RSA CONFERENCE 2014

FEBRUARY 24 – 28 | MOSCONE CENTER | SAN FRANCISCO

Share.
Learn.
Secure.

Capitalizing on
Collective Intelligence

We Are All Intelligence Officers Now

SESSION ID: EXP-F02

Dan Geer

Principal
Geer Risk Services

.We Are All Intelligence Officers Now

.Dan Geer, 28 February 14, RSA/San Francisco

Good morning. Thank you for the invitation to speak with you today, which, let me be clear, is me speaking for myself, not for anyone or anything else. As you know, I work the cyber security trade, that is to say that my occupation is cyber security. Note that I said "occupation" rather than "profession." Last September, the U.S. National Academy of Sciences concluded that cyber security should be seen as an occupation and not a profession because the rate of change is simply too great to consider professionalization.[NAS] You may well agree that that rate of change is paramount, and, if so, you may also agree that cyber security is the most intellectually demanding occupation on the planet.

The goal of the occupation called cyber security grows more demanding with time, which I need tell no one here. That growth is like a river with many tributaries. Part of the rising difficulty flows from rising complexity, part of it from accelerating speed, and part of it from the side effects of what exactly we would do if this or that digital facility were to fail entirely -- which is to say our increasing dependence on all things digital. One is at risk when something you depend upon is at risk. Risk is, in other words, transitive. If X is at risk and I depend on X, then I, too, am at risk to whatever makes X be at risk. Risk is almost like inheritance in a programming language.

I am particularly fond of the late Peter Bernstein's definition of risk: "More things can happen than will."[PB] I like that definition not because it tells me what to do, but rather because it tells me what comes with any new expansion of possibilities. Put differently, it tells me that with the new, the realm of the possible expands and, as we know, when the realm of the possible expands, prediction is somewhere between difficult and undoable. The dynamic is that we now regularly, quickly expand our dependence on new things, and that added dependence matters because the way in which we each and severally add risk to our portfolio is by way of dependence on things for which their very newness makes risk estimation, and thus risk management, neither predictable nor perhaps even estimable.

The Gordian Knot of such tradeoffs -- our tradeoffs -- is this: As society becomes more technologic, even the mundane comes to depend on distant digital perfection. Our food pipeline contains less than a week's supply, just to take one example, and that pipeline depends on digital services for everything from GPS driven tractors to robot vegetable sorting machinery to coast-to-coast logistics to RFID-tagged livestock. Is all the technologic dependency, and the data that fuels it, making us more resilient or more fragile?

In the cybersecurity occupation, in which most of us here work, we certainly seem to be getting better and better. We have better tools, we have better understood practices, and we have more and better colleagues. That's the plus side. But from the point of view of prediction, what matters is the ratio of skill to challenge; as far as I can estimate, we are expanding the society-wide attack surface faster than we are expanding our collection of tools, practices, and colleagues. If your society is growing more food, that's great. If your population is growing faster than your improvements in food production can keep up, that's bad. So it is with cyber risk management: Whether in detection, control, or prevention, we are notching personal bests, but all the while the opposition is setting world records. As with most decision making under uncertainty, statistics have a role, particularly ratio statistics that magnify trends so that the latency of feedback from policy changes is more quickly clear. Yet statistics, of course, require data, to which I will return in a moment.

In medicine, 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, anyway. But if you check in with Bubonic Plague or Typhus or Anthrax, you will have zero privacy as those are "reportable conditions," as variously mandated 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 the government? 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[CF] 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 cyber failures in the form of their data breach laws, 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. If you discover a data breach, do you have an ethical obligation to report it? Should the law mandate that you fulfill such an obligation?

Almost everyone here has some form of ingress filtering in place by whatever name -- firewall, intrusion detection, whitelisting, and so forth and so on. Some of you have egress filtering because being in a botnet, that is to say being an accessory to crime, is bad for business. Suppose you discover that you are in a botnet; do you have an obligation to report it? Do you have an obligation to report the traffic that led you to conclude that you had a problem? Do you even have an obligation to bother to look and, if you don't have or want an obligation to bother to look, do you want your government to require the ISPs to do your looking for you, to notify you when your outbound traffic marks you as an accomplice to crime, whether witting or unwitting? Do you want to lay on the ISPs the duty to guarantee a safe Internet? They own the pipes and if you want clean pipes, then they are the ones to do it. Does deep packet inspection of your traffic by your ISP as a public health measure have your support? Would you want an ISP to deny access to a host, which might be your host, that is doing something bad on their networks? Who gets to define what is "bad?"

