December 1, 2020

Privacy and Security: Values for the Internet

Nick Doty
UC Berkeley, School of Information

Status of This Document *This is a chapter of a published dissertation: Enacting Privacy in Internet Standards.*

3 Privacy and Security: Values for the Internet

Privacy and security are just two values among many that can be enacted within a technical design. Accessibility, accountability, archivability, fairness, free expression, internationalization, justice, neutrality, performance and many other values can all be affected by the particular technical architecture.[1] However, privacy and security have played an outsized role in the history and use of the Internet and the Web, despite the clichéd and largely inaccurate notion that the original design of the Internet ignored security. Because of the decentralized architecture of the Internet and the end-to-end property of its design, security is a challenge to achieve, while being a pre-requisite for the use of the Web for electronic commerce. Because the Internet is, most of all, an information medium that billions of people use to communicate, protecting privacy and control over the flow of personal information is a fundamental task, especially as users contribute more of their personal thoughts in increasingly popular social media applications.

We could imagine an alternative history of the World Wide Web that didn't prioritize these applications — ecommerce, personal communications, social media — one that was more limited to the accessible library of information originally imagined by Tim Berners-Lee. In theory, such a Web might see privacy or security

[1] See the "values in design" concept and the trend towards integration, as described in The Ethics of Engineering, previously.

as less fundamental issues. With fewer commercial applications, confidentiality, integrity and availability may have been less pressing properties for development; if the Web were more a reading platform than one where users generated content themselves, privacy issues, while germane, might be less inherently essential. Arguing for a necessary history of the Web from its origins to its current form is counter to good historiography; in this case, it is also unnecessary. There are reasons to support the notion that the success of the Internet and the Web made it likely that commercial applications would be developed and that without commercial applications the infrastructure would not be as substantial or as popular. As noted previously, from the earliest days of the Internet, email and personal communications were essential drivers of the infrastructure. Similarly, the architecture and history of the Web suggest that user-provided content of some form would be supported, whether through a "read-write Web"[2] or more centralized social media. As interesting as these alternative likely histories are, the fact remains that ecommerce and social media have been large, popular, driving applications of the Web, and applications that are particularly likely to involve security and privacy issues. As such, it's fruitful to look at these values, even as we recognize that other values have also been important to the development of the Web and that different values may support different applications in the future.

To begin with, let's define, or at least scope, some of the basic terminology.

3.1 Definitions and contentions

"Security" can mean many different things to different people and in different cultural contexts. While some might immediately think of the locked doors of a bank vault (an access control view), others might think of the safety of basic needs. The Japanese word "anshin" is used in some contexts as a translation of security, but describes more broadly a sense of peace of mind, tied to confidence, familiarity and knowledge (Okumura, Shiraishi, and Iwata 2013).

In the fields of network or information security, what is considered the classical model defines security as a property of a system that satisfies three objectives: confidentiality, integrity and availability (the C-I-A triad).[3] While critiques and

[2] A concept long-advocated by Tim Berners-Lee and popularized in the Read/Write Web blog (MacManus 2003), where users can contribute to web pages as easily as they browse them.

[3] The original source identifying these objectives as fundamental to security is unknown. An early reference identifying them as the most common goals of a security policy is a report from Dave Clark and David Wilson: "A Comparison of Commercial and Military Computer Security Policies" (1987). Notably, this is the same Dave Clark known for design of the Internet architecture.

extensions to the confidentiality-integrity-availability model are common, most researchers in these fields continue to rely on something similar; this research uses "security" to refer to these objectives in computer/information security unless otherwise noted.

Contentions about the definition of "security" are mild in comparison to the myriad differences over "privacy." A classical definition is that privacy is control over personal information, as presented by Alan Westin, considered a founder in the field (1967). That definition mirrors early definitions of security: an access-control approach based on satisfying a particular privacy or security policy. However, many scholars have noted limitations to this "informational privacy" definition; that it doesn't capture intrusions into our daily lives or substantively capture what is distinctive about violations of that control over information. Many practitioners rely on a concept of "fair information practices" or "fair information practice principles" (FIPPs), drawn from a Department of Health Education, and Welfare report from 1973 (Department of Health, Education and Welfare 1973) and the Organization for Economic Cooperation and Development from 1980 (Organization for Economic Cooperation and Development 1980) and still very present in the Obama administration's proposed Consumer Privacy Bill of Rights ("Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy" 2012). Philosopher Helen Nissenbaum argues for a theory of "contextual integrity" (2004) to explain our privacy views based on the flows of personal information: not absolute access control policies but instead expectations built up from social and legal norms.

