CHAPTER 3

Data and Privacy Governance Concepts

Computers are magnificent tools for the realization of our dreams, but no machine can replace the human spark of spirit, compassion, love, and understanding.

—Louis Gerstner

This chapter will look at the relationship among privacy frameworks and data management, data governance, and data stewardship, highlighting how frameworks such as the OECD Guidelines and GAPP are used for personal information management. Included in this discussion will be a look at Privacy by Design (PbD), which supports and complements privacy engineering (Figure 3-1).

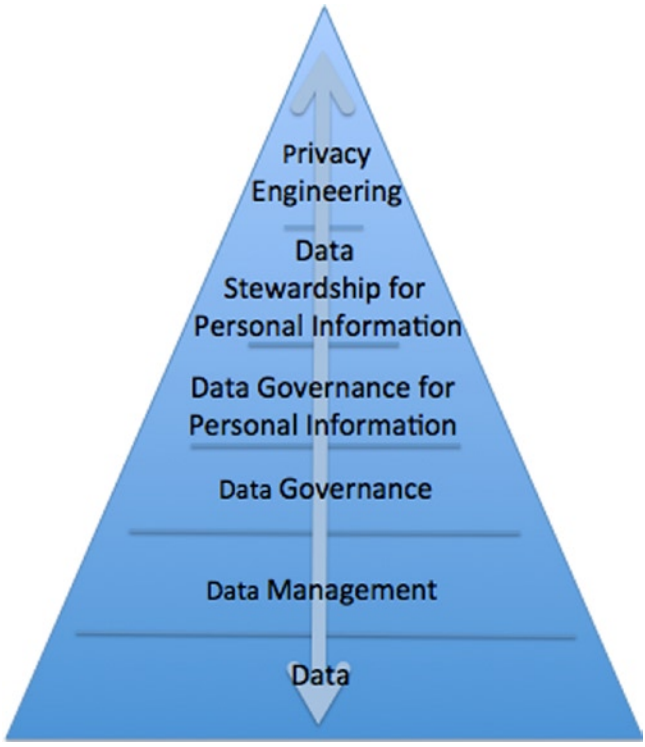

Figure 3-1. *Good privacy engineering is built on a foundation of data management and governance*

Data Management: The Management of "Stuff"

The raison d'etre of any organization, whether a corporation, a nonprofit, or a governmental entity, is to do "stuff;" doing "stuff" requires managing "stuff." Data represents this stuff. Examples include:

- Customers
- Suppliers
- Money
- Resources
- Products
- Customer orders
- Customer order line items
- Inventory

- Policies
- Business rules
- Privacy rules
- Roles
- Intellectual property[1]

The administration of the data that represents the "stuff" of an organization is the science and art of data management, or as it is defined in the DAMA Data Management Body of Knowledge: "Data management is the development, execution, and supervision of plans, policies, programs and practices that control, protect, deliver and enhance the value of data and information assets."[2]

In a structured data management program, data stewards, who are domain or subject matter experts for each of these classes of data, work with data management experts to ensure that procedures, processes, standards, guidelines, and business rules for using such information support the goals and objectives of the enterprise. This is called data governance.

Data Governance

Data governance is a strategic, "top-down" program for data management in which an organization's leadership communicates the core value of data quality and integrity to stakeholders. It includes the development and enforcement of standards and procedures. It requires broad understanding of data entrusted to the organization, the value and use of data, upstream and downstream stakeholders, systems, and processes for all decisions and issue resolution. To be effective, data governance requires data stewardship and data stewards. It also requires executive sponsors and support.

Stewardship is not ownership. A *steward* is a custodian who is responsible for managing something that belongs to someone else. *Data stewardship* is the managing of information on behalf of the "owners" of the data. The data steward is in effect "the feet on the ground," ensuring the data governance standards are adhered to and evolve as necessary.

[1]For any enterprise, we would expect to find over 20 different data models containing at least five unique classes or data entities and the relationships between these classes or data entities. We have built these types of enterprise data models for a number of pharmaceutical companies, communications companies, oil companies, hospitality companies, and government agencies, among others.
[2]"DAMA-DMBOK Guide (Data Management Body of Knowledge) Introduction & Project Status." www.dama.org/files/public/DI_DAMA_DMBOK_Guide_Presentation_2007.pdf.

