CHAPTER 4

Developing Privacy Policies

If at first the idea is not absurd then there is no hope for it.

—Albert Einstein

Don't skip this chapter because the information presented seems obvious or is something you might feel you want to pass off to your legal team. The search for solid engineering requirements starts with solid policy. By policy, we mean the rules that govern, not the Privacy Policy we associate with the web site that is never read.

This is not a chapter about traditional policy creation. The Privacy Policy is the "silk road" (in the classic sense of the ancient Asian Silk Road, not the contemporary online black market web site). It leads the organization to this new world of innovation and privacy engineering. It brings multidisciplinary actors and actions together and combines the best of legal, technical, and process-oriented teams for fair and legitimate processing of personal information (or privacy). This Privacy Policy becomes the basic map or blueprint for the build out. It ultimately should be viewed as the "meta" set of use-case requirements.

This chapter covers the development of policies that will be used as the basis for development of the controls and measures to protect personal information (i.e., privacy standards, guidelines, business rules, and mechanisms). When we discuss policy creation in this context, we are talking about starting with business requirements (a task or series of tasks needed to serve a goal) and functionality goals. Once defined for goals and basic functions, we add requirements driven by applicable law. We then fit and bend our requirements to view the policies we must create through a lens of *functionality* (i.e., each action taken or demanded may be viewed as a requirement specification that must be included in a *system*). That system may be an enterprise, a subunit, end-to-end processing cycle, application, an element of functionality, a person-managed governance activity, among others. There is no exclusive list of what constitutes a system.

Every discussion in this chapter must be considered in this operational, requirement-driven context otherwise it will be easy to slip into traditional "policy" mode. This is *not* a discussion chief privacy officers (CPOs; or whomever is leading the privacy function) will have with every privacy engineer; however, every CPO *must* consider the output of his or her labor in terms of the concrete and measurable requirements and the outcomes discussed here.

Following chapters will show Unified Modeling Language (UML) and systems creation techniques for metadata as a methodology for taking the requirements derived from privacy policies and other technical sources and creating solutions that reflect those requirements. Where neither systems nor features nor privacy enhancing technologies can meet the requirements set forth, governance, training, and leadership "systems" involving the human players in the privacy engineering drama are discussed.

Elements of Privacy Engineering Development

Privacy engineering is the discipline of developing privacy solutions that consist of procedures, standards, guidelines, and mechanisms. Part 2 covers the process of developing privacy solutions, as depicted in Figure 4-1.

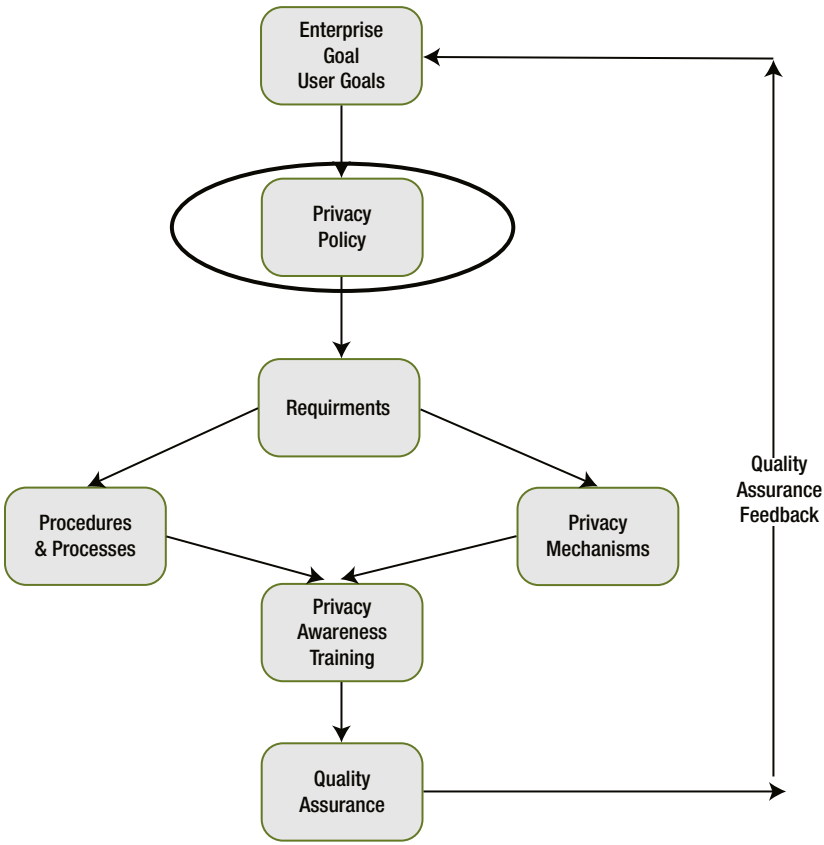

Figure 4-1. *Privacy engineering development process*

The elements of the process of developing a privacy solution, based on a set of privacy policies, are:

- *Enterprise goals*: They must be reflected and aligned with privacy engineering solutions, including their privacy policies, standards, and guidelines. To make this happen, a privacy development team[1] must first understand the goals and objectives of the enterprise in which the solution will operate. For the purposes of this book, "enterprise" includes organizations large and small that manage or otherwise process data. This definition would, of course, include government entities that may be governed by specific or additional rules and regulations and the organizing principles will still apply.
- *User/individual goals*: These must be incorporated to develop effective and flexible privacy policies that will be accepted by the end user and individuals. The team members must understand the goals and objectives (and privacy sensibilities) of the end users and individuals who will participate in the system or become the data subjects for PI managed by the system.
- *Privacy policy*: Development of a privacy policy is discussed in Chapter 4. The policy plays a key role in guiding how privacy engineering is applied.
- *Privacy requirements*: Requirement gathering is critical for effective policy creation and solution development. Chapter 5 describes the application of use cases for requirement collection and introduces a unique use-case metadata model.
- *Privacy procedures and processes*: These are the overall privacy activities (procedures) and their human or automated tasks (processes). Chapters 5 and 6 cover developing and using these as part of the privacy engineering discipline. Mandated standards and recommended guidelines factor into the creation of procedures and processes. It is procedures, processes, standards, and guidelines that translate "policy" into reality.
- *Privacy mechanisms*: These are the automated solutions built with software and hardware to enforce privacy policies. Examples are created for illustration in Chapters 7, 8, and 9 using the development process presented in Chapter 6, including a privacy engineering component and how it can fit within an application system environment.

[1]This team will consist of members from a formal privacy function, business-oriented data stewards, privacy engineers, security analysts, and IT data analysts. Data governance was discussed in Chapter 2. Organizational aspects of privacy engineering will be addressed in Chapter 11.

- *Privacy awareness and readiness preparation*: As part of developing a privacy engineered solution, the team will engage with various stakeholders so they are aware of what the Privacy Policy is and what it does. The privacy team works together with these stakeholders to address how the privacy-engineered solution could affect their roles and responsibilities. This subject is addressed in Chapter 10.
- *Quality assurance*: This is required to ensure that the privacy engineering solution functions properly, as well as satisfies enterprise goals, user goals, and accepted privacy standards within the context they are to operate. Quality assurance for privacy solutions is discussed in Chapter 10.
- *Feedback loop*: This will ensure that the privacy engineering solution is improved continuously as it will periodically quality assess or audit the solution and build in the ability to do so as a technical and procedural requirement.

After reading Part 2, whether you are a privacy professional or an engineer without a privacy background, you should have an understanding of how privacy is engineered into systems.

Privacy Policy Development

Balanced with the enterprise requirements (where the data value of the solution should always exceed its risks when used in context), individual or "user" goals must be considered as part of the final articulation of the "enterprise" goals. The mission, goals, and objectives of the enterprise must be recognized, understood, and analyzed to determine a privacy-engineered solution's requirements. From these, the privacy policies that will govern the privacy engineering solution can be determined. The privacy policy development should be done at two levels: a general level, relevant to all parts of the enterprise, and at an enterprise-specific level, which will often be more specific and detailed than an "enterprise-wide" policy.

Although drafting privacy policies can be the subject of entire legal or organizational tomes, this chapter will go into enough depth so that the principles that comprise privacy policies are sufficiently understandable as the foundational layer of privacy engineering and use-case requirements. These policies enable the management of the principles *as a framework*, which in turn can also lead to:

- The development and deployment of privacy engineered systems
- The exciting missing beast—the framework to build and innovate the privacy engineered data-centric networks, tools, and solutions of the future

What Is a Good Policy?

A policy is considered good based on the manner in which it functions as well as its contextual fit (i.e., how well it balances the needs and objectives of the enterprise with the objectives of the users or customers or employees whose data ultimately flows through that organization). A good policy:

- Arises from well-articulated enterprise goals, which are based on a clear statement of belief or purpose
- Describes what is wanted or intended by the various parties of interest impacted by the enterprise
- Explains why these things are wanted
- Provides positive direction for enterprise employees and contractors
- Provides transparency to the users of systems or individuals interacting with the enterprise
- Is flexible enough so there can be adjustments to changing conditions without changing the basic policy itself
- Is evaluated regularly
- Can be readily understood by all

Policy statements should be written in clear, concise language. A privacy policy should contain everyday words and short sentences and avoid the use of acronyms. If actions are compulsory, "must" should be used. If actions are recommended, "should" should be used. The policy must be practical and easy to implement.