You might prefer one privacy definition over another or find one more often used in a particular setting, but increasingly it seems clear that "privacy" is an *essentially-contested concept* for which we will not and should not settle on a single definition. Following the characteristics laid out by Gallie (1956), privacy is: *appraisive*, a valuable achievement; *complex*, with multiple dimensions including objectives and justifications; *open*, changing in salience over time in response to different technological and social circumstances; and, finally, subject to *progressive competition*, where ongoing debates over the concept can contribute to better understanding privacy (Mulligan, Koopman, and Doty 2016). As a practical matter, this suggests conducting research to anticipate and uncover, rather than foreclose, different approaches to privacy. And as we note in that work, contestation of privacy has important implications for design:

- debates over a single definition of privacy will not be conclusive, and so

it will be more useful to describe particular concepts or privacy goals in a particular context;[4]
- a static list will not be able to anticipate all privacy concerns, and so designers can benefit both from looking very concretely at specific user needs or concerns and at a higher-level from understanding the object of privacy and identifying where it can be supported in the technical architecture; and,
- because contestation will continue, designers should anticipate and accept this openness.

Instead of a singular definition of privacy, then, we end up with meta-analyses of privacy concepts. Dan Solove argues for a Wittgeinsteinian "family resemblance" approach and sets out a large taxonomy of different actions that might constitute privacy violations (2006). Colin Koopman and Deirdre Mulligan devise a privacy analytic to map out theories of privacy on a large number of dimensions, including the purpose of protecting privacy and from whom one is protecting their privacy (Koopman and Mulligan 2013; Mulligan, Koopman, and Doty 2016).

By necessity, then, this research does not rely on a single, narrow definition of "privacy" for its inquiry. Further, as a methodological matter, foreclosing any dispute on the definition or sense of privacy might lead to missing that same dispute within the community or communities in question. How privacy is differently defined by the engineers and other participants in technical standard-setting is itself a research question. Existing work has looked at the models of privacy evident in the work of computer scientists working in security and privacy (Danezis and Gürses 2010) and in the nascent field of privacy engineering (Gürses and Alamo 2016).

Care is taken, as a matter of research method, not to "prime" or load a particular meaning of the term "privacy" during interviews with participants.[5] This method is more than contingently important, because one possibility to be explored is that, because of the openness in response to technological change of values particularly impacted by the Internet, what privacy is may not only be debated among engineers, but materially constructed by them.

[4] See, for example, this two-year discussion of a definition of privacy and whether it's necessary or useful for IETF specification work on the `ietf-privacy` mailing list. In Do Not Track discussions, participants debated whether defining "privacy" was useful for scoping or an unimportant academic matter.

[5] See interview guide in the appendix.

For the purpose of scoping my own inquiry, I focus on privacy as the family of values related to norms and controls over flows of information about people and freedom from intrusions.

3.2 Relationship between privacy and security

Why consider privacy and security together? Aren't these separate values that need to be distinguished in order to determine the distinctive effects and factors related to privacy?

There is some truth to the common cliché that "you can't have privacy without security." That is, systems that are vulnerable to attacks that break the properties of confidentiality or integrity typically can't guarantee users control over how information about them is collected, used or disclosed. This is true in more than the most naive sense in which security is necessary for a system to provide any other value — if a system is not available, then it cannot provide any of its functionality; if a system cannot provide integrity, then it could have been altered to counter some other value for which it was designed. For example, any conception of privacy that includes keeping some information secret or controlling access to a piece of information will be undermined by violations of confidentiality: if a system is vulnerable to threats where an attacker can access information she is not intended to be able to access, then the system is less likely to provide contextual integrity or effective controls over information disclosure.

