

Identity and Privacy Strategies

In-Depth Research Overview



Privacy

Version: 1.0, Apr 03, 2009

AUTHOR(S):

Ian Glazer

iglazer@burtongroup.com

Bob Blakley

bblakley@burtongroup.com

TECHNOLOGY THREAD:

Policy, Privacy, and Personalization

Table Of Contents

Summary of Findings.....	3
Analysis.....	5
Privacy Is Not Secrecy.....	6
Privacy Is Fundamentally Contextual.....	9
Societal Norms and Starting Assumptions.....	9
Relationships and Interactions.....	10
Domains.....	11
Legal Regimes and Jurisdictions.....	11
Sectoral Domains.....	11
Evaluating Contexts.....	11
Asymmetries of Relationship, Value, Expectations, and Power.....	12
Asymmetric Relationships.....	12
Asymmetric Value.....	13
Asymmetric Expectations.....	13
Asymmetric Power.....	13
Balancing Asymmetries with Intermediaries.....	14
A Single Strict Definition of Privacy May Not Be Desirable.....	14
But Privacy Principles Are Useful.....	15
Principles Yield Practice.....	16
The Details.....	18
Characteristics of an Effective Privacy Program.....	18
Privacy Governance Structure.....	18
Documented Principles, Policies, and Practices.....	18
The Few, the Proud, the Privacy Team.....	19
Point of Contact.....	20
Budget.....	20
Overview of the Activities of an Effective Privacy Program.....	21
Data Inventory and Classification.....	21
What, Who, and How to Ask.....	21
Who Owns the Effort and the Result.....	23
Risk Assessment.....	23
Policy Management.....	24
Privacy Impact Assessments.....	25
Privacy Audits.....	26
Education and Awareness.....	26
A Comparison of Privacy Principles.....	27
Building a Set of Contextual Privacy Principles.....	27
Normalized Privacy Principles.....	27
Conclusion.....	29
Notes.....	30
Author Bio.....	31

Summary of Findings

Bottom Line: Privacy is not about data—it's about people. Privacy is not secrecy, and it is not about hiding information. Privacy is concerned with the proper handling of personal information and with respecting the dignity of the individual to whom the information refers. The fundamentally contextual nature of the use of personal information prevents us from formulating a single strict definition of “privacy.” However, privacy principles accommodate this contextuality and guide the development of enterprise privacy practices that can reduce risk and cost.

Context: Over the course of thousands of years, society has established sets of norms and behaviors that respect people's privacy. What societies took thousands of years to do, today's enterprises have attempted to implement in code and business processes in a few decades, and in their rush they have often mistaken privacy for secrecy and people for data. Privacy is contextual; interpreting privacy principles in all the relevant contexts yields effective privacy practices.

Takeaways:

- Privacy is not secrecy:
 - Privacy is not about hiding information; rather, it's about the appropriate handling of information.
 - Privacy is not concerned with building and maintaining anonymity; once you're anonymous, you don't need privacy.
- Privacy is fundamentally contextual:
 - Societal norms and starting assumptions greatly influence expectations about privacy.
 - The nature of the relationships between parties helps to establish context.
 - The sectoral domains and legal jurisdictions of the parties involved establish context as well.
 - The evaluation of these contexts is necessary to understand privacy problems.
- A single strict definition of privacy may be neither desirable nor achievable:
 - A definition of privacy that works in one context will fail in another.
 - Holding fast to a single strict definition will ensure failures of scalability of products and services.
 - Dictating a single definition of privacy can create an adversarial relationship with potential customers, citizens, and partners.
- Privacy principles are useful:
 - They are the mission statement for enterprise privacy teams.
 - They reflect the values of the enterprise.
 - Privacy principles interpret specific aspects of Burton Group's Golden Rule of privacy: *We protect privacy when we consider the dignity of individuals about whom we know things, and when we use what we know about them only in ways which preserve and enhance that dignity.*
 - Privacy principles are the templates for privacy practices, and privacy practices can accommodate the contextual nature of privacy.
- The hard part of designing an effective privacy program is contextualizing privacy principles:
 - The contextualization process involves determining the complete set of contexts in which an organization, its employees, its customers, and its partners operate.
 - The contextualization process views personal information from the perspective of each party in each context.
 - Contextualization of privacy principles is labor intensive, detail oriented, and manual.
 - Contextualization of privacy principles must be revisited periodically; in particular, it must be re-examined when laws, social norms, offerings, processes, or affected populations change.