If you are saying to yourself, "This is beginning to sound like surveillance" or something similar, then you're paying attention. Every one of you who lives in a community that has a neighborhood watch already has these kinds of decisions to make. Let's say that you are patrolling your street, alone, and there have been break-ins lately, there have been thefts lately, there has been vandalism lately. You've lived there for ten years and been on that neighborhood watch for five. You are on duty and you see someone you've never seen crossing the street first from one side then the other, putting a hand on every garden gate. What do you do? Confront them the way a polite neighbor would? Challenge them the way a security guard would? Run home to lock your own doors and draw your drapes? Resign from the neighborhood watch because you are really not ready to do anything strenuous?

Returning to the digital sphere, we are increasing what it is that can be observed, what is observable. Instrumentation has never been cheaper. Computing to fiddle with what has been observed has never been more available. As someone who sees a lot of fresh business plans, I can tell you that these days Step Six is never "Then we build a data center." Step Six, or whatever, is universally now "Then we buy some cloud time and some advertising." This means that those to whom these outsourcing contracts go are in a position to observe, and observe a lot. Doubtless some of what they observe will be problematic, whether on legal or moral grounds. Should a vendor of X-as-a-Service be obliged to observe what their customers are doing? And if they are obliged to observe, should they be obliged to act on what they observe, be that to report, to deploy countermeasures, or both?

As what is observable expands so, naturally, does what has been observed. Dave Aitel says "There's no reason a company in this day and age can't have their own Splunk or ElasticSearch engine that allows them to search and sort a complete history of every program anyone in the company has ever executed."[DA] Sometime in the last five to ten years we passed the point on the curve where it became much cheaper to keep everything than to do selective deletion. When you read the Federal Rules of Civil Procedure with respect to so-called e-discovery, you can certainly conclude that total retention of observed data is a prudent legal strategy. What is less clear is whether you have a duty to observe given that you have the capacity to do so. All of which also applies to what others can observe about you.

This is not, however, about you personally. Even Julian Assange, in his book _Cypherpunks_, said "Individual targeting is not the threat." It is about a culture where personal data is increasingly public data, and assembled en masse. All we have to go on now is the hopeful phrase "A reasonable expectation of privacy" but what is reasonable when one inch block letters can be read from orbit? What is reasonable when all of your financial or medical life is digitized and available primarily over the Internet? Do you want ISPs to retain e-mails when you are asking your doctor a medical question (or, for that matter, do you want those e-mails to become part of your Electronic Health Record)? Who owns your medical data anyway? Until the 1970s, it was the patient but regulations then made it the provider. With an Electronic Health Record, it is likely to revert to patient ownership, but if the EHR belongs to you, do you get to surveil the use that is made of it by medical providers and those that recursively they outsource to? And if not, why not?

Observability is fast extending to devices. Some of it has already appeared, such as the fact that any newish car is broadcasting four unique Bluetooth radio IDs, one for each tire's valve stem. Some of it is in a daily progression, such as training our youngsters to accept surveillance by stuffing a locator beacon in their backpack as soon as they go off to Kindergarten. Some of it is newly technologic, like through the wall imaging, and some of it is simply that we are now surrounded by cameras that we can't even see where no one camera is important but they are important in the aggregate when their data is fused. Anything, and I mean anything, that has "wireless" in its name creates the certainty of traffic analysis.

As an example relevant to rooms such as this, you should assume that all public facilities will soon convert their lighting fixtures to LEDs, LEDs that are not just lights but also have an embedded, chip-based operating system, a camera, sensors for CO/CO2/pollutant emissions, seismic activity, humidity & UV radiation, a microphone, wifi and/or cellular interfaces, an extensible API, an IPv4 or v6 address per LED, a capacity for disconnected "decision making on the pole," cloud-based remote management, and, of course, bragging rights for how green you are which you can then monetize in the form of tax credits.[S] I ask again, do you or we or they have a duty to observe now that we have an ability to do so? It is, as you know, a long established norm for authorities to seize the video stored in surveillance cameras whether the issue at hand is a smash and grab or the collapse of an Interstate highway bridge.[M] What does that mean when data retention is permanent and recording devices are omnipresent? Does that make you the observed or the observer? Do we have an answer to "Who watches the watchmen?"[J]