From the perspective of Internet architecture, security may be more relevant at lower layers – e.g. establishing secure channels of application-agnostic communication – while privacy may be more significant at higher layers – e.g. user controls over information disclosure in particular applications.[6] Braman identifies privacy as a topic of concern from the earliest days of Internet architecture design as described in the first ten years of RFCs, with confidentiality and access control of particular importance for protecting information on hosts or transmitted through the network (2012).

In addition to security as a pre-requisite for (or lower layer to) privacy, there are also cases where privacy and security overlap. One common reason for conflating security and privacy is the assumption that privacy *just is* confidentiality. It's popular to claim that this conflation is simply erroneous; however, if we accept that privacy is plural and essentially contested, it's more difficult to flatly discount

[6] This may be a common perspective, but I'm not sure whether it's published or documented as a design principle. I attribute it to presentations by Alissa Cooper.

such a theory. What we can say is that most typical definitions or theories of privacy include protections beyond confidentiality. That is, privacy-as-confidentiality is an uncommonly *narrow* conception of privacy. That said, those same typical definitions (including both control over personal information and contextual integrity) would count many confidentiality violations as prototypical violations of privacy: many concepts of privacy *include* confidentiality.

For example, the National Institute of Standards and Technology (NIST), in seeking to improve and systemize the engineering practices for privacy, has drafted an evolving set of privacy engineering objectives, to serve a similar functional purpose to the C-I-A triad. In initial drafts, the list of objectives included confidentiality (as defined in related security engineering documentation) to explicitly mark an overlap between privacy and security (NIST 2014).[7]

Distinct from the layer model (privacy on top of security) described above, there might also be cases where a lack of privacy undermines security. Designs for security that rely on trust in participants might have a vulnerability if the personal privacy of an individual is compromised. For example, the confidentiality of classified information depends on the reliability and lack of coercion of those cleared to receive that information; if the intimate details of a person's life are accessed, a blackmailer may be able to obtain government secrets.[8] Similarly, some authentication mechanisms rely on limited flows of information about a person; if an attacker can unexpectedly easily determine your birthdate and addresses of previous residences, they may be able to impersonate you to your bank.

While there are substantive connections between privacy and security in the design of Internet protocols, an additional motivation to consider these values together is their integration in the *practice* of privacy and security engineering work. As later sections will demonstrate, the work of identifying and mitigating privacy concerns and security concerns share techniques (like threat modeling), expertise and people. Even if values can be, conceptually, separately defined and considered, if the engineering efforts are themselves combined, then understanding and improving the practice of privacy and security engineering requires exploring the values together.

As an empirical matter, efforts for coordinated security and privacy review have become more integrated in recent years. One explanation is that, in addition to the inherent connections between accepted security properties and common

[7] Partly in response to public comments, a subsequent draft finding uses "disassociability" instead, with a definition distinct from confidentiality, and more like "unlinkability" (NIST 2015).

[8] h/t Daniel Griffin

conceptions of privacy, the historical context of a changing political and technological atmosphere has shifted privacy to depend more deeply on traditional security objectives. That openness is a piece of privacy's essential contestedness. In a historical review of privacy, we can note how privacy (at least in Western society) was broadly conceived in the late 19th century as a freedom from harassment or publicity – a response to photography and newspapers; and in the mid-20th century as a concern about unfair or unaccountable analysis in newly-available large, computerized databases. I believe we can see a similar shift over shorter time-frames in the conception of online privacy in the 21st century. When a plethora of online tracking mechanisms and corresponding behavioral advertising companies appeared in the early 2000s, the privacy concern of protection from corporate profiling was heightened; after the Snowden revelations in 2013, a shift in effort and attention was made towards privacy from large-scale government surveillance and securing infrastructure. That a concern was heightened during a particular time doesn't imply that it was absent otherwise; government surveillance was not a wholly new concern after 2013 and online corporate data collection remains a privacy issue (not just because the same infrastructure is relevant to government access). Similarly, the privacy torts about unwanted publicity didn't disappear after the 20th century. But these historical shifts and the competing concepts of privacy they highlight are, I argue, reflected in the work on engineering privacy on the Internet and its increasing integration with security.