Designing a Privacy Policy

Some organizations begin taking action on mitigating business risks before an official Privacy Policy is published, but defining the policy should be a high priority. Sadly, many enterprises copy policies they find on other companies' web sites and post what amounts to an ad hoc policy of their own before any due diligence has been exercised with regard to knowing their personnel's, process's, or technology's requirements. It's a sad fact, but a vast majority of enterprises own what we call "complianceware"—stuff that they purchase, license, or otherwise "acquire" just in case there is a data breach or a regulatory inquiry at a later date but that they never actually completely deploy.

An example of this is where an enterprise purchases an identity management suite of products and sets the roles to "employee" or "nonemployee" without regard to a good policy that would illuminate why individuals required access to process data or how the roles or employees themselves should be protected and governed. A good privacy policy should be linked closely to this type of deployment. It will set its requirements before deployment or, better yet, before purchase or development if the identity solution is homegrown.

The next section describes the key considerations for crafting an effective privacy policy as well as how to maintain it.

What Should Be Included in a Privacy Policy?

Policies must be designed to meet a complex set of competing needs:

- Local and international legal, jurisdictional, and regulatory necessities, depending on the scope of the enterprise
- Organization or business requirements
- Permission for the marketing–customer relationship for management or business intelligence
- Brand identity
- Industry standards
- Usability, access, and availability for end users of information systems
- Economic pressure to create value through efficient sharing or relationship building
- Enforceability and compliance
- Ethical obligations
- Realistic technology capabilities and limitations

Everything with a digital heartbeat is connected through dynamically formed relationships governed by privacy, security, and trust policies. This means there may be multiple interactive or cascading privacy policies based on the role of the various parties of interest:

- Customers
- Employees or contractors
- Third parties impacted by the enterprise
- Intellectual property owners
- Data types

Each privacy policy should start with the data type and its anticipated lifecycle and be aligned with the enterprise brand and the enterprise standards of conduct. The policy should add value by managing data:

- Respecting and managing regulatory and industrial standards compliance
- Using personal information and confidential data related to it safely and ethically
- Reconciling differences and leveraging synergies between overlapping or competing enterprise policies and goals for other areas, such as audit or litigation data preservation, records management, and physical and IT security

- Establishing a basis for objective respect and trust between an enterprise and its customers, employees, and other impacted groups

As discussed in Chapters 2 and 3, there are several sets of external standards and guidelines defining privacy requirements, including the OECD guidelines for the protection of privacy and transborder flows of personal data, GAPP, PbD, sectorial and competition laws in the United States, APEC privacy accountability frameworks, and the European Union (EU) Data Protection Directive (and member-states implementation of its requirements).[2] These external guidelines and principles can provide a framework for ensuring that the Privacy Policy will offer compliance within the related jurisdictional area.

It should, of course, be noted in the privacy requirements that:

- Not all laws are granular enough to provide one objective interpretation that must be instantiated
- All rules and regulations can always be harmonized to be free of directly conflicting standards and so-called best practices

What is possible is an objective working framework that will become the policy for the enterprise and, ultimately, the basis for process and technology policies, as described in the sidebar.

INTERNATIONALIZATION: DEVELOPING A GLOBAL PRIVACY POLICY

By Dr. Mark Watts, Head of Information Technology Law, Bristows

Europe is not a country. It isn't. And while this will be blindingly obvious to most people reading this book, it's surprising how often I hear it assumed that Europe is essentially a country, with a single, homogenous data privacy law that sets out the rules applicable across the entire region (50 or so countries). If only life were that simple. If only European privacy rules were *that* simple. Sadly, they're not. And the point here is not to ridicule anyone's understanding of European geography or laws, but rather to make the point that, although when working "internationally" in privacy we all make assumptions—we have to, to rationalize the almost overwhelming legal complexity involved—making the wrong assumptions can quickly cause a project to go astray.

[2] OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data are available at www.oecd.org/document/18/0,3343,en_2649_201185_1815186_1_1_1_1,00.html. A downloadable version of the Generally Accepted Privacy Principles (GAPP), along with additional information about the development and additional privacy resources, can be found at www.aicpa.org/privacy. Information about the European Union's Directive on Data Protection is available at http://ec.europa.eu/justice_home/fsj/privacy/index_en.htm.