- Effective privacy programs have:
 - Formal governance structure
 - Written policies and practices
 - Line-item funding for privacy and the privacy team
 - Formal procedure for receiving and resolving inquiries and complaints
 - A designated point of contact for privacy issues
- Effective privacy programs:
 - Assign individual responsibility for privacy
 - Educate and raise awareness.
 - Evaluate privacy risks and set risk-based priorities using data inventories.
 - Partner with multiple groups within the enterprise to enact privacy practices.

Conclusion: Because privacy is fundamentally contextual, a privacy team's most important work is to contextualize common privacy principles into useable, meaningful privacy practices in the contexts relevant to the organization. Doing this requires collaboration with partners throughout the enterprise, strong executive support, a budget, and a formal privacy governance structure. Success in this effort builds an organization that respects the dignity of customers, employees, citizens, and partners; forms stronger relationships; and reduces risks to its finances and reputation.

Analysis

“A man, to be greatly good, must imagine intensely and comprehensively; he must put himself in the place of another and of many others; the pains and pleasures of his species must become his own.”

—Percy Bysshe Shelley, *A Defense of Poetry*

On January 20, 2009, Heartland Payment Systems, a processor of credit card transactions, issued a press release reading in part as follows:

Payments processor Heartland Payment Systems has learned it was the victim of a security breach within its processing system in 2008. Heartland believes the intrusion is contained.

“We found evidence of an intrusion last week and immediately notified federal law enforcement officials as well as the card brands,” said Robert H.B. Baldwin, Jr., Heartland's president and chief financial officer.

“We understand that this incident may be the result of a widespread global cyber fraud operation, and we are cooperating closely with the United States Secret Service and Department of Justice.”

The release went on to say that “No merchant data or cardholder Social Security numbers, unencrypted personal identification numbers, addresses or telephone numbers were involved in the breach.” However, given that the breach came to light after Visa and MasterCard detected a pattern of fraudulent purchases using credit cards whose transactions had been processed by Heartland, this statement isn't very reassuring.

It's not yet clear exactly how many accounts have been compromised; the consensus guess at this early stage is that more than 100 million accounts are involved. The cost to Heartland in the short term has been a nearly 50% loss in market capitalization as the price of the company's stock dropped on the news. The cost to reimburse banks for reissuance of scores of millions of cards is likely to be even higher; these costs, together with the loss of card issuers' trust, drove CardSystems Solutions into bankruptcy after it experienced a breach leading to the theft of 40 million cards in 2005.

Three days after the Heartland announcement, the online job search firm Monster.com announced that it too had suffered a breach:

As is the case with many companies that maintain large databases of information, Monster is the target of illegal attempts to access and extract information from its database. We recently learned our database was illegally accessed and certain contact and account data were taken, including Monster user IDs and passwords, e-mail addresses, names, phone numbers, and some basic demographic data. The information accessed does not include resumes. Monster does not generally collect—and the accessed information does not include—sensitive data such as Social Security numbers or personal financial data.

Monster has seen this movie before; in August 2007, the company experienced a Trojan-horse attack which netted the attackers 1.3 million Monster user records. Despite actions to improve security that Monster announced at the time (in a [press release](#) whose language is eerily similar to [the press release disclosing the 2009 breach](#)), the company's database was breached again. The company has not, as of the time of writing, released an estimate of the number of accounts compromised.

Legislation is being introduced to tighten the liability of payment processors which suffer breaches, and to forbid retention of transaction data once a transaction has cleared. Heartland and Monster have promised (in Monster's case, again) to tighten their security controls.

While these incidents resulted *from* breaches of security, they resulted *in* breaches of customers' privacy. However, privacy breaches aren't always caused by security failures. Sometimes privacy is breached because of deficient business practices, and sometimes it's breached through the actions of authorized (but malfeasant) insiders. Failures to protect privacy, however they arise, can erode people's trust in a business. An enterprise that respects its customers in all aspects of their interaction will have happier customers; an effective privacy program is one important method for demonstrating that respect. What is often less acknowledged is that without an effective privacy program, an enterprise may find itself, as in the case of CardSystems, without any customers whatsoever. Without effective privacy programs, services such as financial payments processors, online health record management systems, and e-government initiatives will find slow adoption as they struggle to build a critical mass of participants. In these cases, privacy isn't a “nice to have”—it's table stakes.

