Association for Computing Machinery (ACM)
ACM US Public Policy Council (USACM)

usacm.acm.org
facebook.com/usacm
twitter.com/usacm

October 17, 2014

National Privacy Research Strategy
NCO, Suite II-405
4201 Wilson Blvd.
Arlington, VA 22230

Submitted via email to: nprs@nitrd.gov

Re: Request for Information (RFI)-National Privacy Research Strategy, 79 FR 56091, Docket No. 2014-22239

The ACM U.S. Public Policy Council appreciates the opportunity to provide comments regarding a proposed National Privacy Research Strategy to be developed by the agencies of the Networking and Information Technology Research and Development (NITRD) program. Continuous investment is needed to keep pace with changes in technology with impacts on privacy, and a strategy can help guide those investments.

ABOUT ACM AND USACM

With over 100,000 members, the Association for Computing Machinery (ACM) is the world's oldest and largest educational and scientific computing society. The ACM U.S. Public Policy Council (USACM) serves as the focal point for ACM's interaction with U.S. government organizations, the computing community, and the U.S. public in all matters of U.S. public policy related to information technology. Our comments are informed by the research experience of our membership. Should you have any questions or need additional information, please contact our Public Policy Office at 212-626-0541 or at acmpo@hq.acm.org.

General Comments

The Request for Information (RFI) raises several questions that touch on privacy risks and what they mean. It would be very helpful to note in the National Privacy Research Strategy (the Strategy) what is meant by privacy risks. As we see it, privacy risks are sometimes distinct from security risks; and sometimes risks defined as security alone affect privacy more than security. While breaches of information, particularly sensitive information, involve both privacy and security, there are different kinds of risks involved, affecting different entities. In the case of a data breach, the security risks are to the systems involved in the collection and use of the data. Companies and institutions are typically more affected by security risks than individuals. For instance, the loss of a corporate password creates more security risks for the corporation than for the individual. However, individuals are more affected by privacy risks. The privacy risks are more about the consequences of the breach than the breach itself, and are exogenous to the affected computing systems. The exposure of information exposes affected parties to privacy harms. Those harms are also more variable and can be more difficult to estimate than security risks. Exposure of the same data about one person may have different impacts than the exposure of those data about another. The differences and interactions between security risks and privacy risks should be clear within the strategy, and may be an appropriate topic for further research.

Understanding those similarities, differences and interactions between those risks is a common thread in some of our comments below.

Answers to specific questions

1. Privacy Objectives

Scenarios that illustrate critical issues concerning privacy follow:

How do legal and compliance personnel communicate with system developers to ensure that relevant privacy laws are embodied in requirements for computing systems? There may be limits to how much of those laws can be formalized in system requirements. Identifying relevant gaps between what laws expect and what can be formally captured in computing systems will be important to properly manage privacy risks.

More effort is needed to capture and express privacy-relevant concepts in both law and software design in more rigorous and complete manners. System requirements must express legal and ethical privacy obligations in a manner that supports software engineers as they build and design real-world systems. There may be limits to how much of those laws can be formalized in system requirements. Identifying relevant gaps between legal obligations and existing software requirements will be important to properly manage privacy risks. The strategy needs to support research on methods for developing legally compliant requirements, including formal and semi-formal methods. The strategy should also include the application of new methods of requirements compliance to managing and mitigating privacy risks.

How, and under what conditions, do organizations go beyond privacy law to identify the "right thing to do" to ensure their privacy requirements are aligned with individual and group privacy preferences? When organizations try to develop privacy practices that match the preferences of the people and of groups they collect information from, what else guides these organizations, besides legal and/or regulatory requirements? How do they gauge the demands of those individuals and groups?

How do developers evaluate third-party services and components for compliance with privacy requirements? What practices and procedures are used to assess these service and components, and how do they differ (if they differ) from what they do to be compliant with security requirements?

How are stakeholders trained to perceive and evaluate privacy risk? Stakeholder understanding of privacy risks, based on reactions to announcements of data breaches and thefts of files from data storage services, is variable. Without adequate and current knowledge of possible dangers, mistakes and pitfalls and of ways of minimizing those risks, stakeholders, such as developers, system administrators and people, will remain unaware of how systems pose risks to personal privacy.