Perhaps the most common working assumption I see crossing my desk is that the data privacy laws of a particular country are either (i) completely and utterly different from those that apply at "home" (usually the country of the parent company) so none of our existing data privacy policy can possibly apply, or (ii) absolutely identical to those that apply at home and so we don't need any special consideration or handling in the privacy policy; in other words, the international privacy policy can simply be the same as a domestic one. Unfortunately, most of the time, neither "working" assumption works particularly well. A sensible, well-drafted data privacy policy written to meet, say, North American legal requirements will contain much of relevance and application to Europe and beyond because good information handling practices, such as transparency, data quality, and security, are just that—*good* practices that should transcend country borders. But equally, to assume that that's all there is to it and that, say, North American laws can be exported globally would be complacent and would be to ignore significant cultural differences and priorities, not to mention historical sensitivities. Many an international company has come unstuck making this assumption.

For example, assuming the laws that relate to monitoring employee communications in, say, Finland are the same and so just as permissive as those in the United States (an assumption we see a lot) could easily land a company in hot water. Equally, for a European-headquartered company to assume that there are no security breach notification laws in the United States simply because there are so few at home in Europe at the moment can be just as problematic. A privacy policy built on shaky, overly broad assumptions can put a company, even a company that is trying very hard to do the right thing, in breach of applicable law, despite it following its privacy policy to the letter. Perhaps more worryingly, sometimes a breach can occur precisely *because* a company followed a privacy policy—admittedly, a poor privacy policy—to the letter.

Shaky assumptions can lead to another, more subtle but equally problematic risk—the risk of unnecessary overcompliance. Now, this isn't to suggest that companies should develop policies requiring only the minimum amount of compliance required by local law (essentially as little as the company can get away with) but would a company really want to apply the highest common denominator—the strictest standard anywhere—to all of its operations worldwide? Surely not. For example, would it really be wise to export the highly restrictive Finnish laws on monitoring employee communications to every country where a company does business? Most unlikely, because although this approach would ensure compliance with the communication monitoring laws of almost all other countries where the company has employees, it could seriously hamper its business operations in countries with more permissive regimes. This isn't a risk of noncompliance; it isn't a risk of breach. It's a risk of overcompliance that can fetter existing business processes, potentially inhibit sales, and, just as importantly for the privacy professional and

privacy engineer, can damage their internal credibility within the company. All in all, overcompliance can be as much of a problem for the company as undercompliance.

The problem here is not that broad "international" assumptions are being made. They have to be. A global company with operations subject to the data privacy laws of hundreds of different countries cannot realistically be expected to identify every last detailed requirement of every last applicable law because, at least from a regulatory point of view, the world is still a very big place. So developing an international privacy policy (including all procedures, consent statements, contracts, and other supporting documents that go with it) has to involve making certain assumptions. It's just that they have to be the *right* assumptions. You have to know when it's safe to assume (or indeed, force) conformity between countries at a privacy policy level and when to leave enough room to accommodate important local differences in countries' laws.

Where does one start? As good a place as any for most companies is to think carefully about what it actually wants its international privacy policy to do. Is it meant to be some all singing, all dancing document that seeks to set out the various compliance requirements for each of the countries where the company does business? Or is it intended to be something with less lofty ambitions, merely a common set of requirements that will improve compliance everywhere while accepting that in certain countries there will be a "delta" between the requirements of the policy and those of applicable law?

Well-advised companies adopt the second approach, prioritizing the simplicity of a common, global policy that leads to a "good" (and hopefully even "very good") level of compliance everywhere over the more comprehensive and unwieldy, not to mention expensive, approach directed at full compliance everywhere, at least on paper and most likely only on paper. By adopting the second approach, companies are recognizing that there will inevitably be some specific (but hopefully minor) country legal requirements that are not covered by the policy in detail and which may not be complied with to the letter and only in spirit. In an attempt to plug the most significant of any known "gaps" like this, companies often develop country-specific annexes or sections in their privacy policy. An example of this would be a section specific to data collected in Switzerland that extends the privacy policy's requirements to information about legal entities (e.g., companies) as well as individuals (i.e., human beings). To include such an onerous requirement in the main body of the data privacy policy would be to export the Swiss requirement globally unnecessarily, requiring all companies to apply the policy in full to information about legal entities even though it is not legally required where they operate. Including the obligation in an additional annex to the policy and restricting it to data collected in Switzerland enables compliance with the local requirement while limiting its impact geographically.