While this work focuses on privacy and security as fundamental values in tension on the Internet, what we learn from the design for these values can inform, and be informed by, research on the design of other values. In particular, there is much to learn from experiences with accessibility and internationalization; and I hope this research can contribute to work on diversity and freedom from harassment.[9]

3.3 Cases in this work

Following these shifts in the concept of privacy, let us look at two cases, with different conceptions of privacy and where there is a change, or potential change, in the distribution of responsibility for protecting privacy. In each, we can see a "handoff" in the larger socio-technical system and the manner of these shifts can help us uncover what value is being supported and how.

[9] See Directions.

First, I look at the movement to encrypt the Web, including designing, advocating for and deploying new security technology to maintain privacy from network surveillance and intrusion; and, second, I consider Do Not Track, an effort to develop a cooperative user choice mechanism for protecting privacy from online behavioral tracking, which will be the topical focus of my empirical work.

References

Braman, Sandra. 2012. "Privacy by design: Networked computing, 1969–1979." *New Media & Society* 14 (5): 798–814. https://doi.org/10.1177/1461444811426741.

Clark, D D, and D R Wilson. 1987. "A Comparison of Commercial and Military Computer Security Policies." *IEEE Symposium on Security and Privacy* 0: 184–94. https://doi.org/10.1109/SP.1987.10001.

"Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy." 2012. White House. http://www.whitehouse.gov/the-press-office/2012/02/23/fact-sheet-plan-protect-privacy-internet-age-adopting-consumer-privacy-b.

Danezis, George, and Seda Gürses. 2010. "A Critical Review of 10 Years of Privacy Technology." *Proceedings of Surveillance Cultures: A Global Surveillance Society*. http://www.researchgate.net/publication/228538295_A_critical_review_of_10_years_of_Privacy_Technology/.

Department of Health, Education and Welfare. 1973. "Records, Computers and the Rights of Citizens." https://epic.org/privacy/hew1973report/.

Gallie, W. B. 1956. "Essentially Contested Concepts." *Proceedings of the Aristotelian Society*, New Series, 56 (January): 167–98. http://www.jstor.org/stable/4544562.

Gürses, S., and J. M. del Alamo. 2016. "Privacy Engineering: Shaping an Emerging Field of Research and Practice." *IEEE Security Privacy* 14 (2): 40–46. https://doi.org/10.1109/MSP.2016.37.

Koopman, Colin, and Deirdre K Mulligan. 2013. "Theorizing Privacy's Contestability: A Multi-Dimensional Analytic of Privacy."

MacManus, Richard. 2003. "The Read/Write Web." *ReadWrite* (blog). April 20, 2003. https://readwrite.com/2003/04/19/the_readwrite_w/.

Mulligan, Deirdre K., Colin Koopman, and Nick Doty. 2016. "Privacy Is an Essentially Contested Concept: A Multi-Dimensional Analytic for Mapping

Privacy." *Phil. Trans. R. Soc. A* 374 (2083): 20160118. https://doi.org/10.1098/rsta.2016.0118.

Nissenbaum, Helen. 2004. "Privacy as Contextual Integrity." *Washington Law Review* 79 (1): 101–39. http://heinonlinebackup.com/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/washlr79&section=16.

NIST. 2014. "NIST Privacy Engineering Objectives and Risk Model Discussion Draft." April. http://csrc.nist.gov/projects/privacy%7B_%7Dengineering/nist_privacy_engr_objectives_risk_model_discussion_draft.pdf.

———. 2015. "Privacy Risk Management for Federal Information Systems." http://csrc.nist.gov/publications/drafts/nistir-8062/nistir_8062_draft.pdf.

Okumura, Kaori, Yoshiaki Shiraishi, and Akira Iwata. 2013. "Survey on Sense of Security for Registering Privacy Information to Return Refugee Supporting System." In *Symposium on Usable Privacy and Security (SOUPS)*. https://cups.cs.cmu.edu/soups/2013/trustbusters2013/Sense_of_Security_Refugee_Supporting_System_Okumura.pdf.

Organization for Economic Cooperation and Development. 1980. "Guidelines on the Protection of Privacy and Transborder Flow of Personal Data." http://www.oecd.org/internet/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm.

Solove, DJ. 2006. "A Taxonomy of Privacy." *University of Pennsylvania Law Review*, no. 477: 477–560. http://www.jstor.org/stable/10.2307/40041279.

Westin, A.F. 1967. *Privacy and Freedom*. New York: Atheneum.