An effective data governance program requires that:

- Data is created, recorded, and distributed in compliance with standards
- An established metadata gathering process clearly describes requirements and characteristics of the data to be maintained (discussed in Part 2 of this book, and Appendix A contains a variety of metadata)
- There is a metric-driven adherence of all data definition standards
- There is a feedback or notification system to identify inadequacies in the data
- There is a data quality assurance process that monitors the integrity of information within the system
- There is a data management structure that includes data stewardship, a data governance panel, and an executive layer

There are two data steward roles: data producer stewards and data usage stewards. Data producer stewards are responsible for:

- Appropriate data content creation and maintenance of quality.
- Appropriate business rules related to all data elements and attributes for which the data steward has responsibility. A data attribute is a fact or characteristic about a data element or entity.

Data usage stewards are responsible for:

- Appropriate data usage quality, including screens and reports
- Appropriate business rules, including privacy
- Appropriate presentation:
    - Method
    - Design
    - Architecture
    - Aesthetics (ugly user interfaces are avoided)

In addition to the role of the data producer and data usage steward, there is the role of data administrator.

Data administrators are those responsible for:

- Data analysis
- Data acquisition design
- Data organizing or classifying
- Data storage and distribution design

- Data archiving
- Ensuring the implementation of business rules
- Data management (metadata) tool administration (as a data dictionary)

Depending on the size and volume of the data being managed, these roles may be combined or staffed by more than one person.

Benefits of Data Governance

Data management programs that have implemented data governance have benefited from features such as:

- *Common names and definitions:* If existing data is not well named, they cannot be found and therefore cannot be shared.[3] In order to determine whether a data object already exists, common names, based on a standard naming convention, speed the analysis. Common names imply that there is a readily understandable business name and an abbreviated short physical name, based in part on a standard abbreviation list.
- *Consistent data:* A consistent business definition of the data is important so that the knowledge worker can determine whether a data object with a name similar to his or her data requirement is in fact the same data object.
- *Consistent reports:* If data attributes are well named or well defined, then the reports resulting from the analysis or use of the elements are apt to be more consistent because the underlying data is consistent.
- *Less duplication of data:* Consistent names and definitions will facilitate the discovery of redundant data. Data modeling normalization is a process for eliminating duplication.
- *Trust by the business users:* Well-executed data governance and data stewardship should improve quality and reliability, which, in turn, should increase accuracy and trust in the data analysis process.
- *Less data correction:* Better managed data should be more accurate and require less correction.

However, the most important feature and benefit of data governance is that the data is being governed and that there are structured, mindful controls and measures in place to manage the data and ensure that its use is in alignment with the organization's overall goals and requirements. In short, the data is being viewed as an asset and is appropriately and meaningfully curated.

[3]B. Van Halle and C. Fleming, *Handbook of relational database design*, Addison-Wesley, 1989, p. 16.

The Privacy and Data Governance/Stewardship Connection

Although it is not often articulated this way, data privacy is a key part of data governance for personal information. In this context, privacy engineering is engineering data governance for personal information into the design and implementation of routines, systems, and products that process personal information. An enterprise's privacy policy (including rules, standards, guidelines, etc.) "governs" the processing of personal information by an enterprise (and in Chapter 4, the privacy policy is not only viewed as a governance concept but also the meta-set of personal information data protection use-case requirements for privacy engineering).

Understanding how data management frameworks (such as data governance and data stewardship) fit with privacy frameworks (such as GAPP and the OECD Guidelines) is key to organizational development. Such frameworks and guidelines help to create the necessary roles and responsibilities to build and maintain a privacy-aware and ready enterprise. Such understanding will also help to recognize and understand privacy policies at meta-use-case requirements for privacy engineering.

Although the connection between data governance and privacy frameworks should be very close, the closeness is not often recognized nor leveraged by either domain. Too often data privacy teams sit outside enterprise-wide data governance and stewardship initiatives. This is unfortunate. File this under the opportunity not realized category. Ultimately both groups should have a shared goal of ensuring data is curated and cared for as an asset whose value is recognized and cultivated within defined parameters.

Data Privacy Governance Frameworks

The OECD Guidelines, that were discussed in Chapter 2, is one of the better-known privacy governance frameworks. In addition to it, are other global and regional frameworks such as the 1995 EU Data Protection Directive (also known as Directive EU 95/46/EC), the Federal Trade Commission's version of the Fair Information Privacy Principles, (FIPPs), the ISO 2700x series of security standards, and the Generally Accepted Privacy Principles (GAPP), which were created by the American Institute of Certified Public Accountants (AICPA) and the Canadian Institute of Chartered Accountants (CICA) Privacy Task Force.

All these and others are worth knowing and learning about to perfect a privacy engineering tradecraft.