> But—tweaking the facts slightly—what if the parent company developing the privacy policy is, say, a Swiss bank? In this case it may be desirable or even essential to require its global operations to handle data about legal entities as if they were all subject to Swiss data privacy law. This would suggest that the "Swiss" provision should be included in the body of the privacy policy rather than being buried in an annex limited to data collected in Switzerland.
>
> And this is how international privacy works; there are few if any invariably true assumptions that can be built into any global privacy policy. They always have to be considered and reconsidered on the particular facts for the company developing the policy. Done well, the result can be a robust privacy policy with a good degree of conformity from country to country, capable of generating clear technical requirements that give the privacy engineers a chance of coding "privacy." Done poorly, the result can be a policy that's unnecessarily strict, or with too many exceptions, or which is simply too vague to be useful, any one of which can require last minute changes to the Privacy Policy (and consequently any technical requirements based on it), something which, in my experience, coders really don't seem to like.

General-Level Privacy Policy Development

One of the first things to be determined when drawing up privacy policies is which geopolitical regions or jurisdictions impact the enterprise. Privacy policies for a global enterprise, for example, can start the foundational development process by basing a strategy on the OECD Guidelines and GAPP. In some cases, other localized articulations of fair information processing may be the foundational basis for policy creation. For whatever framework is chosen, the policy creators will need to be able to translate how the various principles are managed if the policy is going to be an effective tool for process and privacy-enhanced systems and features in a privacy engineering context.

For example, a policy statement might require that data be collected relevant to services provided by the current enterprise. The general policy would require a well-defined privacy notice to provide for transparency between the collector of data and the data subject as well as to build an enforceable governance structure where the data asset is known as it enters and moves through its predicted lifecycle. An enterprise must be able to articulate and document how much personal information would be collected for specific purposes according to proportionality principle.

A policy statement should cover proportionality requirements: the benefit derived from the processing of the data should be proportional to its impact to privacy of the individual whose data is being processed. To achieve data proportionality at the time of collection, the data subject's perspective needs must be balanced within the enterprise's objectives.

The Privacy Policy should require a storage and archiving strategy. Encryption, obfuscation, or other security tactical requirements should be covered in the Privacy Policy and have associated standards and guidelines for operational implementation.

Allowances for revisions and exceptions should be included in privacy policies to address the fact that policy needs will change. There are occasions when a customer's, employee's, supplier's, or other party of interest's feedback or requirements may lead to the need to modify privacy policies or grant exceptions.

When an enterprise operates internationally, privacy policies should address the transfer of data among various jurisdictions. The underlying strategies should be people-process and technology oriented and include governance mechanisms that must be designed and executed to follow the data wherever they travel.

This is the point at which many initiatives often fail due to the lack of coordination and integration of effort. The lawyers head off to draft elaborate legal documents neatly tucked away behind a small link that says "Privacy Notice" at the bottom of a web page or buried in the terms and conditions statement of an application. The technical teams can rush off to buy products that obscure or encrypt enough data to satisfy the annual return of the audit team and so on among the teams. An institutional anthropologist could build an entire career analyzing the fascinating and often divergent goals of these now forever-parted teams. Anthropologic observations aside, the course of behavior that should be charted is an ongoing dialogue between the key stakeholders so that a privacy policy (i.e., requirements for processing personal information) can evolve and continue to meet the needs of individuals and the organization and keep pace to aid and not hinder innovation.

Enterprise-Specific Privacy Development

The nature and culture of an enterprise business impacts privacy policies and the creation process. For instance, in the United States, the legal approach is often sectorial governed. An example of this is health care in the United States, where the Health Insurance Portability and Accountability Act of 1996 (HIPAA) policies and privacy rules should be incorporated. This type of enterprise will always be extremely open with many third parties, operating in a nonstop high-stakes context (in some cases, life and death). Getting the balance between use, sharing, access, and accuracy will be a supreme consideration. The rights and sensitivities of the data subjects within this context are highly subjective while also the subject of extensive regulation. Although other jurisdictions may not have standalone health data protection statutes, this type of context, and health data specifically, is governed as a protected class—or even an enhanced protected class, as in the European Union, a "sensitive" data class of data worldwide.

A health care-, financial-, or politically sensitive type of context is actually the proving grounds for many other types of businesses. These enterprises require personalization and intimate knowledge of personal information, but also value a certain level of autonomous innovation with data and financial models based on data. Innovating for high-risk data is a bit like the lyrics from the song "New York, New York": "If I can make it there, I'll make it anywhere."

A similar illustration can be drawn for financial data in the United States where the Gramm-Leach-Bliley Act requires financial institutions—companies that offer consumers financial products or services like loans, financial or investment advice, or insurance—to explain their information-sharing practices to their customers and to safeguard sensitive data. These types of data are covered by other comprehensive global laws such as the Personal Information Protection and Electronic Documents Act (PIPEDA) in Canada or

under the Argentine Data Protection Laws but may not be called out under a specific law or called out as "sensitive" data calling for enhanced protections beyond the comprehensive requirements. The point here is that although not all data is created equal (nor do they call for exactly the same type of privacy policy treatment), personal information should be considered a controlled substance, and close partnerships and legal considerations are certainly necessary before we innovate on top of the foundational policy.