By now it is obvious that we humans can design systems more complex than we can then operate. The financial sector's "flash crashes" are the most recent proof-by-demonstration of that claim; it would hardly surprise anyone were the fifty interlocked insurance exchanges for Obamacare to soon be another. Above some threshold of system complexity, it is no longer possible to test, it is only possible to react to emergent behavior. Even the lowliest Internet user is involved -- one web page can easily touch scores of different domains. While writing this, the top level page from cnn.com had 400 out-references to 85 unique domains each of which is likely to be similarly constructed and all of which move data one way or another. If you leave those pages up, then because many such pages have an auto-refresh, moving to a new subnet signals to every one of the advertising networks that you have done so. How is this different than having a surveillance camera in the entry vestibule of your home?

We know, and have known for some time, that traffic analysis is more powerful than content analysis. If I know everything about to whom you communicate including when, where, with what inter-message latency, in what order, at what length, and by what protocol, then I know you. If all I have is the undated, unaddressed text of your messages, then I am an archaeologist, not a case officer. The soothing mendacity of proxies for the President saying "It's only metadata" relies on the ignorance of the listener. Surely no one here is convinced by "It's only metadata" but let me be clear: you are providing that metadata and, in the evolving definition of the word "public," there is no fault in its being observed and retained indefinitely. Harvard Law professor Jonathan Zittrain famously noted that if you preferentially use online services that are free, "You are not the customer, you're the product." Why? Because what is observable is observed, what is observed is sold, and users are always observable, even when they are anonymous.

Let me be clear, this is not an attack on the business of intelligence. The Intelligence Community is operating under the rules it knows, most of which you, too, know, and the goal states it has been tasked to achieve. The center of gravity for policy is that of goal states, not methods.

Throughout the 1990s, the commercial sector essentially caught up with the intelligence sector in the application of cryptography -- not the creation of cyphers, but their use. (Intelligence needs new cyphers on a regular basis whereas commercial entities would rather not have to roll their cypher suites at all, much less regularly.) In like manner commercial firms are today fast catching up with the intelligence sector in traffic analysis. The marketing world is leading the way because its form of traffic analysis is behavior-aware and full of data fusion innovation -- everything from Amazon's "people who bought this later bought that" to 1 meter accuracy on where you are in the shopping mall so that advertisements and coupons can appear on your smartphone for the very store you are looking in the window of, to combining location awareness with what your car and your bedroom thermostat had to say about you this morning. More relevant to this audience, every cutting edge data protection scheme now has some kind of behavioral component, which simply means collecting enough data on what is happening that subsequently highlighting anomalies has a false positive rate low enough to be worth following up.

If you decide to in some broad sense opt out, you will find that it is not simple. Speaking personally, I choose not to share CallerID data automatically by default. Amusingly, when members of my friends and family get calls from an unknown caller, they assume it is me because I am the only person they know who does this. A better illustration of how in a linear equation there are N-1 degrees of freedom I can't imagine. Along those same lines, I've only owned one camera in my life and it was a film camera. Ergo, I've never uploaded any photos that I took. That doesn't mean that there are no digital photos of me out there. There are 3+ billion new photos online each month, so even if you've never uploaded photos of yourself someone else has. And tagged them. In other words, you can personally opt out, but that doesn't mean that other folks around you haven't effectively countermanded your intent.

In short, we are becoming a society of informants. In short, I have nowhere to hide from you.

As I said before and will now say again, the controlling factor, the root cause, of risk is dependence, particularly dependence on the expectation of stable system state. Yet the more technologic the society becomes, the greater the dynamic range of possible failures. When you live in a cave, starvation, predators, disease, and lightning are about the full range of failures that end life as you know it and you are well familiar with each of them. When you live in a technologic society where everybody and everything is optimized in some way akin to just-in-time delivery, the dynamic range of failures is incomprehensibly larger and largely incomprehensible. The wider the dynamic range of failure, the more prevention is the watchword. Cadres of people charged with defending masses of other people must focus on prevention, and prevention is all about proving negatives. Therefore, and inescapably so, there is only one conclusion: as technologic society grows more interconnected, it becomes more interdependent within itself. As society becomes more interdependent within itself, the more it must rely on prediction based on data collected in broad ways, not in targeted ways. That is surveillance. That is intelligence practiced not by intelligence agencies but by anyone or anything with a sensor network.

