October 15, 2014

Stuart Shapiro
Julie Snyder

MITRE Comments on the Draft Privacy Engineering Objectives and Risk Model

Approved for Public Release; Distribution Unlimited. Case Number 14-3585

This page intentionally left blank.

Table of Contents

General Comments

The work NIST is doing in privacy engineering represents a critical paradigm shift in how privacy risks are managed. Every day, systems and their supporting technologies evolve. And seemingly every day, there are news stories published about privacy breaches or how these technologies are potentially eroding privacy and often stories about both. Law and policy are based on the long recognized Fair Information Practice Principles (FIPPs), but privacy risks remain and privacy breaches continue to rise. Why? Because these things alone do not *proactively* address privacy risks at the appropriate level of specificity for a given system. To be effective, systems containing personally identifiable information (PII) must be capable of preventing or minimizing the effect of human error or fallibility and appropriately constraining system actions.

To adequately address privacy risks, systems that manage PII must behave in a privacy-sensitive manner. Systems engineering processes are a largely untapped opportunity to embed privacy requirements into organizational activities in a way that provides major impact and will proactively address privacy risks. For all of the relevant stakeholders to participate in the process in a way that is effective, a common vocabulary to discuss privacy engineering is necessary. We applaud NIST's progress in this area and agree that privacy engineering objectives and a system privacy risk model are the right place to start to construct a holistic discipline of privacy engineering with privacy-specific methods, including ethics. This paper provides our comments on the Draft Privacy Engineering Objectives and Risk Model published for discussion at the September 2014 Privacy Engineering Workshop. Our comments are based on MITRE's experience addressing privacy engineering challenges and our participation in the April and September workshops.

Comments on the Draft

Purpose and Scope

Reference: Page 2

The defined scope of privacy engineering refers to "normal system behavior". "Normal" is a loaded term. By virtue of developing a risk model, organizations should be anticipating consequences of system actions and trying to change what is "normal" today into a new "normal" that is better for privacy. Alternatively, Charles Perrow's Normal Accident Theory (NAT)[1] may provide a useful perspective. The more complex and tightly coupled the system, the more likely "accidents" are to occur; in this sense they are "normal" even if they cannot be anticipated in advance. This notion of "normal" seems consistent with the intent of the definition and NIST should consider how Perrow's framing might be integrated into its approach.

The fact that a privacy attack is malicious does not mean it is automatically transformed into a security issue rather than a privacy issue, though it could be both. Consider, for

[1] Perrow, Charles. *Normal Accidents: Living with High-Risk Technologies* New York: Basic Books, 1984. (Revised edition, 1999).

example, dating website OkCupid's experimentation on its users, including putting the "wrong" people together to see if they would connect.[2] This was arguably a malicious privacy violation, in that it was intentional. A proper human subjects research review process might have resulted in a mitigating protocol or even a decision not to proceed, but security principles were not relevant to this particular issue.

We recommend the following edits as a potential way to address some of these comments, as shown in yellow highlight and ~~strikethrough~~:

> ***Scope:*** *The privacy engineering objectives and risk model are primarily focused on mitigating risks arising from ~~unanticipated consequences of normal system behavior~~ the behaviors of systems and their users.*

Privacy Engineering Definition

Reference: Page 4

The definition of privacy engineering is too narrow and the itemization of possible risks by definition excludes other unforeseen risks. Further, while security is relevant to privacy generally and can be referenced in other parts of the model, the definition and objectives should be broader than any one FIPP. An objective of proportionality—that the benefits of the system be proportional to the privacy risks created by the system—would provide a basis for evaluating the purpose of the system and the propriety of its impact on privacy.

We recommend the following edits as a potential way to address these comments, as shown in yellow highlight and ~~strikethrough~~:

> *Privacy engineering is a collection of methods to support the mitigation of risks to individuals ~~of loss of self-determination, loss of trust, discrimination and economic loss~~ by assuring ~~providing~~ predictability, manageability, and ~~confidentiality~~ proportionality with respect to the handling of personal information within information systems.*

Privacy Engineering Objectives

Reference: Page 8

The image effectively shows how privacy engineering objectives and risk analysis fit into the big picture of system design. Consistent with other images that depict similar notions, it would help the non-technical consumers of the model to show that there is a direct correlation between Requirements and Evaluation Criteria. See below for an example.

