Accountability as a Driver of Innovative Privacy Solutions

Joan Feigenbaum[1]
Yale Computer Science Department
http://www.cs.yale.edu/homes/jf/

October 2010

Abstract: The standard technical approach to privacy in particular and computer security in general is *preventive*: Before someone can access confidential data or take any other action that implicates privacy or security, he should be required to prove that he is authorized to do so. As the scale and complexity of online activity has grown, it has become increasingly apparent that the preventive approach is inadequate. It is our thesis that a paradigm shift to *accountability*, rather than prevention, as an organizing principle for privacy in online interaction could spark much needed innovation.

Introduction

In the offline world of face-to-face interactions, physical objects, and analog communications, it is both obvious and widely accepted that people are physically able to break the law and, more generally, to violate rules and norms. With respect to offline privacy, it is often technically feasible for people to eavesdrop on others' private conversations, to violate confidentiality agreements, or to use sensitive information for purposes other than those for which it was collected. When people resist the temptation to violate privacy in these ways, it is typically because they know that doing so would be wrong or because they do not want to take the risk of getting caught and being answerable to the authorities or to their victims.

By contrast, the utopian dream of many cryptography and security researchers is an electronic world in which people are unable to break the rules. With proper use of encryption, authentication, digital signatures, and other security technology, it should be technically infeasible for people to read others' confidential communication, access others' computers and networks, distribute others' copyright material, etc. without permission. Thus, our basic technical approach to online privacy and security has been a *preventive* one: Before someone can take an action that might violate a privacy or security policy, he is required to prove that he is authorized to take it. Just as importantly, people are exhorted to protect their sensitive data and other valuable electronic resources either by refusing to reveal them ("don't give your social-security to anyone except your employer, your accountant, or the Social Security Administration") or guarding them with passwords, firewalls, etc. Although the full utopian dream of an electronic world in which bad things can't happen has never been realized, this

[1] Supported in part by NSF grant 1016875.

preventive approach works reasonably well in some environments – particularly those in which the amount of sensitive information is relatively small, the approved uses of it are few and predictable, and the people involved are all part of a stable organization in which they play well defined roles.

In today's world of Internet commerce, social networking, web-accessible health records, personalized search, and many other ways to engage socially, economically, and intellectually with numerous strangers online, preventive mechanisms are grossly inadequate. As a result, a growing faction in the cryptography and security community has embraced greater reliance on *accountability* mechanisms: When an action occurs, it should be possible to determine (perhaps after the fact) whether a rule has been violated and, if so, to punish the violators in some way. In the case of sensitive personal information, rather than exhorting people not to share it with others even when they will clearly benefit in the short run from doing so, we should develop mechanisms for holding those who receive such information accountable for the ways in which they use it. Doing so would make the online world more like the offline world, in which potential violations of security and privacy are often deterred by the prospect of negative consequences rather than prevented by truly unbreakable locks.

In the following sections, we provide examples of scenarios that call for an accountability approach, discussion of technical barriers to realizing accountability in networked interactions (the existence of which demonstrates the need for innovation!), and an observation about the term ``accountability.''

Examples

Ex. 1: Contextual integrity of information. Weitzner et al.[2] provide a crisp example of the inadequacy of today's "hide-it-or-lose-it" approach to sensitive information. Consider a woman whose online activity (searches, chat-room participation, blog posts, etc.) reveals the fact that she has a disabled child; when she applies for a job and is rejected, she suspects that, although the employer did not ask about the health of her children, he concluded from her online activity that she would be a heavy user of family-medical benefits. If we believe that the employer's action is unfair, what should be done about it? Should the mother have been encouraged (and empowered by technology) to hide her identity while engaging in online activity related to her child's disability, or should the employer be held accountable for misuse of personal information about a job applicant? The accountability approach recognizes the mother's first-amendment right to freedom of speech and freedom of assembly; she should not be forced to conceal her desire to inform herself and to help her child in order to have a fair crack at a job for which she is qualified. Analogous (and even stronger) arguments can be made about race, gender, and religion. During a face-to-face interview, a job applicant will inevitably reveal his or her race and gender, and clothing might reveal religious affiliation; nonetheless, anti-discrimination laws forbid the use of this information in hiring

[2] Daniel J. Weitzner, Harold Abelson, Tim Berners-Lee, Joan Feigenbaum, James A. Hendler, and Gerald J. Sussman, "Information accountability," *Communications of the ACM*, 51(6):82-87, 2008.

decisions. As more and more daily activity is mediated by computers and networks, it is becoming more difficult to hide such demographic information online; rather than relying solely on "privacy-enhancing technology" to hide it, we seek accountability mechanisms that ensure its proper use.