HOW THE FRAMEWORKS ALIGN

You can see from Table 3-1 how the various frameworks cited align. One of the most comprehensive is GAPP, which was designed to create a set of principles that would encompass the key points of the existing frameworks.

Table 3-1. *How Key Privacy Frameworks Align*

GAPP	OECD Guidelines	FTC FIPPS	EU Directive	ISO 27002	APEC
Management				Operations Management	Preventing Harm
Collection	Collection Limitation		Proportionality	Information Acquisition	Collection Limitations
Quality	Data Quality				Integrity of Personal Information
Notice	Specification of Purpose	Notice/Awareness	Transparency		Notice
Use, Retention, Disposal	Use Limitation		Legitimate Purpose	Asset Management	Uses of Personal Information
Security for Privacy	Security Safeguards	Integrity/Security		Security	Security Safeguards
Access	Openness	Access/Participation		Access Control	Access and Correction
Choice/Consent	Individual Participation	Choice/Consent		Asset Management	Choice
Monitoring and Enforcement	Accountability	Enforcement/Redress	Supervisory authority	Compliance	Accountability
Disclosure to Third Parties			Transfer of personal data to third parties		

Generally Accepted Privacy Principles (GAPP)

According to the American Institute of Certified Public Accountants (AICPA), which developed the Generally Accepted Privacy Principles:

> *Generally Accepted Privacy Principles (GAPP) have been developed from a business perspective, referencing some, but by no means all, significant local, national and international privacy regulations. GAPP operationalizes complex privacy requirements into a single privacy objective that is supported by 10 privacy principles. Each principle is supported by objective, measurable criteria that form the basis for effective management of privacy risk and compliance in an organization.*[4]

The following are the 10 GAPP:

1. *Management:* The entity defines, documents, communicates, and assigns accountability for its privacy policies and procedures.
2. *Notice:* The entity provides notice about its privacy policies and procedures and identifies the purposes for which personal information is collected, used, retained, and disclosed.
3. *Choice and consent:* The entity describes the choices available to the individual and obtains implicit or explicit consent with respect to the collection, use, and disclosure of personal information.
4. *Collection:* The entity collects personal information only for the purposes identified in the notice.
5. *Use, retention, and disposal:* The entity limits the use of personal information to the purposes identified in the notice and for which the individual has provided implicit or explicit consent. The entity retains personal information only as long as necessary to fulfill the stated purposes or as required by law or regulation and thereafter appropriately disposes of such information.
6. *Access:* The entity provides individuals with access to their personal information for review and update.
7. *Disclosure to third parties:* The entity discloses personal information to third parties only for the purposes identified in the notice and with the implicit or explicit consent of the individual.

[4]See www.aicpa.org/InterestAreas/InformationTechnology/Resources/Privacy/GenerallyAcceptedPrivacyPrinciples/DownloadableDocuments/10261378ExecOverviewGAPP.pdf

8. *Security for privacy:* The entity protects personal information against unauthorized access (both physical and logical).

9. *Quality:* The entity maintains accurate, complete, and relevant personal information for the purposes identified in the notice.

10. *Monitoring and enforcement:* The entity monitors compliance with its privacy policies and procedures and has procedures to address privacy-related complaints and disputes.

We will show in later chapters how frameworks like the OECD Guidelines and GAPP are used as a basis for developing the enterprise's privacy policies, processes, procedures, standards, guidelines, and mechanisms.

ISO2700X: HOW SECURITY STANDARDS SUPPORT PRIVACY

By Joel Weise, Director of Security and Compliance, Hootsuite

The ISO 27001:2005 "Information technology—Security techniques—Information security management systems—Requirements" and the complementary ISO 27002:2005 "Information technology—Security techniques—Code of practice for information security management" standards provide a very good framework for defining, creating, and managing a comprehensive security architecture and governance framework that supports not only security but also privacy. Some of the primary advantages are that these are mature standards, internationally recognized and well harmonized with other local and national standards such as the US NIST Special Publication 800-53 "Recommended Security Controls for Federal Information Systems and Organizations." Further, when utilized, the standards can enable compliance to privacy laws, demonstrate an organization's commitment to privacy and minimize, or limit the opportunity for breaches that could affect security and privacy of data, people as well as supporting technology and governance.