Internal vs. External Policies

Data protection standards such as the OECD Guidelines and GAPP, among others, require that privacy policies should be published both internally in enterprises and externally (actually, externally, it is usually a statement or notice of an enterprise practices that is posted, not the actual policy) to give notice to users of systems, customers, or other data subjects interacting with the enterprise. Failure to comply with the enterprise public notices can lead to:

- *Dissatisfied customers*: Customers and other users will expect compliance to the privacy protection actions as indicated within the notice. It may be considered an implied contract. If there is a breach, users will tend to look to safer sites. If a user discovers identity theft that seems to have come from personal information collected by an enterprise, that user will take it out on the enterprise maintaining the site that failed them.
- *Regulatory investigations*: Where an enterprise has not lived up to its notice commitments, regulators from one or more jurisdictions will likely investigate the problems and may take either criminal or civil actions or both against both the enterprise and, conceivably, against employees within the enterprise.
- *Bad publicity*: Forty-six US states, the District of Columbia, plus other US territories have security breach notification laws that involve personal information. There are comparable laws throughout the world. The media keep a lookout for such notifications and determine when breaches are significant. Any breach scares people, and serious breaches equal bad publicity.
- *Litigation*: Potential liability in privacy-related lawsuits has been increasing steadily in recent years. This expanding legal exposure has been fueled by plaintiffs' class action lawyers targeting privacy litigation as a growth area. Moreover, federal and state government agencies, as well as data protection agencies throughout Europe and Asia, are becoming increasingly aggressive in their efforts to investigate and respond to privacy and data security concerns and incidents. The Federal Trade Commission (FTC) is imposing stricter standards on businesses, while state attorneys general are pursuing enforcement actions and conducting high-profile investigations in response to data breaches and other perceived privacy violations.

- *Harm to brand*: For most enterprises, the equity invested in their brands is an invaluable but fragile asset. When privacy protection problems occur, the reaction of the enterprise is crucial to the maintenance of a very positive brand.
- *Weak innovation*: Effective innovation comes from making improved products that deliver what people want. To find what customers and potential customers want requires the collection of data. An enterprise that does not protect the privacy of data will weaken the ability to collect the data needed to determine where innovation is required.
- *Employee distrust*: Just as customers can be turned off when privacy notice failures occur, employees can begin to distrust their enterprise when their data is not protected as the privacy notice promise.

An enterprise should consider creating training based on internal privacy rules that are more granular, specific, and more restrictive than externally posted notices. These internal policies should be coordinated with a human resources policy team to ensure that staff and business partners know exactly what to do, how to get help when they need it, and how and when these may be enforced and encouraged.

These policies must all be reflected and are instantiated in product and systems development as discussed further in Chapters 5 and 6.

ENGINEERS AND LAWYERS IN PRIVACY PROTECTION: CAN WE ALL JUST GET ALONG?

By Dr. Annie I. Antón, Professor in and Chair of the School of Interactive Computing at the Georgia Institute of Technology

Peter Swire, Nancy J. and Lawrence P. Huang Professor, Scheller College of Business, Georgia Institute of Technology

In March 2013 we participated in a panel titled "Re-Engineering Privacy Law" at the International Association of Privacy Professionals Privacy Summit. The topic of the panel closely matches the topic of this book, how to bring together and leverage the skill sets of engineers, lawyers, and others to create effective privacy policy with correspondingly compliant implementations. As a software engineering professor (Antón) and a law professor (Swire), we consider four points: (1) how lawyers make simple things complicated; (2) how engineers make simple things complicated; (3) why it may be reasonable to use the term "reasonable" in privacy rules but not in software specifications; and (4) how to achieve consensus when both lawyers and engineers are in the room.

1. *How lawyers make simple things complicated.* A first-year law student takes Torts, the study of accident law. A major question in that course is whether the defendant showed

> "reasonable care." If not, the defendant is likely to be found liable. Sometimes a defendant has violated a statute or a custom, such as a standard safety precaution. More often, the answer in a lawsuit is whether the jury thinks the defendant acted as a "reasonable person." The outcome of the lawsuit is whether the defendant has to pay money or not. We all hope that truth triumphs, but the operational question hinges on who can prove what in court.

The legal style is illustrated by the famous *Palsgraf* case.[3] A man climbs on a train pulling out of the station. The railroad conductor assists the man into the car. In the process, the man drops a package tucked under his arm. It turns out the package contains fireworks, which explode, knocking over some scales at the far end of the platform. The scales topple onto a woman, causing her injury.