How are and how can systems be designed to enable a range of privacy preferences? There is little support for negotiation by privacy-aware people who would engage in services. The current opt-in/opt-out framework is typically binary in execution, rarely offering full participation by individuals with high

privacy sensitivity. Adoption of privacy-enhancing technologies or provisions to purchase privacy-enhanced versions of products has been advocated in the economics of security and privacy literature as optimal for both providers and consumers. However, this has not been proven in the marketplace.

How are interfaces designed to ensure that users can state their privacy preferences at scale? Every device and service has different product features and can have differing terminology for describing sensitive and non-sensitive data. That can make it difficult for people to express their privacy preferences in a global sense – for every service and/or device that they may participate in or use. Privacy is also contextual, so people may have different preferences depending on the service or device in question. Being able to set privacy preferences in a global sense would be helpful in situations where a change in an individual's circumstances would affect their privacy preferences. Changing those preferences for every device, service and context is burdensome. Another challenge of scale concerns the tradeoffs involved in each instance where information is released or collected. It can be challenging to effectively assess the benefits and risks when specifying privacy preferences, as the complexity of those tradeoffs aren't easily communicated.

What are the boundaries between personal and work (or personal and official) technologies and services? A workplace may allow employees to use their own device for work purposes. Employees may download non-work applications or programs onto office technology or via office networks. Public officials are often required to use official communications services (email, texting, etc.) for official business and may maintain separate private accounts for personal use. There are certainly security implications for commingling outside software and hardware with items developed and/or vetted within an organization. But such commingling may also lead to the sharing of information (intentionally or unintentionally) on private equipment or services that is meant for non-private uses, and vice versa.

2. Assessments

There should be more than one privacy paradigm to inform how we assess privacy research results. The following privacy paradigms can affect research (particularly when approaching privacy from disciplines other than computing).

Privacy as secrecy: Privacy modeled as a refusal to disclose. The goal is considered complete confidentiality.

Privacy as contextual:[1] Reflects the increasing amounts of devices and other means for disclosing information, sometimes unknowingly. Because technologies are parts of human activity in ways not previously possible (or considered), there are many new situations in which information could be disclosed that were not previously considered.

[1] Nissenbaum, H. *Privacy in Context: Technology, Policy and the Integrity of Social Life*. (Palo Alto, California: Stanford University Press, 2010).

Privacy as autonomy:[2] Considers privacy in the context of activities done with or without state surveillance. It includes the potential chilling effect of observation on these activities, and covers the 'right to be left alone.'

Economic privacy:[3] The collection of personal information can also lead to price and other forms of discrimination in providing goods and services.[4]

These and other paradigms should inform how we search for, and assess, novel technical solutions to privacy. Some argue that traditional conceptions of privacy are inadequate for current technologies.[5] In health care, considering privacy as autonomy can have value because releases of information in that field often depend on known treatment relationships (such as doctor/patient). Should patients withhold information out of concern over inappropriate disclosure, there would be consequences for their own health and for medical research.

With respect to surveillance technology, individuals resort to reserving or shielding their thoughts and communications when under unwanted surveillance through a variety of strategies.[6] How do we assess whether a surveillance technology increases reservations by causing people to disengage from public life? What are the long-term consequences to democracy? Technology plays a role in monitoring, detecting and altering the behavior of systems with respect to these paradigms and we shouldn't exclude the consequences of technology on society from the research strategy — they should be part of the strategy to assess technology. Several paradigms drive privacy risk and too much emphasis has been placed on secrecy, which is a native and comfortable territory for security researchers. By increasing the perspectives that inform privacy research we can better describe the differences between privacy and security risks, and identify other important areas where privacy requires responses distinct from what might be done from the perspective of security.

For specific tools and techniques for assessment, besides the various conceptions of Fair Information Practice Principles (FIPPs), the strategy should also consider a dataflow-based lexicon, different privacy

---

[2] Cohen, Julie E. *A Right to Read Anonymously: A Closer Look at "Copyright Management" in Cyberspace*, 28 CONN. L. REV. 981 (1995).