The overall value of the standards is to elaborate an information security management system (ISMS) as noted in ISO 27001:2005 and based on the security control objectives as noted in ISO 27002:2005. The ISMS uses a continuous improvement approach so that it is flexible and can change as new laws, technology, and threats emerge. The standards further allow for the foundation of a framework that can be audited so that its effectiveness can be measured. Such a foundation is critical to supporting security and privacy efforts in an organization. According to the standards, "The ISMS is designed to ensure the selection of adequate and proportionate security controls that protect information assets and give confidence to interested parties." This goal is fundamental to how the ISMS functions and addresses both security and privacy. The overall benefit of the standards is that they are used to enable the design, configuration, implementation, and use of controls

that reflect best practices, and, most important, it allows for interoperability and a lingua franca so that different organization, security, and privacy professionals as well as auditor and legal authorities can analyze the use of those controls.

When considering security and privacy controls, one must always consider the costs of such controls. It is important that controls be balanced against their actual and intangible costs. For example, it would not be reasonable to implement a $100 control to address a risk that is only worth $10. A security practitioner must always evaluate controls within the business context of the environment in which they will be implemented. In addition to an actual value, one must consider the intangible costs of controls. For example, even if a $100 control is used to address a risk valued at $1,000, the security practitioner must consider intangible costs such as the impact the moral, productivity, and general perception of security. If a control negatively impacts the organization, even in such intangible ways, those should be taken into consideration.

The ISO 27002:2005 standard has 11 different sections. Table 3-2 outlines each of these areas as they apply to privacy.

Table 3-2. *Standards that Apply to Privacy*

Standard Topic Area	Overview	Privacy Objective
Policy	The policy is a high-level statement about information security and privacy. It lays down the key information security and privacy directives for an organization.	The policy should reflect the privacy compliance objectives of the organization and reference applicable standards, legal and regulatory mandates, and relevant industry-best practices.
Organizing Information Security	An information security governance structure should span the entire business and technical components of the organization.	The organizational governance structure should include specific individuals and functions that have privacy as their primary mandate.
Asset Management	Asset management is a means for an organization to identify, organize, and manage their information resources.	The maintenance of privacy for data assets is an organizational imperative because many assets include a privacy component.

(*continued*)

Table 3-2. (*continued*)

Standard Topic Area	Overview	Privacy Objective
Human Resources Security	The organization should manage user access rights as well as undertake suitable security awareness training and educational activities. These are all necessary to ensure the human element actively participates in the overall security effort.	In order to ensure employee personal information is secure, protected, and used appropriately, privacy needs to be instilled in an organization's culture through training and awareness activities.
Physical and Environmental Security	Valuable IT equipment should be physically protected against malicious or accidental damage or loss including damage or loss due to environmental factors such as an inadvertent loss of power or overheating.	Maintaining privacy in an organization's physical space is also important as is security and privacy of data assets.
Communications and Operations Management	Controls for systems and network management include a broad range of capabilities from network management to operational procedures.	In the IT world, privacy can only be enabled when appropriate system and network controls are utilized to ensure the security, availability, and reliability of operational resources.
Access ControlCommunications and Operations Management	Access control includes user access controls for IT systems, including, operating systems, networks, and applications and data.	Access control is critical for the support of privacy in any environment where data and processing resources may contain personal information.

(*continued*)

Table 3-2. (*continued*)

Standard Topic Area	Overview	Privacy Objective
Information Systems Acquisition, Development, and Maintenance	This section details the policies covering everything from cryptography to processes for specifying, building or acquiring, testing, implementing, and maintaining IT systems.	Maintaining the privacy of data is predicated upon implementing and supporting an IT infrastructure that works as advertised. Without that assurance, it is not possible to state that an organization is capable of maintaining the privacy of data.
Information Security Incident Management	Incident management covers procedures required to manage incidents consistently and effectively.	Knowing that intrusions can exacerbate vulnerabilities, maintaining the privacy of data relies upon a comprehensive incident management function. It also alerts you to breaches so you can remedy them as quickly as possible.
Business Continuity Management	This section describes the relationship between IT disaster recovery planning, business continuity management, and contingency planning.	To the extent that personal information is retained in backups, then disaster recovery and business resumption processes must ensure the continued control over those assets.
Compliance	Compliance includes not only compliance with legal requirements, but also with security and privacy policies and standards.	Compliance to relevant security and privacy policies is integral to ensuring privacy as this enables users a means to validate adherence to those policies.

Impact of Frameworks on the Privacy Engineer

Privacy engineers must understand the OECD Guidelines, GAPP, and the other frameworks, as well as their organization's own privacy policies, standards, and guidelines sufficiently to understand their purpose and limitations. In doing so, any creative innovation should have a tie into a rationalized set of existing requirements. This will, in turn, make it easier to implement such an innovation or manage change effectively as a logical leap forward in achieving the ultimate goal of efficiently, effectively, and ethically protecting information about people.