From teaching the case, here is the outline of a good law student answer, which would take several pages. The answer would address at least four issues. For each issue, the student would follow IRAC (Issue, Rule, Analysis, Conclusion) form, discussing the issue, the legal rule, the analysis, and the conclusion: (1) Was the man negligent when he climbed on the moving train? (2) When the railroad conductor helped the man up, was the conductor violating a safety statute, thus making his employer, the railroad, liable? (3) When the man dropped the fireworks, was it foreseeable that harm would result? (4) Was the dropping of the package the proximate cause of knocking over the scales? In sum, we seek to determine whether the railroad is liable. The law student would explain why it is a close case; indeed, the actual judges in the case split their decision 4-3.

Engineers design and build things. As such, they seek practical and precise answers. Instead of an IRAC form, engineers seek to apply scientific analytic principles to determine the properties or state of the "system." The mechanisms of failure in the *Palsgraf* case would be analyzed in isolation: (1) The train was moving, therefore, the policy of only allowing boarding while the train is stopped was not properly enforced, thereby introducing significant safety risk into the system. (2) The scales were apparently not properly secured, thus a vibration or simple force would have dislodged the scales, introducing safety risk into the system. Is the railroad liable? An engineer would conclude the compliance violation and unsecured scales means that it would be liable. The engineering professor would congratulate the engineering student for the simple, yet elegant, conclusion based on analysis of isolated components in the system. In engineering, simplicity is the key to elegance.

[3] *Palsgraf v. Long Island Railroad Co*., 248 N.Y. 339 (N.Y. 1928).

The lawyer may agree in theory that simplicity is the key to elegance, but law students and lawyers have strong reasons to go into far more detail. The highest score in a law school exam usually spots the greatest number of issues; it analyzes the one or two key issues, but also creates a research plan for the lawyers litigating the case. For example, the railroad has a safety rule that says the conductor shouldn't help a passenger board when the train is moving, but surely there are exceptions? In the actual case (or the law school exam), the lawyer would likely analyze what those exceptions might be, especially because finding an applicable exception will free the railroad from liability. The good exam answer may also compare the strange chain of events in *Palsgraf* to other leading cases, in order to assess whether the plaintiff can meet her burden for satisfying the difficult-to-define standard for showing proximate cause.

In short, lawyers are trained to take the relatively simple set of facts in *Palsgraf* and write a complex, issue-by-issue analysis of all the considerations that may be relevant to deciding the case. The complexity becomes even greater because the lawyer is not seeking to find the "correct" answer based on scientific principles; instead, the lawyer needs to prepare for the jury or judge, and find ways, if possible, to convince even skeptical decision-makers that the client's position should win.

2. *How engineers make simple things complicated.* A typical compliance task is that our company has to comply with a new privacy rule. For lawyers, this basically means applying the Fair Information Privacy Principles (FIPPs), such as notice, choice, access, security, and accountability. The law is pretty simple.

The engineer response is: How do we specify these rules so that they can be implemented in code? Stage one: specify the basic privacy principles (FIPPs). Stage two: specify commitments expressed in the company privacy notice. Stage three: specify functional and nonfunctional requirements to support business processes, user interactions, data transforms and transfers, security and privacy requirements, as well as corresponding system tests.

As an example, some privacy laws have a data minimization requirement. Giving operational meaning to "data minimization," however, is a challenging engineering task, requiring system-by-system and field-by-field knowledge of which data are or are not needed for the organization's purposes. Stuart Shapiro , Principal Information Privacy & Security Engineer, The MITRE Corporation, notes that an implementation of data minimization in a system may have 50 requirements and 100 associated tests. Input to the system is permitted only for predetermined data elements. When the system queries an external database, they are permitted only to the approved data fields. There must be executable tests—apply to test data first and then confirm that data minimization is achieved under various scenarios.

For the lawyer, it is simple to say "data minimization." For the engineer, those two words are the beginning of a very complex process.

3. *Why it may be reasonable to use the term "reasonable" in privacy rules.* Swire was involved in the drafting of the HIPAA medical privacy rule in 1999–2000. Antón, the engineer, has long chastised Swire for letting the word "reasonable" appear over 30 times in the regulation. Words such as "promptly" and "reasonable" are far too ambiguous for engineers to implement. For example, consider HIPAA §164.530(i)(3): "the covered entity must *promptly* document and implement the revised policy or procedure." Engineers can't test for "promptly." They can, however, test for 24 hours, 1 second, or 5 milliseconds. As for reasonable, the rule requires "*reasonable* and appropriate security measures"; "*reasonable* and appropriate polices and procedures" for documentation; "*reasonable* efforts to limit" collection and use "to the minimum necessary"; a "*reasonable* belief" before releasing records relating to domestic violence; and "*reasonable* steps to cure the breach" by a business associate.