An effective privacy program can be a competitive differentiator. In the course of researching this overview, Burton Group found that a number of organizations are paying closer attention to the privacy policies and programs of their partners. Some of the enterprises that were interviewed indicated their desire to strengthen the privacy requirements in their contract language for their partnership agreements. By being able to demonstrate superior privacy practices, a service provider can differentiate itself from its competitors; and in today's economic situation, every valid point of differentiation is important. What start as differentiators will become minimum antes in the future—effective privacy programs will become a requirement of all partnership agreements.

Effective privacy programs reduce the risk of accidental disclosures and breaches, and in doing so, reduce the chance that the enterprise will suffer financial and reputation loss. These programs also can serve as competitive advantages. But designing an effective privacy program is hard, as Monster.com's example demonstrates. The difficulty of designing an effective privacy program stems mostly from the difficulty of following Shelley's advice to “imagine intensely and comprehensively” and to “put oneself in the place of another and of many others.” To handle privacy successfully, an enterprise must step outside itself and put itself in the place of the individuals whose private information it handles.

In this overview, Burton Group sets forth the essential elements of a privacy program that imagines intensely and comprehensively. The elevator-pitch version of these essential elements is:

- Privacy is not about data, it's about people. Successful organizations instill in their personnel the core belief that when they are handling personal information, they are not shoveling bits—they're affecting people's lives.
- Protecting privacy is not a technology problem; it's a social problem. The core of privacy is not keeping data secret; it is ensuring that people's dignity is respected at all times.
- Privacy is fundamentally contextual. No single set of rules or practices protects privacy in all circumstances.
- Because privacy is a social phenomenon, requiring awareness of human dignity in a particular context, privacy cannot be protected by an automated system. People must always be in the privacy loop.
- Despite privacy's contextuality, there is a “golden rule” of privacy.
- A set of core privacy principles can be derived from the Golden Rule, and these principles are widely accepted.
- Transforming core privacy principles into an effective privacy program is an ongoing effort that requires a dedicated organization, support from the top, and continuous improvement.

Privacy Is Not Secrecy

Discussions of privacy (and of the impact of technology on privacy) have tended to focus on “information privacy” and, more specifically, on anonymity and the secrecy of personal information.

But privacy is not secrecy. In fact, the need for privacy arises only when secrecy is no longer an option. People rely on each other to protect the privacy of sensitive information, which, for one reason or another, they have chosen to reveal rather than to conceal.

It's useful to draw explicit distinctions between confidentiality, secrecy, and privacy. Confidentiality is the obligation to hold information in confidence—that is, not to disclose information to parties who should not have it. Confidentiality obligations can arise for a variety of reasons. Information can be confidential because it is sensitive and personal, and therefore private. Information can also be confidential because it is the trade secret of a business or because its release would adversely affect national security. Burton Group discussed data confidentiality in the *Security and Risk Management* technical position “[Information Confidentiality](#).”

Secrecy is a means of protecting confidentiality. If I have an obligation to keep information confidential within the walls of my business, I can encrypt it to provide secrecy protection, but I can also do other things. I can, for example, stamp it with a rubber stamp to ensure that other employees, if they come across the information, will understand that it is confidential and behave accordingly.

Privacy is the respect we pay to other people's dignity by treating their personal information with sensitivity. Private information must often be kept confidential, but confidentiality is only part of the privacy obligation. We are also obliged to use private information properly. So, for example, a doctor should use information about his patient's medical condition only to treat the patient—and not as a source of after-dinner jokes—even if the other people attending the dinner are also doctors and are also treating the patient, and there is therefore no concern about improper disclosure of the information. In a business context, the doctor also should not use information about the life-threatening illness of a corporate chief executive officer (CEO) to place trading orders for the CEO's company's stock (this would be a violation of insider-trading rules as well as of the CEO's privacy). And using information in appropriate ways also includes an obligation to communicate in a dignified way with the individual whose information is being used. Society would be outraged if a medical testing laboratory delivered lab results over the telephone by playing a recording saying, “You have brain cancer. Thanks for using our company's services.”