If data is processed in a way that honors or adheres to the OECD Guidelines or GAPP, or one of the other frameworks, then chances are, under most data privacy regimes, it will likely be considered to be fair and legitimate processing as most privacy laws are based on the FIPPs in some fashion (and these other frameworks essentially follow the FIPPs). However, as noted later, each specific case or legal regime can and often does interpret the FIPPS, adherence, and individual level of competency differently.

In Part 2 of this book, we will discuss how privacy rules are developed based on privacy policies, processes, procedures, standards, guidelines, and best practices that are derived in part from these frameworks. These privacy rules will be used to implement mechanisms that are used within systems satisfying privacy requirements.

Frameworks Are Not the Same as Laws

How each enterprise addresses privacy requirements at a deeper more granular level is a decision that is based on many factors such as size, jurisdiction, risk profile, internal policies and public positions, and, most important, what kind of personal information is involved (i.e., how much and how sensitive) and whose data it is.

To get to this level of granularity in understanding requirements, you should work with legal resources with privacy domain expertise and look at the specific laws and regulations that govern the space in which you are working, as well as applicable internal policies and requirements.

For this reason, the techniques for privacy engineering that will be discussed in this book and the issues that they will address are going to be characterized at a framework level, not based on a specific statute or regulation level.

UBIQUITOUS COMPUTING REQUIRES GLOBAL PRIVACY LAW AWARENESS

By Francoise Gilbert, Founder and Managing Director of IT Law Group and author and editor of *Global Privacy and Security Laws*

As citizens, we might feel allegiance to a particular region where our ancestors were born and our family roots were formed, but these boundaries are artificial. When looking at the earth from the 10,000-foot level, states merge into one another seamlessly. Clouds that fly over country borders ignore the passport control booths.

Like their geophysical cousins, the clouds in which our electronic files are stored and processed know no borders. Our smartphones, tablets, laptop computers, smart watches or glasses and the underlying technology into which we plug our equipment allow us to be connected at all times, from anywhere to, to anyone.

Data, like the genie, have jumped out of their bottle. They are taking a path of their own that does not stop at the edge of the device that was used to collect them or at the political border of the country in which that device is operated. With interconnectivity and ubiquitous computing available to us, we can, while seated on a bench in the middle of Golden Gate Park in San Francisco, access or modify files that are processed in Argentina by a payroll service established in France. These files may be simultaneously backed-up in Singapore and replicated for disaster recovery purposes in New Zealand. They may pertain to the employees of an Australian company who telecommute to work from South Africa.

This might look like a law school exam hypothetical. It happens increasingly in the 21st-century world of virtual companies or virtual employees where intangible intellectual property is frequently the most valuable asset of a business. Which privacy or data protection law applies to this hypothetical? Which state or country has jurisdiction over a particular dataset?

Ask five different judges, and you are likely to receive five different answers. The laws of several countries might apply, and more than one court could assert jurisdiction: That of the country where the data controller is located; that of the countries where the servers that process or store the data are located; that of the country where the data subject is physically located, or where his employer is established to do business, or where his payroll is generated.

Countries are very protective of their citizens and want to apply their laws—or are asked by plaintiff to apply their laws—to matters that may take place within their boundaries or affect their citizens. See, for instance, the current Article 3—Territorial Scope-- of the draft EU Data Protection Regulation, which is expected to supersede the 1995 EU Data Protection Directive. This provision might allow the application of the EU Data Protection laws to the hypothetical above, due to the fact that the payroll company is established in the EU, even though the data subjects are located in South Africa and their employer in Australia. Article 3 provides in part (emphasis added):

> This Regulation applies to the processing of personal data in the context of the activities of an establishment of a controller or a processor in the Union, *whether the processing takes place in the Union or not*.
>
> This Regulation applies to the processing of personal data of data subjects residing in the [European] Union by a controller or processor *not established in the Union*, where the processing activities are related to: (a) the offering of goods or services, irrespective of whether a payment of

> the data subject is required, to such data subjects in the Union; or (b) the monitoring of such data subjects.
>
> This Regulation applies to the processing of personal data by a *controller not established in the Union*, but in a place where the national law of a Member State applies by virtue of public international law.

We cannot rely on the law of a single country as the framework in which to develop policies, practices, and procedures or evaluate the risk to which data might be exposed. Ubiquitous computing, business process outsourcing, and cloud computing are available to all companies. Size no longer matters. The proverbial flower shop around the corner may have its accounting or payroll data processed or stored on another continent, in the same manner as a Fortune 10 company can.