[2] http://www.bbc.com/news/technology-28542642

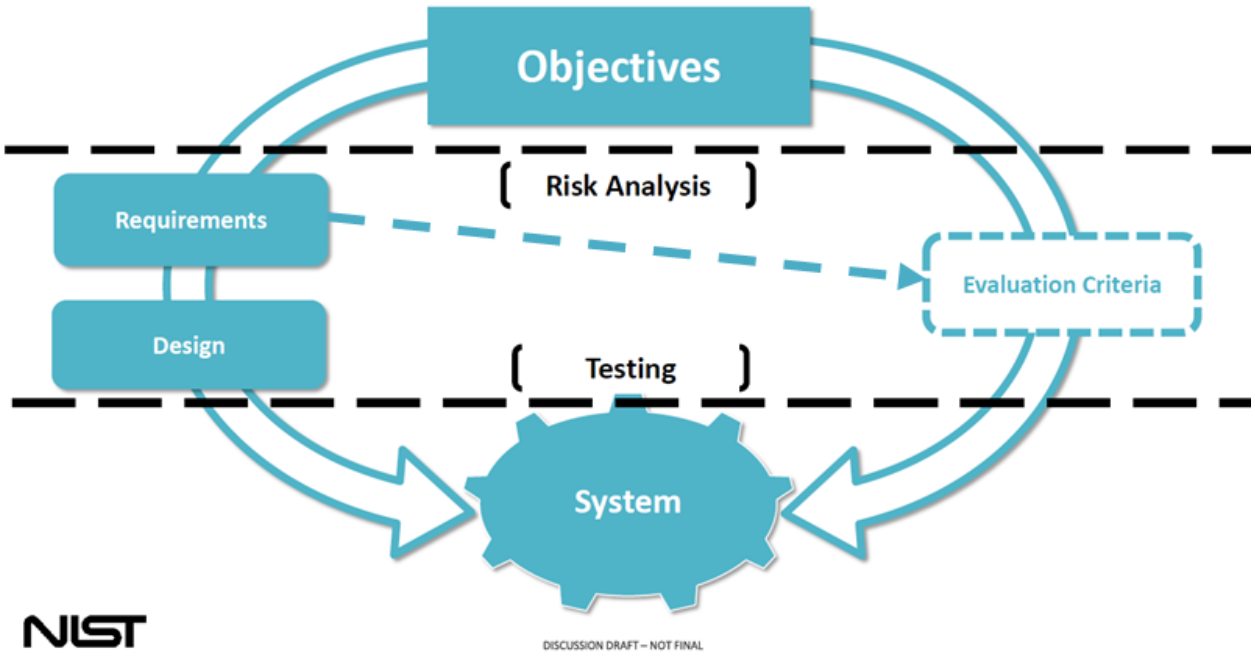

Figure 1. Recommended Edits to Privacy Engineering Graphic

Predictability

Reference: Page 9

The reference to "assumptions" in the definition raises the question of whose assumptions: The individual's? The data controller's? Society's? This should be specified, as well as how differing assumptions might be reconciled.

Manageability

Reference: Page 10

As defined, manageability seems too narrow and too security-like to achieve what is intended. We recommend replacing it with "Data Stewardship" or some similar concept that embodies the following notions:

- Ethics
- Confidentiality
- Data quality
- Maintaining clarity of intent
- Awareness of environment realities
- Meeting the values articulated by the FIPPs

Confidentiality

Reference: Page 11

As indicated above, we believe that confidentiality, at least in a security sense, should be acknowledged but not serve as an objective in its own right within this framework.