The core problem of privacy is not, therefore, how to protect the secrecy of personal information; it is how to ensure that those to whom the information is disclosed handle it appropriately.

In contexts other than technology, this has been understood for millennia. The Hippocratic Oath defined both privacy practices and the sanctions for violating them almost 2,500 years ago:

What I may see or hear in the course of the treatment or even outside of the treatment in regard to the life of men, which on no account one must spread abroad, I will keep to myself, holding such things shameful to be spoken about. If I fulfill this oath and do not violate it, may it be granted to me to enjoy life and art, being honored with fame among all men for all time to come; if I transgress it and swear falsely, may the opposite of all this be my lot.¹

The photo in Figure 1 shows an old privacy-enhancing technology; it's the confessional of the Cathedral of St. Vitus in Prague.



Figure 1: *Confessional of the Cathedral of St. Vitus, Prague*

The purple curtain conceals the penitent so that casual passersby cannot tell whose sins are being enumerated. The curtain is secrecy at work; the penitent uses it (as on the Internet he might use a Microsoft Windows CardSpace information card) to keep private information from falling into the hands of those who don't need to have it—but the curtain is the lesser of two protections the confessional offers. The greater protection is the priest's oath to respect the “seal of confession.” The *Catholic Encyclopedia*² describes this rule and the penalties for its violation:

Let him [the confessor] beware of betraying the sinner by word or sign or in any other way whatsoever . . . we decree that he who dares to reveal a sin made known to him in the tribunal of penance shall not only be deposed from the priestly office, but shall moreover be subjected to close confinement in a monastery and the performance of perpetual penance.

This rule was already well established when the quoted language was adopted by the Fourth Lateran Council in 1215. The *Catholic Encyclopedia* goes on to note that:

. . . by a decree of the Holy Office (18 Nov., 1682), confessors are forbidden, even where there would be no revelation direct or indirect, to make any use of the knowledge obtained in confession that would displease the penitent, even though the non-use would occasion him greater displeasure.

Without the seal of confession, the booth's purple curtain is just a decoration. Without its equivalent, today's “privacy enhancing technologies” are also just decorations.

The main business of privacy today is not to build anonymity and pseudonymity technologies; it is not to settle opt-in versus opt-out arguments about user consent; and it is not to make identity systems “user centric” by flashing personal information in front of people's eyeballs before sending it to others.

The main business of privacy today is to build our privacy oaths, take them, and honor them; in other words, to find out what people want us to do with information about them, to promise to do that, and to build a set of practices which help us to keep that promise.

Privacy Is Fundamentally Contextual

Privacy is not a static object with a discrete set of attributes and actions. It is neither directly observable nor measurable. Privacy's protean nature is the source of its value and the source of its challenges.

Privacy is, fundamentally, contextual. Any question about privacy must be understood in the context of:

- The starting assumptions and principles of the parties
- The relationship between the parties
- The interaction between the parties among which private information is shared
- The domain (e.g., sector, nation, etc.) in which the parties are interacting
- The societal norms to which the parties adhere

Minor variations in any one of these contextual aspects of the situation can lead to major differences in the privacy practices that should be applied. As in chaos theory—the butterfly flapping its wings—seemingly minor differences in relationships, societal norms, starting assumptions, or other aspects generate radically different definitions of privacy.

The following sections describe some of the contextual elements of privacy. In understanding contexts like these, product designers, program managers, and even architects will be better able to assess the privacy-related risks of a new endeavor.

Societal Norms and Starting Assumptions

The starting assumptions of the parties in a relationship and the societal norms they are steeped in influence their expectations of privacy. Societal norms shape social behavior as well as expectations. One norm is that a doctor will not take advantage of a patient based on a piece of information the doctor has learned. (While this norm is also codified in various ways into tort law, the following example focuses on its existence solely as a norm.) A patient knows this norm and knows the doctor knows the norm, and in that knowledge, the patient discloses symptoms of a condition to the doctor. If the doctor attempts to blackmail the patient because the patient has a potentially embarrassing condition, society (as well as the medical board and others) would view the doctor's actions as violating the patient's privacy.