Privacy professionals must be aware, and keep abreast of, the legal developments regarding information privacy or security laws in all the countries in which the personal data in their clients' custody are or might be located. It is only with this global knowledge and legal awareness that they will be able to properly evaluate and anticipate the legal constraints to which these data might be subject.

Although most of the world's data protection laws take an approach to the protection of personal information, personal space, and intimacy that is loosely based on similar fair information privacy principles (whether they are expressed in the OECD Guidelines, the APEC Privacy Framework, or other document), the devil is in the detail. Each country's legal framework is different. When these principles are implemented, each country has its own view and its own sensitivity to a particular topic.

Keeping abreast of these developments is difficult and time consuming. It is not that simple to know and appreciate a country's vision of privacy and what is necessary to achieve compliance in that particular country. It is a major mistake to take a one-size-fits-all approach or ignore the legal and cultural nuances among countries, even neighboring ones, or the historical foundation that have resulted in a certain legal system or certain local customs or behaviors. A formality that does not exist here may be required there and may be attached to prison terms elsewhere in cases of delinquency.

Privacy is a cross-functional and complex concept. Unlike tax, real property, or corporate law, privacy laws do not have hundreds of years of history in the making. Nevertheless, all over the world, there is more to privacy than what judges or legal scholars have designed. The social aspects and the individual, cultural, or ethnic sensitivities are also part of the foundation. Before becoming regulated, privacy has evolved in great parts outside courts, being shaped slowly by reactions to significant or traumatic events.

Privacy concepts and privacy laws may result from societal pressures, changes in mores and habits, reaction to government abuses, or may respond to technology advances. In each country, they are a reflection of the country's culture, history, and sensitivity. At times, the religious and philosophical beliefs of its citizens may have also influenced the way in which a country designed and implemented (or not) data protection principles and protected (or not) the privacy rights of its citizens. Developing a global privacy program requires an appreciation and understanding of these nuances and sensitivities.

The world of privacy and data protection is uniquely complex. As the field evolves, and, concurrently ubiquitous computing is becoming the norm, it is indispensible to take a global approach to privacy and data protection while remaining aware of the significant discrepancies between the laws, regulations, guidelines, and sensitivities that exist and will remain at the micro level in each country or state.

Privacy by Design

Privacy by Design (PbD) is a concept popularized by Ann Cavoukian, the commissioner for information and privacy for the province of Ontario, Canada. It was developed to ensure that privacy was protected and that people gained control over their information and the information of their enterprises. In 2011, at their 32nd annual conference, the international Data Protection and Privacy Commissioners recognized PbD as an "essential component of fundamental privacy protection."[5]

It teaches the following seven "Foundational Principles":[6]

1. *Proactive not* Reactive; *Preventative* not Remedial
2. Privacy as the *Default Setting*
3. Privacy *Embedded* into Design
4. Full functionality—*Positive-sum,* not *Zero-sum*
5. End-to-End Security—*Full Lifecycle Protection*
6. *Visibility and Transparency*—Keep it Open
7. *Respect* for User Privacy—Keep it *User-Centric*

[5]Resolution on Privacy by Design, 32nd International Conference of Data Protection and Privacy Commissioners, Jerusalem, Israel. www.justice.gov.il/NR/rdonlyres/F8A79347-170C-4EEF-A0AD-155554558A5F/26502/ResolutiononPrivacybyDesign.pdf

[6]Foundational Principles, Privacy by Design. www.privacybydesign.ca/index.php/about-pbd/7-foundational-principles/

NEXT-GENERATION PRIVACY FOR A NEXT-GENERATION WORLD: *PRIVACY BY DESIGN* RESOLUTION

By Ann Cavoukian, PhD, Information and Privacy Commissioner, Ontario, Canada

In October 2010, a landmark resolution was unanimously passed by the International Privacy Commissioners and Data Protection Authorities at their annual conference, recognizing *Privacy by Design* (PbD) as an "essential component of fundamental privacy protection." The Resolution also:

> Encouraged the adoption of the principles of *Privacy by Design* as part of an organization's default mode of operation; and
>
> Invited Data Protection and Privacy Commissioners to promote *Privacy by Design*, foster the incorporation of its Foundational Principles in privacy policy and legislation in their respective jurisdictions, and encourage research into *Privacy by Design.*

Since then, PbD has become a global operation, having been translated into 35 languages. Public policymakers in the United States, Europe, and Australia have issued proposals to express PbD in reformed information privacy governance and oversight regimes. More than a concept, PbD has become a legal and regulatory requirement in major jurisdictions around the world. With the world evolving so rapidly, privacy protections must also evolve in equal measure.

