From:	Brooks, Sean W.
To:	Soutar, Colin (US - Arlington); Jankowski, Meredith (US - Arlington); O"Shaughnessy, Logan (US - Arlington)
Cc:	Garcia, Michael; Lightman, Suzanne; Nadeau, Ellen M.
Subject:	Comments from EFF
Date:	Wednesday, October 15, 2014 3:18:38 PM

Sent to us personally, but they've Ok'd the posting of these online.

_____

From: Seth David Schoen <schoen@eff.org>
Sent: Wednesday, October 15, 2014 1:20:36 PM
To: Brooks, Sean W.
Cc: schoen@eff.org; Lefkovitz, Naomi B.; Lightman, Suzanne
Subject: Re: Thank You

Dear Sean, Naomi, and Suzanne:

Thanks for the opportunity to participate in the recent discussions about the draft privacy engineering objectives and risk model and the invitation to comment on them. These informal comments are submitted on behalf of the Electronic Frontier Foundation by Seth Schoen and Lee Tien.

I. The privacy risk equation

As we discussed at the workshop, the privacy risk equation gives the appearance of being quantitative and even of allowing users to quantify risk in an actuarial sense. It's possible to conceive of the equation

Personal Information Collected or Generated * Data Actions Performed on that Information * Context = System Privacy Risk

in a semi-quantitative way by thinking of the System Privacy Risk as something like "expected total loss or harm", and each factor as measuring the amount of personal information collected, the duration or frequency of use or retention, the risk of adverse event per use or per time period, and the expected magnitude of loss or harm per adverse event. We don't think that NIST intended to suggest that such quantitative calculations can readily be made, particularly given uncertainties about the quantitative measurement of risk or harm to individuals. It might be useful to clarify how quantitative the risk equation is meant to be; if it is indeed not intended in a quantitative sense, it might be clearer not to call it an "equation" and simply describe it as a set of factors that interact to determine the likelihood and the severity of the privacy risks that a system will create.

It can still be extremely valuable for the developers of a system to explicitly think through the factors that affect the likelihood, extent, and severity of privacy risks from that system, even if they can't quantify those risks precisely.

By contrast, there are models of privacy and privacy protection in which fully quantitative equations occur that can yield explicit numerical solutions. For example, differential privacy (pursued by Dwork et al. since 2006) provides a fully quantitative analysis of the privacy impact of some particular actions. However, the scope of differential privacy is much narrower because it defines privacy only in a particular interactive context involving the analysis of aggregate data

sets, and because it requires the deliberate introduction of errors to create uncertainty about the content of a database. Although research in differential privacy shows what a formal definition of privacy and an associated quantitive model can look like, we don't think that NIST intended either of those things here -- and if we're right, it might be worth more clearly disclaiming them.

A fairly significant source of disagreement and uncertainty at the San Jose workshop centered on whose perspective should be taken when talking about and analyzing privacy risks: risks, harms, or inappropriate actions to whom and from whose point of view? It should be abundantly clear that different parties can disagree about whether data uses and outcomes are appropriate or inappropriate, harmful or harmless.

It may not be possible to eliminate such disagreements. Consider the 2004 disclosure by Yahoo! China of information about Shi Tao, a journalist who was a user of its services, to Chinese authorities. As a result of these disclosures, the authorities identified Mr. Shi as a source for U.S. reporting on efforts to limit Chinese journalists from covering the 15th anniversary of the 1989 Tiananmen protests; he was brought to trial for his role, convicted, and subsequently served over eight years in prison. It would be possible to view Yahoo! China's disclosures as proper as a matter of applicable law and policy, and as foreseeable to users at the time as a result of terms of service as well as prior practice. This view is uncommon in the United States, where it's typical to view Chinese government restrictions on journalists as a major unjustice deserving of the widest possible exposure and criticism. Indeed, many in the United States criticized the company not only for complying with Chinese law enforcement inquiries but for maintaining records that would be responsive to those inquiries.

II. Perspective and subjectivity

We believe that the most sensible resolution to this problem is to try to take data subjects' point of view and to try to assess what privacy, appropriate or problematic uses, privacy harms, and so on mean to the data subjects.

Ignoring the data subjects' perspective as either irrelevant or unknowable risks reducing privacy to a matter of disclosure or regulatory compliance alone. This doesn't really capture what's meant by privacy in everyday speech, among privacy practioners, or in academic analysis. Instead, the notions of privacy we deal in every day are intrinsically tied to what data subjects believe and experience.

Some workshop participants were extremely resistant to formulations that relied on analyzing subjective beliefs, expectations, values, or experiences of an inevitably diverse population. We agree that there is no surefire way to do this mechanically, and that there is no guarantee that any attempt to predict or summarize the beliefs of a population of users can capture the beliefs and attitudes of each of them.

People can vary widely, but that may be all the more reason to try to think broadly about their possible concerns. A service may be used by particularly vulnerable populations, like domestic violence victims, who have quite different risk profiles from other users. Something that seems normal and innocuous to some users (like having their home phone number listed in a telephone directory) may be seriously concerning and threatening to others. Police officers or judges, too, may feel

very strongly about the privacy of some contact information that others readily regard as public. Accounting for this kind of diversity can be challenging.