Societal norms are often nearly invisible to the parties in a relationship. Each party expects that the other party in the relationship understands and acts in accordance with the same norms. It is unlikely that a party in a relationship (e.g., a doctor's patient) would have specific privacy expectations that enumerate all the societal norms at play—each party simply expects social norms to be upheld even if they are unstated. And herein lies a cautionary word to service providers who expand their offerings to new regions: A service provider whose offerings have privacy implications needs to take care to understand the societal norms of the region into which it is expanding, or it will face the consequences of violating those norms. In some cases, moving into new regions will require segregating data along geographic lines. The government of Canada, for example, requires companies handling certain kinds of data about Canadian citizens to ensure that the data will not be transmitted into the United States because of concerns about access to the data under the terms of the USA PATRIOT Act.

Like societal norms, starting assumptions are rarely fully and explicitly enumerated. Starting assumptions guide behavior, occasionally in non-rational ways. Consider a recent college graduate and a retiring baby boomer. The recent graduate's starting assumptions about privacy stand in stark contrast to the baby boomer's. The recent graduate, a “digital native,” may consider the default privacy settings of a social networking site good enough; after all, the social networking site is run by responsible older people, right? When the graduate discovers that what was thought to be a semi-private space (“just friends”) turns out to be a public space, she is shocked and outraged. The retiring baby boomer, being a “digital tourist,” is far more suspicious of online interactions. If he eventually decides to “do it in the Facebook,”³ more likely than not, he won't consider the privacy settings either. When someone unexpected calls, having recently read something the baby boomer shared, he'll shrug his shoulders, saying, “I wouldn't have posted it online if I thought it needed to be private; I have no privacy online anyway. . . .” The graduate and the baby boomer have different starting assumptions; therefore, while their behaviors are similar, their reactions (their sense that their privacy has been injured) are quite different.⁴

Relationships and Interactions

The nature and structure of the relationships between parties is the largest contributor of context to the evaluation of a privacy-related situation. Relationships codify the “terms and conditions” under which parties interact, become more intimate, and even sever their association. Relationships provide the rules of the game when parties interact; understanding these rules is the first step to understanding the contextual nature of privacy for a given situation.

For example, consider Elizabeth. Elizabeth has a relationship with her healthcare provider: She is a patient. This relationship allows the healthcare provider to gather information about Elizabeth and to share that information with other parties with Elizabeth's consent. Elizabeth, as part of the terms of the relationship agreement, can inspect this information and instruct the provider about her preferences for sharing the information with third parties. In contrast, consider Elizabeth's relationship with the same healthcare provider when she acts on behalf of her elderly father. The terms and conditions of this relationship are governed, at least in part, not by Elizabeth, but by her father and his instructions to the healthcare provider. Without understanding the relationships between Elizabeth, her father, and the healthcare provider, we cannot accurately interpret Elizabeth's actions as an invasion of her father's privacy. Healthcare providers can't simply throw up their hands in situations like this; they need to know enough about Elizabeth's relationship with her father to decide what uses of Elizabeth's father's information are permissible and what uses are forbidden. The healthcare provider needs, in Shelley's words, to “put [itself] in the place of another [Elizabeth's father] and of many others [Elizabeth and her father, and maybe her lawyer, too]” in order to do its business.

Although studying relationships at rest is interesting, it does not provide as much contextual enlightenment as understanding relationships in motion. Considering Elizabeth again, the definitions of her relationships to the healthcare provider and to her father are contextually interesting, to a point. However, when Elizabeth wants to obtain information from the healthcare provider on behalf of her father, when the parties interact, the situation is far more interesting from a privacy perspective. It is during these interactions that friction, cognitive dissonance, and mismatched behavioral norms are observed.

Further, interactions between parties have a value to each party. Those values add another component to the contextual nature of privacy. In Elizabeth's case, she places a high value on receiving medical information she has requested, and her healthcare provider similarly assigns a high value in the interaction. Because of the high-value nature of the interaction, both parties have an interest in keeping the interaction private. To that end:

- The healthcare provider will take extra steps to ensure that the person requesting Elizabeth's information is Elizabeth.
- Elizabeth will attempt to verify that she is interacting with her healthcare provider and not a third party.