Spoken of in this manner, official intelligence agencies that hoover up everything are simply obeying the Presidential Directive that "Never again" comes true. And the more complex the society they are charged with protecting becomes, the more they must surveil, the more they must analyze, the more data fusion becomes their only focus. In that, there is no operational difference between government acquisition of observable data and private sector acquisition of observable data, beyond the minor detail of consent.

David Brin was the first to suggest that if you lose control over what data can be collected on you, the only freedom-preserving alternative is that everyone else does, too.[DB1] If the government or the corporation or your neighbor can surveil you without asking, then the balance of power is preserved when you can surveil them without asking. Bruce Schneier countered that preserving the balance of power doesn't mean much if the effect of new information is non-linear, that is to say if new information is the exponent in an equation, not one more factor in a linear sum.[DB2] Solving that debate requires that you have a strong opinion on what data fusion means operationally to you, to others, to society. If, indeed, and as Schneier suggested, the power of data fusion is an equation where new data items are exponents, then the entity that can amass data that is bigger by a little will win the field by a lot. That small advantages can have big outcome effects is exactly what fuels this or any other arms race.

Contradicting what I said earlier, there may actually be a difference between the public and the private sector because the private sector will collect data only so long as increased collection can be monetized, whereas government will collect data only so long as increased collection can be stored. With storage prices falling faster than Moore's Law, government's stopping rule may thus never be triggered.

In the Wikipedia article about Brin, there is this sentence, "It will be tempting to pass laws that restrict the power of surveillance to authorities, entrusting them to protect our privacy -- or a comforting illusion" thereof.[W] I agree with one of the possible readings of that sentence, namely that it is "tempting" in the sense of being delusional. Demonstrating exactly the kind of good intentions with which the road to Hell is paved, we have codified rules that permit our lawmakers zero privacy, we give them zero ability to have a private moment or to speak to others without quotation, without attribution, without their game face on. In the evolutionary sense of the word "select," we select for people who are without expectation of authentic privacy or who jettisoned it long before they stood for office. Looking in their direction for salvation is absurd. And delusional.

I am, however, hardly arguing that "you" are powerless or that "they" have taken all control. It is categorically true that technology is today far more democratically available than it was yesterday and less than it will be tomorrow. 3D printing, the whole "maker" community, DIY biology, micro-drones, search, constant contact with whomever you choose to be in constant contact with -- these are all examples of democratizing technology. This is perhaps our last fundamental tradeoff before the Singularity occurs: Do we, as a society, want the comfort and convenience of increasingly technologic, invisible digital integration enough to pay for those benefits with the liberties that must be given up to be protected from the downsides of that integration? If risk is that more things can happen than will, then what is the ratio of things that can now happen that are good to things that can now happen that are bad? Is the good fraction growing faster than the bad fraction or the other way around? Is there a threshold of interdependence beyond which good or bad overwhelmingly dominate?

We are all data collectors, data keepers, data analysts. Some citizens do it explicitly; some citizens have it done for them by robots. To be clear, we are not just a society of informants, we are becoming an intelligence community of a second sort. Some of it is almost surely innocuous, like festooning a house with wireless sensors for home automation purposes. Some of it is cost effectiveness driven, like measuring photosynthesis in a corn field by flying an array of measurement devices over it on a drone. I could go on, and so could you, because in a very real sense I am telling you nothing you don't already know. Everyone in this and other audiences knows everything that I have to say, even if they weren't aware that they knew it.

The question is why is this so? Is this majority rule and the intelligence function is one the majority very much wants done to themselves and others? Is this a question of speed and complexity such that citizen decision making is crippled not because facts are hidden but because compound facts are too hard to understand? Is this a question of wishful thinking of that kind which can't tell the difference between a utopian fantasy, a social justice movement, and a business opportunity? Is this nowhere near such a big deal as I think it is because every day that goes by without a cascade failure only adds evidence that such possibilities are becoming ever less likely? Is the admonition to "Take care of yourself" the core of a future where the guarantee of a good outcome for all is the very fact that no one can hide? Is Nassim Taleb's idea that we are easily fooled by randomness[TF] at play here, too? If the level of observability to which you are subject is an asset to you, then what is your hedge against that asset?

This is not a Chicken Little talk; it is an attempt to preserve if not make a choice while choice is still relevant. As The Economist in its January 18 issue so clearly lays out,[TE] we are ever more a service economy, but every time an existing service disappears into the cloud, our vulnerability to its absence increases as does the probability of monopoly power. Every time we ask the government to provide goodnesses that can only be done with more data, we are asking government to collect more data.