Additional Thoughts on Objectives

In addition to protecting individuals from privacy breaches and other risks, one of the key benefits of a privacy engineering risk model is that it will improve the quality of new applications, technologies, and programs. Additionally, building privacy into these things requires a new level of creativity in the areas of requirements, architecture, design, and verification.

System Privacy Risk Model

Assessing System Privacy Risk

Reference: Page 13

The first bullet on this page states:

> *To understand the magnitude of privacy risk within an information system, the proposed model focuses on the risk or likelihood of problematic data actions occurring that could result in privacy harm to individuals.*

This statement suggests the risk is data action risk rather than risk of harm, which raises an interesting question, namely what's the real adverse event of concern? Fault analysis explicitly allows for chaining, but the focus is always on the end event, not an intermediate one. Alternatively, this could be framed in terms of first and second order risks. The "could result" glosses over the existence of a fault/risk chain that runs from the system's data actions to the harms affecting individuals; this should be directly addressed.

Although this privacy risk model is focused on engineering, it needs to acknowledge the business process inputs to systems and the outputs from systems back into the process. If the latter is ignored, problematic data actions and resulting privacy risks may be missed. One of the tenets of the Privacy by Design movement is that privacy assurance should be a default mode of operation. To be effective, privacy controls must not only be considered in the context of a specific system, but also in the context of all system activities that influence how systems are built and operate. Figure 2 below represents the functional layers of system activities.

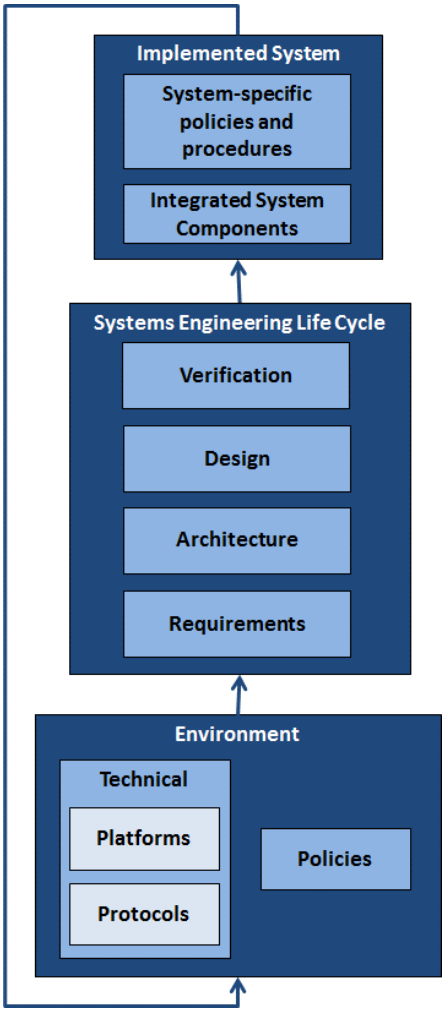

Figure 2. System Activities Stack

Privacy must be addressed in each layer of engineering and operational activities. This ensures privacy is an organic part of the system development and deployment processes that is consistently addressed throughout all activities for all systems, rather than an unnatural "bolt on" activity singularly addressed by system owners.

System Privacy Risk Equation

Reference: Page 14

The proposed System Privacy Risk Equation discards much of the standard risk vocabulary that includes threat, vulnerability, and impact (or consequences). Many other domains successfully employ that triad when talking about risk. Privacy should not dispense with it absent some compelling reason to do so. As discussed in the model, problematic data actions seem to constitute privacy vulnerabilities. Threat modeling of both malicious and non-malicious threats would help identify potential attack vectors that could exploit a vulnerability, thereby establishing a risk.

As proposed, the equation does not depict how the risk value is calculated (additive?, multiplicative?, some other operation?). Does "Data Actions Performed on that Information" refer to problematic actions, non-problematic actions, or any actions? Should

the reader assume that context includes likelihood and impact severity? If the risk triad will not be used, it will help to show how this relates to it and discuss why we need a different and presumably better articulation.

