the view of experts at least as good as those in this audience that, in the aggregate, risk is accumulating and in much the same way that burnable timber accumulates here on the eastern slope of the Rockies.

Because the Index is composed not of one question but of twenty, each asking about one or another source of risk such as malware, hacktivism, counterparty interconnections, and so forth, we have found that the steady rise in the Index is not dominated by any one sub-component nor is the ordering of the influence of the sub-components on the overall Index stable and unchanging. One month, it is counterparties. Another month, it is the impact of diverting security budgets to compliance. In yet others, it is malware that is undetectable by any of the array of commercial products for malware detection. And so forth. These are your peers speaking and they are saying that the risk is growing. They also make comments many of which are, in so many words, talking about irresistible commercial pressures.

Against such a formal, metrics-based backdrop, I can confidently say that "we" are not running fast enough to stay in the same place. If those respondents and I are not fooling ourselves, preserving fallback is prudent, if not essential. That does not mean, per se, that the preservation of manual means is an easy out. It may well be that, as various Department of Defense thinkers have come around to saying, our goal can no longer be intrusion prevention but must now turn towards intrusion tolerance. As before, this is easier said than done, but if we are to practice evidence-based medicine on the body Internet, we must acknowledge that expensive therapy is not always the answer. Cost-effective medicine cannot be practiced if every human life is infinitely valuable. Perhaps one of you can come up with a cyber analog to "quality-adjusted life years" and help us all decide when to treat, when to palliate, and when to just plain avoid.

As you well know, 100% availability can be achieved by either driving the mean time between failures to infinity such that nothing ever breaks or driving the mean time to repair to zero such that failure consequences are de minimus. I am old enough to remember that rebooting a machine for prophylactic purposes was what systems administrators did. These days, most people view a reboot as proof of failure. I don't agree.

Summing up for the second time, 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. If the critical infrastructures are those physical and cyber-based systems essential to the minimum operations of the economy and government, and if aggregate risk, as described by leading cyber security operational management, is growing steadily, then do we put more of our collective power behind forcing security improvements that will be sharply diseconomic, or do we preserve fallbacks of various sorts in anticipation of events that, if the Index of Cyber Security can be believed, seem more likely to happen as time passes? Does the old Yankee saying "Use it up, wear it out, make it do, or do without" have any guidance for us? Is centralizing authority the answer, or is avoiding further dependence until we can fix things the better strategy? Can we imagine starting over in any real sense, or is balkanization not just for nations but for critical sectors as well? Is the creative destruction that is free enterprise now to be focused on remaking what are normally the steadying flywheels of American society, by which I mean government and our most capital-intensive industries? Does the individual who still prefers to fix things he or she already has to be celebrated, or are those individuals to be herded into National Health Information Networks, Smart Grids, and cars that drive themselves?

In closing, remember that the Internet was built by academics, researchers, and hackers—meaning that it embodies the liberal cum libertarian cultural interpretation of "American values," namely, that it is open, non-hierarchical, self-organizing, and leaves essentially no opportunities for governance beyond protocol definition. Anywhere the Internet appears, it brings those values with it (treating censorship as a routing failure, for example). Other cultures, other governments, know that these are our strengths and that we are dependent upon them; hence as they adopt the Internet they become dependent on those strengths and thus on our values. A greater challenge to their sovereignty does not exist. The challenge to our sovereignty is dual—it is the choice of whether to commit our critical infrastructures to the Internet in the entire or to discard our fallbacks along with those who practice them, to bet the farm on a roll of the geopolitical dice.

There is never enough time. Thank you for yours.

References

[1] Merriam-Webster Unabridged.

[2] Interconnection of PSInet and UUNet by CIX, summer 1991.

[3] "Digital Differences," Pew Research Center, April 13, 2012: tinyurl.com/d7eqo7v.

[4] "How Tiny Estonia Stepped Out of USSR's Shadow to Become an Internet Titan," *The Guardian*, April 15, 2012: tinyurl.com/7srar5z.

[5] P. Kilbridge, "Computer Crash—Lessons from a System Failure," *New England Journal of Medicine* vol. 348, no. 10, pp. 881–82, March 6, 2003: tinyurl.com/75fjmbb.

[6] NIST, High Integrity Software System Assurance, section 4.2: tinyurl.com/canwggd.

[7] Presidential Decision Directive 63, May 22, 1998: tinyurl.com/4974j.

[8] S. Branigan and B. Cheswick, "The Effects of War on the Yugoslavian Network," 1999: tinyurl.com/cu9nd5u.

[9] L. Barabasi and R. Albert, "Emergence of Scaling in Random Networks," *Science*, vol. 286, October 15, 1999, pp. 509–12.

[10] Simon S.Y. Shim, "The CAP Theorem's Growing Impact," "The CAP Theorem's Growing Impact," *IEEE Computer* 45 (2)21–22, February 2012.

[11] D. Florencio and C. Herley, "Sex, Lies and Cyber-Crime Surveys," WEIS 2011: tinyurl.com/3zsspah.

[12] D. Florencio and C. Herley, "The Cybercrime Wave That Wasn't," *New York Times*, April 14, 2012: tinyurl.com/8ylrf7b.

[13] D. Geer, "People in the Loop: Are They a Failsafe or a Liability?" *Suits & Spooks*, February 8, 2012: tinyurl.com/7cavobr.

[14] LANGSEC: Language-Theoretic Security: http://www.cs.dartmouth.edu/~sergey/langsec/.

[15] The White House Blog, "Securing Our Digital Future," May 29, 2009: http://tinyurl.com/l6z3fl.

[16] The Index of Cyber Security: http://cybersecurityindex.org.