[3] Odlyzko, Andrew. "Privacy, Economics and Price Discrimination on the Internet." Proceedings of the 5th international Conference on Electronic Commerce. Association for Computing Machinery, 2003.

[4] The Obama Administration's 2014 Big Data report highlighted the possibility of data-enabled discrimination. http://www.whitehouse.gov/sites/default/files/docs/big_data_privacy_report_5.1.14_final_print.pdf

[5] Solove, Daniel J. *The Digital Person: Technology and Privacy in the Information Age*. NYU Press, 2004.

[6] Lyon, David. *Surveillance Studies: An Overview*. Polity, 2007.

risk models[7] (including those suggested by the paradigms discussed above), the tools being developed by the Identity Ecosystem Steering Group (IDESG),[8] and the range of technologies on the market today.

A critical aspect for assessing the effectiveness of privacy-enhancing technologies is related to quality attributes or non-functional requirements. While methods, tools and techniques may improve privacy, they may not sufficiently address usability or performance. Usability is itself an active area of research,[9] covering issues such as time delays from routing inherent in privacy-enhancing technologies, or multi-party computation being both technical and risk communication challenges.

3. Multi-disciplinary approach

Privacy is part of a socio-technical system, depending on both technical and social components in order to be effectively preserved (or violated). However, computer science majors in undergraduate programs are generally not trained in how to conduct human subjects research. Graduate students in Ph.D. programs may take one course in statistics, which is dwarfed by the 4-6 courses that a psychology Ph.D. student must take to study human beings. To understand how people experience privacy, computer science researchers need to be teamed with law, public policy, and social science disciplines, including risk analysis, anthropology and psychology. This includes interdisciplinary research, but also new training for computer science students at all levels. The people who build the software must know how to "go native" to better understand how end users experience privacy risk.

An important thing to do with this multi-disciplinary research is to make sure that computing developers and other computer scientists are able to recognize when additional disciplinary expertise is needed; and others need to learn how to communicate the results and technical implications of experiments that are based in fields outside of computer science.

4. Privacy Architecture

There are at least two kinds of architecture worth noting in responding to this question: enterprise architecture, which includes the IT resources and business processes and how these two phenomena are coordinated to achieve company or agency goals; and software architecture.

In enterprise architecture, there is a need for more traceability to understand how privacy requirements influence social and technical dimensions of an organization: how do training and audits complement

---

[7] The ACM US Public Policy Council detailed both the privacy risk models and the dataflow-based lexicon in comments to the Federal Trade Commission on February 18, 2011. http://usacm.acm.org/images/documents/FTCprivacyResponseFinal.pdf

[8] Identity Ecosystem Steering Group (IDESG). http://www.idecosystem.org

[9] The Symposium on Usable Security and Privacy (SOUPS) is in its second decade; and the Internet Society affiliated Useable Security (USEC) event is integrated with its symposium on security for networked and distributed systems.

technical solutions in privacy? How do organizations make privacy trade-off decisions when designing a system that leverages personal data to achieve company or agency goals? What are the enterprise-endogenous risks to privacy and what controls do we need to shape privacy across the entire data ecosystem?

Software architecture affords the ability to encapsulate system qualities, such as privacy, in ways that can be used to reduce the burden on developers. This can be achieved automatically by an architecture, meaning the developer is less likely to make mistakes (e.g., not releasing a resource lock causing another process to wait indefinitely for access). Are there novel solutions in architecture that could embody data minimization, collection and use limitation, or individual participation (i.e., the FIPPs)? It should be possible to encapsulate privacy analysis – the privacy properties of a system – in an architecture while hiding details as well.

With respect to existing architectural paradigms, there can be sticky privacy policies, in which privacy requirements "follow" the data (not to be confused with storing the policy with the data, which makes changing policies very difficult and something to avoid). This is technically difficult when data can be transformed and described in different ways, and when different companies use different standards. However, the goal is worthwhile and well suited to further research. There appears to be too little research in software architecture on enforcing privacy requirements.