D. Crocker
Brandenburg InternetWorking
November 5, 2010

Joint Privacy Workshop: Position Comments by D. Crocker

crocker-privacy-position

[This is offered as a personal contribution to the Privacy Workshop, to be held at MIT, discussing: "How Can Technology Help to Improve Privacy on the Internet?"]

1. Privacy Space

More than other security-related functions, privacy is a peculiarly social concern. All network services have a human component, since they ultimately must satisfy human producers, consumers and operators, however indirectly. Privacy is unusual because it depends so heavily on the vagaries of human preference, human interaction and human comprehension at its core. Criteria for acceptable types and degrees of privacy are subject to human whim; today's openness is often followed by tomorrow's personal regret. A demand for strictness this morning might turn into an indignant demand for looseness in the afternoon. Worse, the granularity of control that is needed is daunting, entailing not accounts or files, but often specific fields.

A cornerstone to privacy discussions is the term "consent". Unfortunately this tends to hide a multitude of design and policy sins, rather than to describe the required informed understanding and agreement that is intended or, at least, claimed. System behaviors often are sufficiently complex to exceed the ability of average users to anticipate their effects. Explicit consent is usually sought at the wrong time in a transaction or requiring too much effort or knowledge. Worse, social and legal constraints vary widely.

Serious effort at improving technology's role with respect to privacy therefore requires focusing primarily on the human and social factors first, and then designing the technology in response. The technological concerns that dominate are meta-issues: indirection, speed and scale. Actions and effects often are not directly visible to affected users and often are the result of cascading systems interactions; they take place too quickly to constrain post hoc, and they often are propagated to the entire Internet, making problematic disclosure visible to billions of people.

So an essential observation about having technology help to improve privacy on the Internet is that the current uses of technology tend to DEGRADE privacy on the Internet. The complexity of technical systems on the Internet impedes privacy-related enforcement, coordination and diagnosis. An effort to make improvements requires juggling system complexity, human usability constraints, and broader public policy requirements. Although the vagaries of public policies are perhaps the most visibly tricky, proper engineering for the human factors is demonstrably beyond the current state of the art... At the least, this should direct efforts toward simplified, basic functions and benefits.

2. Privacy Work

The most significant improvement for working on Internet privacy is going to be development of a sufficient powerful yet sufficiently simple framework for discussing and specifying the lifecycle aspects of a service's privacy requirements, mechanisms and effects. Given such a framework, it will be possible to include reasonable consideration of privacy issues into normal technical development efforts. Here, again, the human factor dominates: engineers are people too and the framework must be comfortable for use within normal engineering discussions:

- What kinds of privacy issues does a service touch? For example, what information does it touch or produce that might have privacy concerns? This obviously includes various data fields, but often will include interactions, such as what individuals are exchanging messages with each other.
- What are the control points with possible privacy concerns? Data creation, transmission, receipt and retention are obvious examples. More challenging are critical phases of human use of the service: registration and login are obvious examples. Unsolicited notifications of privacy-related events can be warranted, much like having a credit agency verify that a transaction is authorized.
- How does privacy function as a process within a service? How is it configured, assessed and monitored?

By virtue of its complexities, ambiguities and social components, the effort to incorporate privacy-related work into engineering activities is often not very productive. The results are too vague, too simplistic or too rigid. The challenge, then, is to find tractable, pragmatic approaches that faciliate privacy considerations during engineering efforts, without creating an onerous burden.

It is worth exploring the possibility of adapting the credit card service model to basic privacy discussions when developing Internet services:

- It is a long-standing and reasonably mature model, covering a number of very sensitive issues, including bits of privacy.
- Maturation of the model has been incremental and difficult; so it reflects a plausible balancing of constraints and requirements.
- It is well enough understood and seems to be sufficiently simple to be tractable for engineers and even for users.
- It is life-cycle oriented, touching all phases of the relevant relationships, transactions, problems and repairs.
- Because the model is anchored in extensive real world practice, attempts to apply it for more general privacy issues seems likely to surface its failings in terms of pragmatics rather than abstractions.
- Unless the model proves entirely hopeless, enhancements that respond to its limitations can be incremental.

3. Collaborative Privacy

Most efforts at systemic privacy seem inclined to treat the user as expert, but secondary. Expert in terms of comprehension, but secondary in terms of initiative. While it is not reasonable to expect users to read complex legal tracts in real-time, it is worth noting that basic privacy assertions are part of daily life. Adding support for these simple assertions could be helpful.

For example in email, it is common for an author to informally mark a message as private. The recipient is trusted to understand and honor this request. It would be quite simple to have messaging tools allow adding this stricture, at author request, and then have recipient tools take note of the request, such as requesting confirmation when a recipient attempts to forward a private message. Note that this is mechanically similar to the existing Web ability to mark pages

as requesting no caching or to be skipped by crawlers.

An elaboration could be longer-term preferences that an author might have for various classes of exchanges. All mail is to be classed as private; or Mail to a particular recipient is private.

Today, human actors have no means for expressing their preferences or policies concern privacy, except at the whim or requirement of the service provider. It is perhaps time to enable assertions at the initiative of the user.

The primary requirement, here, is to look for established privacy practices in the non-technical world and then find ways to emulate them. The balance will be simplicity for the users, while still providing meaningful benefit.

Author's Address

Dave Crocker
Brandenburg InternetWorking
675 Spruce Drive
Sunnyvale, CA 94086
USA
Phone: +1.408.246.8253
EMail: dcrocker@bbiw.net