Let me ask a yesterday question: How do you feel about traffic jam detection based on the handoff rate between cell towers of those cell phones in use in cars on the road? Let me ask a today question: How do you feel about auto insurance that is priced from a daily readout of your automobile's black box? Let me ask a tomorrow question: In what calendar year will compulsory auto insurance be more expensive for the driver who insists on driving their car themselves rather than letting a robot do it? How do you feel about public health surveillance done by requiring Google and Bing to report on searches for cold remedies and the like? How do you feel about a Smart Grid that reduces your power costs and greens the atmosphere but reports minute-by-minute what is on and what is off in your home? Have you or would you install that toilet that does a urinalysis with every use, and forwards it to your clinician?

How do you feel about using standoff biometrics as a solution to authentication? At this moment in time, facial recognition is possible at 500 meters, iris recognition is possible at 50 meters, and heart-beat recognition is possible at 5 meters. Your dog can identify you by smell; so, too, can an electronic dog's nose. Your cell phone's accelerometer is plenty sensitive enough to identify you by gait analysis. The list goes on. All of these are data dependent, cheap, convenient, and none of them reveal anything that is a secret as we currently understand the term "secret" -- yet the sum of them is greater than the parts. A lot greater. It might even be a polynomial, as Schneier suggested. Time will tell, but by then the game will be over.

Harvard Business School Prof. Shoshanna Zuboff has had much to say on these topics since the 1980s, especially her Three Laws:[ZS]

. Everything that can be automated will be automated

. Everything that can be informated will be informated

. Every digital application that can be used for surveillance and control will be used for surveillance and control

I think she is right, but the implication that this is all outside the control of the citizen is not yet true. It may get to be true, but in so many words that is why I am standing here. There are a million choices the individual person, or for that matter the free-standing enterprise, can take and I do not just mean converting all your browsing over to Tor.

Take something mundane like e-mail: One might suggest never sending the same message twice. Why? Because sending it twice, even if encrypted, allows a kind of analysis by correlation that cannot otherwise happen. Maybe that's too paranoid, so let's back off a little. One might suggest that the individual or the enterprise that outsources its e-mail to a third party thereby creates by itself and for itself the risk of silent subpoenas delivered to their outsourcer. If, instead, the individual or the enterprise insources its e-mail then at the very least it knows when its data assets are being sought because the subpoena comes to them. Maybe insourcing your e-mail is too much work, but need I remind you that plaintext e-mail cannot be web-bugged, so why would anyone ever render HTML e-mail at all?

Take software updates: There is a valid argument to make software auto-update the norm. As always, a push model has to know where to push. On the other hand, a pull model must be invoked by the end user. Both models generate information for somebody, but a pull model leaves the time and place decisions to the end user.

Take cybersecurity technology: I've become convinced that all of it is dual use. While I am not sure whether dual use is a trend or a realization of an unchanging fact of nature, the obviousness of dual use seems greatest in the latest technologies, so I am calling it a trend in the sense that the straightforward accessibility of dual use characteristics of new technology is itself a growing trend. Leading cybersecurity products promise total surveillance over the enterprise and are, to my mind, offensive strategies used for defensive purposes. A fair number of those products not only watch your machine, but take just about everything that is going on at your end and copies that to their end. The argument for doing so is well thought out -- by combining observational data from a lot of places the probability of detection can be raised and the latency of countermeasure can be reduced. Of course, there is no reason such systems couldn't be looking for patterns of content in human readable documents just as easily as looking for patterns of content in machine readable documents.

Take communications technology: Whether we are talking about triangulating the smartphone using the cell towers, geocoding the Internet, or forwarding the GPS coordinates from onboard equipment to external services like OnStar, everyone knows that there is a whole lot of location tracking going on. What can you do to opt out of that? That is not so easy because now we are talking not about a mode of operation, like whether to insource or outsource your e-mail, but a real opt-in versus opt-out decision; do you accept the tracking or do you refuse the service? Paraphrasing Zittrain's remark about being a customer or being a product, the greater the market penetration of mobile communications, the more the individual is either a data source or a suspect.