Evolving Privacy Contexts

Privacy is often said to be in "crisis" today as a result of numerous developments:

> Leapfrogging information and communications technology developments;
>
> The advent of social, cloud, mobile, and ambient computing;
>
> Evolving cultural norms; and
>
> A global patchwork of outdated privacy laws.

The information privacy solution requires a combination of data minimization techniques, credible safeguards, meaningful individual participation, and robust accountability measures, informed by an enhanced and enforceable set of universal privacy principles adapted to modern realities.

PbD evolved from early efforts to express Fair Information Practice principles directly in the design and operation of information and communications technologies, resulting in *Privacy Enhancing Technologies* (PETs). Over time, the broader systems and processes in which PETs were embedded and operated were also considered. These include organizational practices and networked information ecosystems. PbD principles emphasize proactive leadership, systematic methods, and demonstrable results.

Proactive Not Reactive; Preventative Not Remedial

PbD principles have changed the global privacy conversation by shifting emphasis away from reactively detecting and punishing privacy offenses after they occur to minimizing risks and preventing harms *before* they occur. "Build it in early" is now a common message from data protection authorities around the world.

PbD principles aspire to the highest global standards of practical privacy possible—to go beyond compliance and achieve visible evidence of leadership, regardless of jurisdiction. Good privacy doesn't happen by itself; it requires proactive leadership and continuous goal setting at the earliest stages.

Global leadership begins with explicit recognition of the benefits and value of adopting strong privacy practices, early and consistently (e.g., preventing data breaches and harms from arising). This implies:

A clear commitment, at the highest levels, to prescribe and enforce high standards of privacy, generally higher than the standards set out by global laws and regulation;

A demonstrable privacy commitment that is shared by organization members, user communities, and stakeholders in a culture of continuous improvement;

Establishing methods to recognize poor privacy designs, to anticipate poor practices and outcomes, and to correct any unintended or negative impacts, well before they occur, in proactive, systematic, and innovative ways; and

Continuous commitment and iterative processes to identify and mitigate privacy risks.

The preventative and systematic approach to engineering privacy is often associated with privacy-enhancing technologies, particularly in Europe. Although PbD is often best illustrated through specific technologies (the more user-centric the better), it is the *organization* that has become a more central and effective focus for applying PbD Principles, especially in view of the requirement to comply with privacy and data protection laws.

Being proactive and preventative requires a clear understanding of the strategic risks, challenges, and rewards of applying strong privacy throughout an organization and across information systems, in a comprehensive manner.

Privacy Embedded into Design

Privacy promises are not enough—they must be implemented in systematic and verifiable ways. Information and communications technologies, systems, and networks are highly complex and dynamic in nature. Data processing is interdependent and tends to be opaque in nature, requiring more trust than ever from stakeholders and users for sustainability. These are not ideal conditions for ensuring that accountability, data protection, and individual privacy will thrive.

Privacy commitments and controls must be embedded into technologies, operations, and information architectures in holistic, integrative, and creative ways:

Holistic, because broader contexts must be considered to properly assess privacy risks and remedies;

Integrative, because all stakeholders should be consulted in the development dialogue; and

Creative, because embedding privacy rights and controls, at times means reinventing the choices offered because existing alternatives are unacceptable.

A systematic, principled approach to operationalizing privacy should be adopted, one that relies on accepted standards and process frameworks, amenable to external reviews and audits. All fair information practices should be applied with equal rigor, at every design step.

Wherever possible, detailed privacy impact and risk assessments should be carried out, documenting the privacy risks and measures taken to mitigate those risks, including consideration of alternatives and the selection of metrics.

The privacy impacts of the resulting technologies, processes, and information architectures should be demonstrably minimized and not easily degraded through use, misconfiguration, or error.

In the United States, the Federal Trade Commission (FTC) has begun to require some organizations to put in place comprehensive, auditable privacy programs. In the European Union, “prior checking” and other due diligence requirements are becoming mandatory for organizations to demonstrate compliance with privacy laws.

Full Functionality: Positive-Sum Not Zero-Sum

Privacy is not an absolute value. To design practical, yet effective, privacy controls into information technologies, organizational processes, or networked architectures, privacy architects need to acknowledge many legitimate (and, yes, sometimes competing) goals, requirements, and interests and accommodate them in optimized, innovative ways.