Contrast the value of this interaction with the following scenario. Elizabeth wants to know what her healthcare provider's office hours are. This is a fairly low-value interaction to both parties, and neither party has privacy concerns. To that end:

- The healthcare provider will take no steps to identify Elizabeth.

- Elizabeth will take only minimal steps to ensure she is interacting with the healthcare provider.

Domains

Besides relationships, understanding the domain in which parties interact is crucial to evaluating the privacy aspects of a situation. This overview uses the term “domain” to describe a set of parties who share a like set of privacy-related customs, obligations, and requirements. Domains can be as limited as a collection of online first-edition book retailers or as broad as the society of the United States. Examples of domains include:

- The healthcare sector
- Outsourced IT services
- Government
- An online social network
- A legal regime and its jurisdiction (e.g., “Grenada” or “the European Union”)

Legal Regimes and Jurisdictions

It is impossible to consider the contextuality of privacy-related situations without considering the backdrop of legal regimes. Domain members operate in a framework of laws (and often in multiple frameworks). The privacy requirements of different frameworks can be in opposition to one another. The European Union (EU), following in the tradition of John Locke,⁵ has constructed a set of privacy laws that categorize private information as property and then apply property laws to that information. However, the current legal framework in the United States does not treat private information as property. Given that domain members and other parties tend to obey laws, it is critical to explore the legal regimes to which all parties are subject when trying to understand the context of a situation from a privacy perspective. It is important to note that the consideration of legal regimes and frameworks must extend beyond laws to regulations, directives, interpretations, and judicial decisions.

Sectoral Domains

Affinities of privacy practice often exist among different enterprises in the same industry or sector. All enterprises in a sector may be subject to a similar set of privacy obligations, which can include regulatory requirements as well as voluntary industry practices. The U.S. Health Insurance Portability and Accountability Act (HIPAA) is an example of such a regulatory requirement for the healthcare sector; the privacy principles of the Network Advertising Initiative (NAI) are an example in the online advertising sector. Even within a single sector, such as government, subdomains may exist to account for sufficiently different sets of requirements. The U.S. federal government can be broken into at least two subdomains to recognize the special requirements of intelligence agencies.

While domains have unique obligations and requirements, parties in relationships with domain members bring with them differing sets of expectations. These differing sets of expectations help shape the interactions between parties. Elizabeth may have one set of expectations when interacting with her healthcare provider and another set of expectations when buying a book from an online retailer. These expectations bolster the parties' mental model for working with members of particular domains. For example, it would be in line with Elizabeth's expectations if her healthcare provider asked for her Social Security number (SSN), whereas it would be highly unusual and unwarranted for the online retailer to do the same.

Evaluating Contexts

Because of the wide variety of domains, legal regimes, types of relationships, and societal norms, this overview cannot provide a prescriptive specific approach for evaluation. A successful evaluation process is an organic one that seeks the input of multiple constituents, from general counsel to the marketing department to users and data subjects. It is important to note that the order in which contexts are presented in this overview does not imply a precedence or hierarchy of importance. Furthermore, this order does not imply a correct order of evaluation. As new projects and products are proposed, some contexts may naturally become more important than others. In the example of debuting a service to a new region, understanding societal norms of the target market would likely be the most important context. In the case where an enterprise is about to form a partnership with another company in a different sector, evaluating the sectoral context may rise to the fore. Overall, the evaluation of contexts requires multiple viewpoints and will serve the enterprise well when its privacy teams attempt to formulate [privacy impact assessments](#) (PIAs).

Asymmetries of Relationship, Value, Expectations, and Power

When both parties in a relationship acknowledge their relationship, when the relationship has similar value to both parties, and when both parties have a similar set of expectations, the privacy issues may be fairly mundane. Elizabeth knows her healthcare provider and her provider knows her. They both treat their interactions as highly valuable and both expect and strive to maintain a level of privacy that respects their relationship. However, when asymmetries exist in a privacy-related situation, the situation becomes far more complex and often results in one party's model of privacy running afoul of the other's model. These asymmetries take four forms:

- Asymmetric relationship, in which one party has no knowledge of or does not acknowledge the other party and thus does not recognize the relationship itself
- Asymmetric value, in which one party derives high value from interactions, but the other party derives low value
- Asymmetric expectations, in which one party expects the other party to behave in ways in which the other party does not expect or intend to behave (particularly with respect to private information)
- Asymmetric power, in which one party has disproportionate ability to cause damage to the other party without itself being subject to punishment or retribution

Asymmetric Relationships

Being ignorant of or denying the existence of a relationship creates an asymmetry. This type of asymmetry is often exemplified in people's mailboxes. Receiving a catalog from a company with whom a person has never shopped is a relatively harmless asymmetric relationship. Often, people throw the catalog out, shrug their shoulders, and say, "Someone must have sold my name to this other company." Far more potentially harmful asymmetric relationships can exist.