Ex. 2: Surveillance. Law-enforcement and intelligence agencies need warrants in order to eavesdrop on US persons but not on foreign nationals. Of course, it is infeasible even to determine the endpoints of many Internet traffic streams, much less to determine whether the sender and recipient are US persons. If one takes a purely preventive approach to privacy, it seems that there is an irreconcilable conflict between eavesdropping (on Internet communications, anyway) and fourth-amendment protection. Taking an accountability approach, the conflict could be reconciled if there were a way to grant law-enforcement or intelligence agents temporary access for the purpose of determining whether the sender and recipient are US persons; an agent who determined that a warrant is needed would be required to go to court to get one and not to use the traffic he eavesdropped on temporarily for any other purpose. How to hold agents accountable for proper use of such traffic is not (technologically) obvious but not necessarily impossible.

Ex. 3: Web search. Search plays a central role in a user's Internet activity, and thus users reveal a great deal of personal information to search engines. Google, in particular, is well aware that some people are disturbed by that fact and has striven mightily to convince them that it handles this information properly. Rather than rely solely on this type of "self regulation," it would be highly desirable to have mechanisms for holding search companies accountable for their uses of personal information.

Technical Challenges

One basic challenge confronting researchers in this area is terminological confusion; although everyone agrees that "accountability" sounds like a good thing, there is no standard technical definition of the term, and indeed different research communities use it to mean different things. The following example is illustrative.

"Accountability" is clearly related to "identifiability" and "anonymity," but the precise relationships of these concepts have yet to be determined. In the aforementioned paper of Weitzner et al., which is based on extensive experimental work by the MIT Decentralized Information Group[3], it is assumed that each entity in an "accountable information system" has a persistent identity. If an event in such a system gives rise to a privacy dispute, the parties to that event are subject to an adjudication procedure the results of which can be enforced; anonymity is not a goal in these systems. By contrast, work on accountability in the cryptographic-protocol literature often strives to enable anonymous interaction and stresses that, although all entities have persistent identities, their identities need not be revealed to all others with whom they interact. Novel protocols for, e.g.,

[3] http://dig.csail.mit.edu

electronic cash[4] and anonymous broadcast[5] provide accountability precisely by preserving anonymity of all parties that follow the protocol and exposing the identities of those who deviate from it. Whether the notion of accountability is meaningful in systems in which the parties do not have persistent identities is an open question.

A second basic challenge presented by the accountability approach to privacy and security is the difficulty of determining information provenance and enforcing information policies. For example, the privacy policies of many organizations assure users of their websites that information provided by those users will not be shared with third parties. Unfortunately, there is no standard method for checking whether an organization's computer systems comply with those policies. When a third party does improperly obtain a user's personal information, it is often infeasible to reconstruct the path that the personal information traveled on its way into the wrong hands and thus infeasible to determine which organization improperly released it. More extensive network monitoring and data retention could be useful in establishing information provenance but simultaneously destructive of privacy.

Conclusion

As explained in the previous section, one of the difficulties confronted by cryptography and security researchers working in this area is terminological confusion. Even if substantial progress is made on the design and implementation of accountability mechanisms, the resulting technology could face substantial barriers to adoption precisely because of this confusion. In common English parlance, "holding someone accountable" entails identifying him and forcing him to answer for his actions. When arguing for an infrastructure that supports accountability, one is often misunderstood to be advocating an infrastructure that prevents anonymity and pseudonomy and therefore inherently destroys privacy. It might be more effective to describe the goal as design and implementation of technology that "deters misuse of sensitive information by imposing negative consequences on those who misuse it" rather than technology that "holds those who misuse sensitive information accountable for their actions."

Regardless of the terminology that we ultimately settle on, it is clear that privacy and innovation need each other in today's technological environment. Sensitive information about people and organizations is being created, stored, and exchanged in ever-increasing amounts, and people are justifiably unwilling to forgo the beneficial uses of all this information sharing in order to avoid the harmful ones. We need innovative solutions to the problem of accountability in networked information exchange.

[4] Jan Camenisch, Susan Hohenberger, and Anna Lysyanskaya, "Balancing accountability and privacy using e-cash (extended abstract)," In *Proceedings of SCN*, pages 141-155, 2006.

[5] Henry Corrigan-Gibbs and Bryan Ford, "Dissent: accountable anonymous group messaging," in *Proceedings of the 17th ACM Conference on Computer and Communication Security*, pages 340-350, 2010.