The PbD Principle of *Full Functionality* requires going beyond privacy declarations and best efforts to *demonstrate* how data processing and other objectives have been, and are being, satisfied in a doubly-enabling, win-win model. External accountability and leadership are enhanced by applying this principle, which emphasizes transparency and measurable outcomes of multiple functionalities:

> When embedding privacy into a given information technology, process, system, or architecture, it should be done in such a way that full functionality is not impaired, and that all legitimate interests are accommodated and requirements optimized;
>
> Privacy is often positioned in a zero-sum manner; that is, having to compete with other legitimate interests, design objectives, and technical capabilities in a given domain. PbD rejects this approach; it embraces legitimate non-privacy objectives and accommodates them in an innovative, positive-sum manner; and
>
> All interests and objectives must be clearly documented, desired functions articulated, metrics agreed upon and applied, and unnecessary trade-offs rejected, in favor of finding a solution that enables multi-functionality.

Additional recognition is deserved for creativity and innovation in achieving all objectives and functionalities in an integrative, positive-sum manner. Organizations that succeed in overcoming outmoded zero-sum choices demonstrate global privacy leadership.

This principle challenges policymakers, technologists, and designers, among others, to find ways to achieve better privacy in a given technology, system, or domain than is currently the case and to document and demonstrate achievements that become best practices.

There are many examples of positive-sum "transformative" technologies that achieve multiple objectives in tandem in a privacy-enhancing manner. For example, Biometric Encryption (BE) achieves positive identification without the need for centrally stored templates. BE has been successfully deployed across Ontario gaming facilities to identify gamblers requesting to be barred from entering the premises. The positive-sum PbD principle has also been successfully applied in a

wide range of areas: road toll pricing, smart meters, whole-body image scanners, RFID-enabled systems, geolocation-enabled services, and many other technologies and services.

The creation, recognition, and adoption of PETs as a means to achieve PbD operational goals is being actively promoted by the European Commission, not only as a major ongoing research funding initiative under the Framework Programme, but notably in the context of the EU review of, and proposed amendments to, the Data Protection Regulation.

Current work by international data protection authorities to define accountability is also establishing common definitions and best practices that help advance organizational PbD practices. Similar work is also under way in international standards groups to define privacy implementation, assessment, and documentation methods. The preparation, use, and publication, whether mandatory, contractual, or voluntary, of privacy impact assessments and privacy management frameworks are also on the rise. We are seeing the growth of standardized privacy evaluation, audit, and assurance systems, innovative co-regulatory initiatives, certification seals and trust marks, and other criteria. Enhanced diligence and accountability measures are consistent with the PbD emphasis on demonstrating results. The publication of successful case studies adds illustrative and educational value for others to emulate. Perhaps the most exciting chapters on achieving PbD results have yet to be written, as public policymakers on both sides of the Atlantic Ocean actively propose weaving the PbD framework and principles into the fabric of revised privacy laws, and in strengthened systems of regulatory oversight—the best is yet to come.

Like privacy engineering, PbD teaches that privacy is also a business issue. The building of consumer trust will provide a competitive advantage. Just one data breach interferes with this trust. PbD, like privacy engineering, recognizes that both physical design and information technology design are crucial to develop an effective privacy program. The privacy designer needs to carefully construct physical security to protect the privacy of both data facilities and paper records. Information technology design can enhance privacy by the use of PETs (discussed in detail in Chapter 6) like a uniqueness identifier with no specific meaning and by utilizing encryption correctly. Security and privacy work together and do not work at cross purposes. It is important that privacy be embedded into the IT system as part of the design process, baked in so it will not interfere with the business purpose of the system but will actually enhance the business objectives.

How Privacy Engineering and Privacy by Design work Together

Privacy engineering is a concept for which PbD is a facilitator. PbD provides valuable design guidelines that privacy engineers should follow. In turn, privacy engineering adds to and extends PbD. It provides a methodology and technical tools based on industry guidelines and best practices, including the Unified Modeling Language.

In the rest of this book, we will discuss the methodologies and the various modeling processes to develop privacy mechanisms that can be used independently or can be plugged into new and existing enterprise systems to enhance their ability to implement enterprise privacy policies.

Conclusion

This chapter explained how privacy and other data management frameworks overlap and can be leveraged as an overall governance framework for personal information. Data management teams and privacy functions have common goals: the health, hygiene, and well-being of the data under their respective custodianship. While there may be different approaches to data management and different privacy frameworks, there are strong points of similarity that can be harmonized to arrive at a functional set of policies and requirements for an enterprise. Chapter 4 will discuss how these Privacy Policies are developed and how an organization's privacy policy can be coordinated as the "meta" document for use case requirements.