Consider the case of Dwayne Cross. From 2002 to 2007, Mr. Cross worked at the Bureau of Consular Affairs at the U.S. State Department, and during that time he used his access to State Department computer systems to view the passport files of celebrities, musicians, athletes, and politicians, including then-presidential candidates Barack Obama and John McCain. This case presents two different asymmetric relationships. First, it is highly likely that of the more than 150 people whose records Mr. Cross examined, none of them knew or knew of Mr. Cross; they had no relationship with him. Second, of those 150 people, very few, if any, knew that the State Department maintained such a system of records; in their eyes, they had no relationship with the State Department.

Asymmetric relationships have a dull, pervasive menace to them. In cases where a person does not know that information is being collected about them and by whom, how that information is being shared and with whom, and what that information is being used for, there is a substantial risk of injury to the dignity of the person whose information is being collected. Because the collector does not acknowledge a relationship with the data subject, the collector does not feel any obligation to the data subject. But this view is shortsighted. Just as forming a relationship causes information to be exchanged between parties, it is also true that collecting information about another party creates a relationship with that party. Society's norms eventually force data collectors to acknowledge their relationships with data subjects and to behave in "social" ways. Many organizations that suffer data breaches have learned this lesson the hard way, after being forced by law or simple customer pressure to offer credit monitoring or other expensive forms of compensation to data subjects with whom they had no formal privacy agreements.

Asymmetric Value

A woman driving across country stops to buy soda and pays with cash; she considers the transaction with the merchant to be worth the 75 cents she pays. The merchant has little reason to expect future transactions and considers the transaction to be worth the 75 cents as well. Both parties are satisfied with the outcome of the transaction. That satisfaction—in fact, the whole transaction—would have been in jeopardy had the merchant considered the transaction to be worth \$75 and required the driver to provide personal information (e.g., a credit history) in order to buy the soda.

As seen in the previous example, when all parties assign similar value to a transaction, all parties are satisfied with the transaction and few privacy implications arise. But, when one of the parties values the transaction differently and sets conditions on the interaction based on that value, the other party is not the least bit satisfied. This asymmetry of value can lead one party to:

- Cancel the transaction, in which case both parties are unsatisfied
- Provide false information, in which case one party is unsatisfied but doesn't know it yet
- Complete the transaction but vow never to interact with the other party again, which doesn't fully satisfy both parties

Solutions to the problems of relationship asymmetry are discussed in the *Identity and Privacy Strategies* overview "[A Relationship Layer for the Web . . . and for Enterprises, Too](#)" in the section titled "The Future of Identity Is Relationship."

Asymmetric Expectations

When one party has a set of expectations for a domain and those expectations are met, there is little to note in the context of privacy. However, when those expectations are not met, typically in cases where a domain member is over-exuberant in its desire to collect, disclose, or use information, then noteworthy friction and tension arise. In the example in which the online retailer asked Elizabeth for her SSN so that she could buy a book, the retailer's behavior certainly did not map well to Elizabeth's set of expectations for a member of the online retail domain.

Because people have different sets of starting assumptions and expectations, they can be thought of as members of different domains as well. Elizabeth, a member of the U.S. citizen domain, and Roland, a member of the Estonian citizen domain, likely have different expectations for the use of national identity cards. Estonia has an advanced national identity system and, through societal norms, Roland likely views his national identity card as a mandatory tool to get things done, including voting and paying for public transportation. Because the United States does not have a national identity card and thus lacks the social norms associated with it, Elizabeth would consider requests for her passport within the United States to be strange and uncomfortable. An enterprise that wants to expand from Estonia to the United States must consider the differences between these domains, so as to not get caught by asymmetries of expectations.