The engineer's critique is: How do you code for "promptly" and "reasonable"? The lawyer's answer is that the HIPAA rule went more than a decade before being updated for the first time, so the rule has to apply to changing circumstances. The rule is supposed to be technology neutral, so drafting detailed technical specs is a bad idea even though that's exactly what engineers are expected to do to develop HIPAA-compliant systems. There are many use cases and business models in a rule that covers almost 20% of the US economy. Over time, the Department of Health and Human Services can issue FAQs and guidance, as needed. If the rule is more specific, then the results will be wrong. In short, lawyers believe there is no better alternative in the privacy rule to saying "reasonable."

The engineer remains frustrated by the term "reasonable," yet accepts that the term is intentionally ambiguous because it is for the courts to decide what is deemed reasonable. If the rule is too ambiguous, however, it will be inconsistently applied and engineers risk legal sanctions on the organization for developing systems not deemed to be HIPAA compliant. In addition, "promptly" is an unintentional ambiguity that was preventable in the crafting of the law. By allowing engineers in the room with the lawyers as they decide the rules that will govern the systems the engineers must develop, we can avoid a lot of headaches down the road.

4. *How to achieve happiness when both lawyers and engineers are in the same room.* Organizations today need to have both lawyers and engineers involved in privacy compliance efforts. An increasing number of laws, regulations, and cases, often coming from numerous states and countries,

> place requirements on companies. Lawyers are needed to interpret these requirements. Engineers are needed to build the systems.

Despite their differences, lawyers and engineers share important similarities. They both are very analytic. They both can drill down and get enormously detailed in order to get the product just right. And, each is glad when the other gets to do *those* details. Most engineers would hate to write a 50-page brief. Most lawyers can't even imagine specifying 50 engineering requirements and running 100 associated tests.

The output of engineering and legal work turns out to be different. Engineers build *things*. They build systems that work. They seek the right answer. Their results are testable. Most of all, it "works" if it runs according to spec. By contrast, lawyers build *arguments*. They use a lot of words; "brief" is a one-word oxymoron. Lawyers are trained in the adversary system, where other lawyers are trying to defeat them in court or get a different legislative or regulatory outcome. For lawyers, it "works" if our lawyers beat their lawyers.

Given these differences, companies and agencies typically need a team. To comply, you need lawyers *and* engineers, and it helps to become aware of how to create answers that count for both the lawyers and the engineers. To strike an optimistic note, in privacy compliance the legal and engineering systems come together. Your own work improves if you become bilingual, if you can understand what counts as an answer for the different professions.

We look forward to trying to find an answer about how to achieve happiness when both lawyers and engineers are in the room. Antón presumably is seeking a testable result. Swire presumably will settle for simply persuading those involved. However, we both agree that the best results come from collaboration because of the value, knowledge, and expertise that both stakeholder groups bring to the table.

Policies, Present, and Future

Policies have to be living documents that can be readily changed as a business changes or as the regulatory environment changes; however, they should not be changed lightly or at whim. There is overhead associated with policy changes, especially in the privacy space. For instance, a change in policy may indicate a change in use of data, which then may require an enterprise to provide notice of the change to whomever's data is affected and get permission for the new uses of the data. Even without a pressing need for change, it is important to review policies on a regular basis, perhaps annually, to determine if change is necessary.

A good policy needs to be forward looking and, at the same time, accurate to the current state. It should be sufficiently detailed as to give direction and set parameters, but not so detailed as to be overly specific or to require excessive change. Each enterprise will need to find the balance between what is communicated as "policy" and what is

communicated as an underlying standard or guideline for meeting the requirements of the policy. Key stakeholders should review policies and practices at least annually to see if revisions are warranted.

Engineered privacy mechanisms can ease the change and improvement of the policies, especially with the specific procedures, standards, guidelines, and privacy rules that need to change if there are policy revisions. The privacy component discussed in Chapters 6, 7, 8, and 9 addresses this crucial need.

Conclusion

Privacy policies are powerful tools in the overall privacy engineering process. Privacy professionals, lawyers, and compliance teams can use them to communicate expected behaviors and leverage them to create accountability measures. In the process of policy creation, internal and external—including systems' users and regulators—requirements and expectations must be gathered. These same requirements and expectations in the traditional lexicon can also be leveraged as engineering requirements in the privacy engineering model and execution sense. We will explore how such requirements fit into a system's model in Chapters 5 and 6. In the remaining chapters of Part 2, we will continue to call on these policy requirements in the context of discrete tools and features that rest in the privacy engineering toolkit.