Take wearable computing: Google Glass is only the most famous. There've been people working on such things for a long time now. Folks who are outfitted with wearable computing are pretty much identifiable today, but this brief instant will soon pass. You will be under passive surveillance by your peers and contacts or, to be personal, some of you will be surveilling me because you will be adopters of this kind of technology. I would prefer you didn't. I am in favor neither of cyborgs nor chimeras; I consider our place in the natural world too great a gift to mock in those ways.

When it comes to ranking programs for how well they can observe their surroundings and act on what they see without further instructions, Stuxnet is the reigning world heavyweight champion. Unless there is something better already out there. Putting aside the business of wrecking centrifuges, just consider the observational part. Look at other malware that seems to have a shopping list that isn't composed of filenames or keywords but instead an algorithm for rank-ordering what to look for and to exfiltrate documents in priority order. As with other democratizations of technology, what happens when that kind of improvisation, that kind of adaptation, can be automated? What happens when such things can be scripted?

For those with less gray hair, once upon a time a firewall was something that created a corporate perimeter. Then it was something that created a perimeter around a department. Then around a given computer. Then around a given datum. In the natural world, perimeters shrink as risk grows -- think a circle of wildebeeste with their horns pointed outward, the calves on the inside, and the hyenas closing in. So it has been with perimeters in the digital space, a steady shrinking of the defensible perimeter down to the individual datum.

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. Suppose that observed data does kill both privacy as impossible-to-observe and privacy as impossible-to-identify, then what might be an alternative? If you are an optimist or an apparatchik, then your answer will tend toward rules of 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 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 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.

And how can you tell if the code you are running is collecting on you or, for that matter, if the piece of code you are running is collecting on somebody else? If your life is lived inside the digital envelope, how do you know that this isn't The Matrix or The Truman Show? Code is certainly getting bigger and bigger. A nameless colleague who does world class static analysis said that he "regularly sees apps that are over 2 GB of code" and sees "functions with over 16K variables." As he observes, functions like that are machine written. If the code is machine written, does anyone know what's in it? The answer is "of course not" and even if they did, malware techniques such as return-oriented-programming can add features after the whitelist-mediated application launch. But I'm not talking here about malware, I am talking about code that you run that you meant to run and which, in one way or another, is instrumented to record what you do with it. Nancy Pelosi's famous remark[NP] about her miserable, thousand page piece of legislation, "We have to pass the bill so that you can find out what is in it" can be just as easily applied to code: it has become "We have to run the code so that you can find out what is in it."

That is not going to change; small may be beautiful but big is inevitable.[BI] A colleague notes that, with the cloud, all pretense of trying to keep programs small and economical has gone out the window -- just link to everything because it doesn't matter if you make even one call to a huge library since the Elastic Cloud (or whatever) charges you no penalty for bloat. As such, it is likely that any weird machine[SB] within the bloated program is ever more robust.

Mitja Kolsek was who made me aware of just how much the client has become the server's server. Take Javascript, which is to say servers sending clients programs to execute; the HTTP Archive says that the average web page now makes out-references to 16 different domains as well as making 17 Javascript requests per page, and the Javascript byte count is five times the HTML byte count.[HT] A lot of that Javascript is about analytics which is to say surveillance of the user experience (and we're not even talking about Bitcoin mining done in Javascript that you can embed in your website.[BJ])

So suppose everybody is both giving and getting surveillance, both being surveilled and doing surveillance. Does that make you an intelligence agent? A spreading of technology from the few to the many is just the way world works. There are a hundred different articles from high-brow to low- that show the interval between market introduction and widespread adoption of technology has gotten shorter as technology has gotten more advanced. That means that technologies that were available only to the few become available to the many in a shorter timeframe, i.e., that any given technology advantage the few have has a shorter shelf-life. That would mean that the technologies that only national laboratories had fifteen years ago might be present among us soon, in the spirit of William Gibson's famous remark that the future is already present, just unevenly distributed. Or maybe it is only ten years now. Maybe the youngest of you in this room will end up in a world where what a national lab has today is something you can look forward to having in only five year's time. Regardless of whether the time constant is five or ten or even fifteen years, this is far, far faster than any natural mixing will arrange for even distribution across all people. The disparities of knowledge that beget power will each be shorter lived in their respective particulars, but a much steeper curve in the aggregate.

Richard Clarke's novel _Breakpoint_ centered around the observation that with fast enough advances in genetic engineering not only will the elite think that they are better than the rest, they will be.[RC]
I suggest that with fast enough advances in surveillance and the inferences to be drawn from surveillance, that a different elite will not just think that it knows better, it will know better.
Those advances come both from Moore's and from Zuboff's laws, but

more importantly they rest upon the extraordinarily popular delusion that you can have freedom, security, and convenience when, at best, you can have two out of three.