Context

Reference: Page 16

Many of the bullets describing context implicitly assume that the user is the data subject, which is often not the case.

Data Actions

Reference: Pages 17-24

Thinking in terms of "data actions" is a useful concept. The model must account for the fact that some of the data actions that are considered problematic in most circumstances may in fact be essential to the intended purpose of certain systems, which may be appropriate in limited circumstances.

Related Observations from the September 2014 Workshop

Defining Privacy Risk

One of the key challenges for this model is the fact that the default privacy risk model is a FIPPs-based risk model. The FIPPs certainly have their place in articulating our privacy values, but if they were all we needed to get privacy right, there would not be such a desperate need for the important work NIST, MITRE, and others are doing in the area of privacy engineering. In our observations, how the typical privacy professional talks about privacy risk differs greatly from true risk management as a discipline. More often than not, privacy risk is discussed in unstructured terms. It may help NIST to discuss how privacy risk fits in to the enterprise risk management picture, including the following types of risk: compliance, financial, schedule, and reputational. Setting the stage this way may make for a more natural segue into harms and other risk-related concepts. Providing illustrative examples of how a FIPPs-only model can fail may also help. Further, depicting how the privacy challenges faced today can slow and even stifle innovation may resonate well with those that will ultimately be responsible for understanding and using the privacy engineering risk model.

Fear of Regulation

As with the Cybersecurity Framework, we observed that some organizations have an intense fear of additional regulation by the Federal Trade Commission. The concern over excessive regulation is legitimate, but it is only one piece of a larger, complex puzzle to consider. While new standards, including de facto standards, may bring compliance burdens and associated costs in the short term, over time those burdens and costs may pale in comparison to the resource drains of continued real and perceived privacy incidents. Privacy engineering, by minimizing these, in the long term will enable resource expenditures on innovation that exhibits both creativity and privacy sensitivity. This is significant when considering some of the recent research, such as the 2014 TRUSTe Privacy

Index survey[3] that found 89% of Internet users say they avoid companies that do not protect their privacy.

Defining Harms

There was a strong reaction to the definitions of harm, partially due to regulatory concerns and partially due to the negative connotations and confusion around what the terms mean. During the first privacy engineering workshop in April 2014, we observed that even seasoned privacy professionals were not familiar with academic works on harm. This could be a reason there was such a notable reaction to the terminology. Many organizations approach privacy from a compliance perspective, which means their activities and risk model (whether formal or informal) are likely based on the FIPPs. Thinking in terms of harms as defined by NIST in the draft privacy engineering risk model simply does not feel natural yet and they may not have had cause to discover the body of work discussing privacy harms and risk.

At the workshop, some suggested the term "outcomes" rather than harms. Another alternative would be "effects", which could prove useful when thinking in terms of cause-effect relationships regarding system characteristics and functionality and the impact they have on privacy. Indeed, "impacts" could also replace references to harms. The former is strongly rooted in risk analysis while the latter is strongly rooted in engineering; either would strengthen connections with risk analysis and engineering. No matter what they are called, though, defining these types of concerns are a critical component of a risk discussion.

Conclusion

There are significant hurdles ahead, as with charting new territory in any field. We encourage NIST to continue striving toward material progress in the area of privacy engineering. The positive impacts to our society and our economy will be significant over the long term.

About MITRE

The MITRE Corporation is a not for profit company that runs Federally-Funded Research and Development Centers (FFRDCs), including one supporting the National Cybersecurity Center of Excellence (NCCoE), for the U.S. government. MITRE's FFRDCs serve agencies in a variety of areas that impact the public in direct and indirect ways, such as national security; aviation safety and administration; tax administration; homeland security; healthcare; benefits services; cybersecurity; and other missions. We are pleased to respond to NIST's Privacy Engineering Objectives and Risk Model based on our broad perspective gained from serving a variety of government missions, and from the unique perspective of a systems engineering company that combines a strong research base with an informed awareness of the larger policy and contexts in which government operations are conducted.

[3] http://www.truste.com/us-consumer-confidence-index-2014/