Nonetheless, we think that the core goal of attempting to take the data subject's point of view is essential in order to avoid hollowing out the concepts of "privacy" and "privacy engineering". (Two habits risk being sheltered by such a narrowed concept: first, the idea that any practice is appropriate, or any privacy risk is mitigated, or any privacy harm definitionally no longer exists, if the practice has been clearly disclosed; second, the idea that that any practice is appropriate or is not a privacy harm if that practice is a part of an enterprise's business model.) More broadly, people who are processing data have a natural tendency to take their own side and to assume their own activities are benign, legitimate, and appropriate. It's useful for privacy assessment purposes, then, to explicitly distinguish between privacy risks as seen from a firm's point of view and privacy risks as seen from the data subjects' point of view.

Significantly, a firm's business model should not generally be relevant to assessing privacy risk in this sense.

We also think it can be important to consider data subjects' potentially limited knowledge and understanding, and the ways in which not all data subjects are prepared to evaluate what is really likely to happen to their data even under the terms of a published policy. For example, empirical investigations have found that most data subjects assume that firms they deal with have substantial pre-existing confidentiality obligations, or that the existence of any privacy policy at all is evidence of such obligations.

The workshop made clear that the privacy engineering goals that NIST is trying to promote are based around improving understanding of risk, and that NIST is not proposing to require any entity to do anything in particular in response to risks that it identifies. In other words, NIST is emphasizing that privacy risks need to be understood clearly, but not proposing how they should be mitigated, accepted, or rejected by any particular organization. To this end, the draft should more clearly acknowledge the importance of assessing risks from the perspective of, and relative to the goals of, data subjects. That is the perspective from which the most serious harms will occur and the perspective that requires the most conscious effort for data-processing organizations.

III. Surveillance

We think that the draft's use of the term "surveillance" to describe data uses that are in some sense disproportionate is idiosyncratic and rather different from the way this term is used in other contexts. We worry that this choice of words may hinder understanding and we suggest finding a different term.

A further difficulty is that a main connotation of this term today is cases where privacy information was "pulled" by some entity, often a government agency, hungry to know everything about everyone. But some uses that may be inappropriate, disproportionate, or risky were actually "pushes" by an organization like a medical provider or Internet firm that wanted other organizations, researchers, or even the general public to have access to data. It's hard to reconcile intentional yet inappropriate publication of personal data with our common notions of

"surveillance", and this may be another reason that the term is best replaced with another.

IV. Role of information security

The draft and presentations emphasized that there is a degree of overlap between privacy and security, definitionally and as a matter of the measures to be used to protect both.

> Confidentiality is preserving authorized restrictions on information access and disclosure, including means for protecting personal privacy and proprietary information. This definition is the same definition as the security objective Confidentiality in NIST SP 800-53 Revision 4.

We agree; in fact, we fear that the degree of overlap has been understated or underemphasized, at least thinking in terms of real-world breaches and threat environments.

Although the draft does state that appropriate security measures need to be devised and implemented as a part of privacy engineering -- in order to achieve the goal of "Confidentiality" -- it regards the choice and significance of these measures as basically the domain and territory of a separate framework. That may be correct in a bureaucratic sense, but it risks leaving privacy engineers (or others evaluating their work) with the wrong impression.

The reality we see is one in which security breaches are overwhelmingly common as a cause or risk factor for privacy harms: a reality in which a huge number of widely-publicized and significant privacy breaches were accomplished entirely through technical means that defeated security controls, and the rate and scale of such breaches seems to be accelerating. (There may also be even more pervasive or serious breaches that have never been detected or reported.) These breaches are a huge part of real-world losses of trust and other harms experienced by data subjects and a basic part of the motivation for increased interest in privacy protections.

That means that privacy engineers shouldn't assume that they can simply leave security considerations to others and count on effective measures being taken. Security breaches today often completely undermine and circumvent what were otherwise well-thought-out and well-specified privacy controls. It also means that privacy engineering should explicitly take into account what will happen if security controls are breached or defeated, and consider measures such as limiting data retention that can mitigate the damage of security breaches.

Privacy and security engineers must work together -- and ideally the model will explicitly call for them to do so. Privacy engineers need to understand security risks, and security engineers need to understand privacy risks. And they should consider together the risks that will occur if technical controls fail, and ways of mitigating the harms that will occur in this case.

A system that transmits information about its users unencrypted must recognize the associated security risk; a system that incorporates obsolete and unsupported operating systems must recognize the associated security risk. And that's true even if these systems otherwise have strong procedural or policy controls related to privacy. There are some obvious examples where choices of technology directly impact risk, but

there's also a general obligation to understand and study and examine technical controls and their efficacy.

--

Seth Schoen <schoen@eff.org>
Senior Staff Technologist https://www.eff.org/
Electronic Frontier Foundation https://www.eff.org/join
815 Eddy Street, San Francisco, CA 94109 +1 415 436 9333 x107