At the same time, it is said that the rightful role of government is to hold a monopoly on the use of force. Is it possible that in a fully digital world it will come to pass that everyone can see what once only a Director of National Intelligence could see? Might a monopoly of force resting solely with government become harder to maintain as the technology that bulwarks such a monopoly becomes democratized ever faster? Might reserving force to government become itself an anachronism? That is almost surely not something to hope for, even for those of us who agree with Thomas Jefferson that the government that governs best is the government that governs least. If knowledge is power, then increasing the store of knowledge must increase the store of power; increasing the rate of knowlege acquisition must increase the rate of power growth. All power tends to corrupt, and absolute power corrupts absolutely,[LA] so sending vast amounts of knowledge upstream will corrupt absolutely, regardless of whether the data sources are reimbursed with some pittance of convenience. Every tax system in the world has proven this time and again with money. We are about to prove it again with data, which has become a better store of value than fiat currency in any case.

Again, that power has to go somewhere. If you are part of the surveillance fabric, then you are part of creating that power, some of which is reflected back on you as conveniences that actually doubles as a form of control. Very nearly everyone at this conference is explicitly and voluntarily part of the surveillance fabric because it comes with the tools you use, with what Steve Jobs would call your digital life. With enough instrumentation carried by those who opt in, the person who opts out hasn't really opted out. If what those of you who opt in get for your role in the surveillance fabric is "security," then you had better be damnably sure that when you say "security" that you all have close agreement on precisely what you mean by that term.

And this is as good a place as any to pass on Joel Brenner's insight:[JB]

> During the Cold War, our enemies were few and we knew who they were. The technologies used by Soviet military and intelligence agencies were invented by those agencies. Today, our adversaries are less awesomely powerful than the Soviet Union, but they are many and often hidden. That means we must find them before we can listen to them. Equally important, virtually every government on Earth, including our own, has abandoned the practice of relying on government-developed technologies. Instead they rely on commercial off-the-shelf, or COTS, technologies. They do it because no government can compete with the head-spinning advances emerging from the private sector, and no government can afford to try. When NSA wanted to collect intelligence on the Soviet government and military, the agency had to steal or break the encryption used by them and nobody else. The migration to COTS changed that. If NSA now wants to collect against a foreign general's or terorist's communications, it must break the same encryption you and I use on our own devices... That's why NSA

> would want to break the encryption used on every one of those media. If it couldn't, any terrorist in Chicago, Kabul, or Cologne would simply use a Blackberry or send messages on Yahoo! But therein lies a policy dilemma, because NSA could decrypt almost any private conversation. The distinction between capabilities and actual practices is more critical than ever... Like it or not, the dilemma can be resolved only through oversight mechanisms that are publicly understood and trusted -- but are not themselves ... transparent.

At the same time, for-profit and not-for-profit entites are collecting on each other. They have to, even though private intelligence doubtless leads directly to private law. On the 6th of this month, the Harvard Kennedy School held a conference on this very subject; let me read just the first paragraph:[HKS]

> In today's world, businesses are facing increasingly complex threats to infrastructure, finances, and information. The government is sometimes unable to share classified information about these threats. As a result, business leaders are creating their own intelligence capabilities within their companies.

In a closely related development, the international traffic in arms treaty known as the Wassenaar Agreement, was just amended to classify "Intrusion Software" and "Network Surveillance Systems" as weapons.[WA]

So whom do you trust? Paul Wouters makes a telling point when he says that "You cannot avoid trust. Making it hierarchical gives the least trust to parties. You monitor those you have to trust more, and more closely."[PW] As I've done with privacy and security, I should now state my definition of trust, which is that trust is where I drop my guard, which is to say that I only trust someone against whom I have effective recourse. Does that mean I can only trust those upon whom I can collect? At the nation state level that is largely the case. Is this the way Brin's vision will work itself out, that as the technology of collection democratizes, we will trust those we can collect against but within the context of whatever hierarchy is evolutionarily selected by such a dynamic?

It is said that the price of anything is the foregone alternative. The price of dependence is risk. The price of total dependence is total risk. Standing in his shuttered factory, made redundant by coolie labor in China, Tom McGregor said that "American consumers want to buy things at a price that is cheaper than they would be willing to be paid to make them." A century and a half before Tom, English polymath John Ruskin said that "There is nothing in the world that some man cannot make a little worse and sell a little cheaper, and he who considers price only is that man's lawful prey." Invoking Zittrain yet again, the user of free services is not the customer, he's the product. Let me then say that if you are going to be a data collector, if you are bound and determined to instrument your life and those about you, if you are going to "sell" data to get data, then I ask that you not work so cheaply that you collectively drive to zero the habitat, the lebensraum, of those of us who opt out. If you remain cheap, then I daresay that opting out will soon require bravery and not just the quiet tolerance to do without digital bread and circuses.

To close with Thomas Jefferson:

> I predict future happiness for Americans, if they can prevent
> the government from wasting the labors of the people under the
> pretense of taking care of them.

There is never enough time. Thank you for yours.

[NAS] "Professionalizing the Nation's Cyber Workforce?"
www.nap.edu/openbook.php?record_id=18446

[PB] _Against the Gods_ and this 13:22 video at
www.mckinsey.com/insights/risk_management/peter_l_bernstein_on_risk
...Bernstein was himself quoting Elroy Dimson and Paul Marsh from
their 1982 paper, "Calculating the Cost of Capital"...

[PHI] Personal Health Information, abbreviated PHI

[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

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

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

[DA] "What is the next step?," Dave Aitel, 18 February 2014
seclists.org/dailydave/2014/q1/28

[S] Sensity's NetSense product, to take one (only) example
www.sensity.com/our-platform/our-platform-netsense

[M] For example, the 2007 collapse of I-35 in Minneapolis.

[J] "Quis custodiet ipsos custodes?," Juvenal, Satire VI ll.347-348

[DB1] _The Transparent Society_, David Brin, Perseus, 1998
[DB2] "The Myth of the 'Transparent Society'," Bruce Schneier
www.wired.com/politics/security/commentary/securitymatters/2008/03/securitymatters_0306
[DB3] "Rebuttal," David Brin
www.wired.com/politics/security/news/2008/03/brin_rebuttal

[W] minor quotation from
en.wikipedia.org/wiki/The_Transparent_Society

[TF] _Fooled by Randomness_, Nassim Taleb, Random House, 2001

[TE] "Coming to an office near you," The Economist, 18 January 2014

cover/lead article, print edition

[ZS] "Be the friction - Our Response to the New Lords of the Ring," 6 Jun 2013

www.faz.net/aktuell/feuilleton/the-surveillance-paradigm-be-the-friction-our-response-to-the-new-lords-of-the-ring-12241996.html

[NS] National Strategy for Trusted Identities in Cyberspace, 2011

www.nist.gov/nstic

[NP] 2010 Legislative Conf. for the National Association of Counties

[BI] "Small Is Beautiful, Big Is Inevitable," IEEE S&P, Nov/Dec 2011

geer.tinho.net/ieee/ieee.sp.geer.1111.pdf

[SB] LANGSEC: Language-theoretic Security

www.cs.dartmouth.edu/~sergey/langsec/

[HT] Trends, HTTP Archive

www.httparchive.org/trends.php

[BJ] Bitcoin Miner for Websites

www.bitcoinplus.com/miner/embeddable

[RC] _Breakpoint_, Richard Clarke, Putnam's, 2007

[LA] "All power tends to corrupt and absolute power corrupts
absolutely. Great men are almost always bad men, even when they
exercise influence and not authority: still more when you superadd
the tendency or the certainty of corruption by authority."

-- Lord John Dalberg Acton to Bishop Mandell Creighton, 1887

[JB] "NSA: Not (So) Secret Anymore," 10 December 2013

joelbrenner.com/blog

[HKS] Defense and Intelligence: Future of Intelligence Seminars

belfercenter.ksg.harvard.edu/events/6230/intelligence_in_the_private_sector

[WA] "International Agreement Reached Controlling Export of Mass

and Intrusive Surveillance," 9 December 2013

oti.newamerica.net/blogposts/2013/international_agreement_reached_controlling_export_of_mass_and_intrusive_surveillance

[PW] "You Can't P2P the DNS and Have It, Too," Paul Wouters, 9 Apr 2012

nohats.ca/wordpress/blog/2012/04/09/you-cant-p2p-the-dns-and-have-it-too

=====

this and other material on file under geer.tinho.net/pubs