Asymmetric Power

When one party in a relationship has disproportionate power (because it has more resources, more information, and more influence than other parties), it can damage the interests of other parties without worrying too much about the consequences.

Large corporations, because they have lawyers, lobbyists, and insurance, can collect personal information, use it or disclose it improperly, and still survive. They can, in other words, create what economists call “externalities”—situations in which their actions cause harm to others but they are not required to be accountable for fixing or paying for the harm. It's tempting to say that individuals, small companies, and other less-powerful entities should avoid entering into relationships with more-powerful entities who are likely to harm them, but this is not always practical. It's hard to exist in the modern industrialized world without a telephone and a credit card, but it's also hard to find a phone company or credit card issuer which is small enough to exist on equal-power terms with individuals and small businesses.

Privacy power asymmetries are found today in the complex, take-it-or-leave-it terms of Gramm-Leach-Bliley disclosures mailed by banks to their customers, in the Terms of Service and Privacy Policy statements found on websites, in the government information collection provisions of laws passed in the wake of terrorist incidents, and in many other contexts. But asymmetric power is fragile; Facebook recently found that it had the right to change its terms of service, but it did not have the ability.

Concerted action by large numbers of users, subscribers, or citizens can equalize power asymmetries; organizations holding privacy-sensitive information should avoid exploiting power asymmetries to do things that traumatize people to the extent that they take concerted action to pass laws and regulations constraining the organization's behavior (or to the extent that they simply take their business elsewhere).

Privacy regulatory regimes which require demonstration of harm before applying sanctions create power asymmetries by placing presumptions of innocence on large, wealthy organizations, and placing burdens of proof on individuals; these regulations may encourage businesses to improve their privacy practices but they seldom produce compensation or satisfaction to individuals whose privacy is compromised, because individuals can't demonstrate harm without at least hiring a lawyer. In recognition of the power asymmetry imposed by requirements to demonstrate harm, recent privacy regulations (California AB 211, for example) mandate compensation for privacy breaches even in the absence of demonstrated harm.

Balancing Asymmetries with Intermediaries

It is unrealistic to expect that an enterprise could foresee and accommodate every party's starting assumptions, expectations, and assignment of transactional values. That being said, there is a way to reconcile the needs and expectations of parties in asymmetric relationships—through the introduction of intermediaries. An intermediary can serve as a translator of social norms, behaviors, and expectations from one domain to another. It can normalize the value assigned to transactions. Further, it can better manage its relationships with a variety of enterprises in multiple domains, allowing a person to have a single relationship—between the person and the intermediary—and leave the management of the other relationships to the intermediary, thus mitigating some of the effects of asymmetric relationships. For more information on intermediaries, see the *Identity and Privacy Strategies* overview “[A Relationship Layer for the Web . . . and for Enterprises, Too](#)” and the blog post on [Identity Oracles](#).

A Single Strict Definition of Privacy May Not Be Desirable

Societies grow and change over time, and with them, perspectives change. As societies mature, notions of privacy and societal norms associated with those societies change as well. Consider that in ancient Greece, public nudity was commonplace, and being seen in the nude was not considered a violation of privacy. In today's society, public nudity is often an arrestable offense, and being seen in the nude is most likely a violation of privacy. Societies and their associated norms are not homogenous. The variances in customs, traditions, and morals are not only found in societies around the world but also in other domains such as industries and professions. These variances are not just historical curiosities, but necessities. Adopting a common definition of privacy would unnaturally force homogeneity and lead to the normalization of these variances.

Attempts to define “privacy” fail because any strict definition runs aground on the rocks of context; definitions prescriptive enough to be applied without significant thought necessarily assume a particular context and will be bad fits for other contexts. A strict definition would therefore not only freeze societal norms at an arbitrary point in space and time, but would also freeze behaviors associated with that set of norms. Furthermore, a globally common definition would necessarily require one culture to dictate the cultural norms associated with its privacy principles on other, differently organized, cultures.

But Privacy Principles Are Useful

However, organizations must eventually define privacy in every situation in which it matters; it will be helpful to have at least a starting point from which they can build the (possibly many) contextual definitions they require.

Burton Group believes the starting point of privacy is respect for the dignity of individuals. Burton Group posits the following Golden Rule:

We protect privacy when we consider the dignity of individuals about whom we know things, and when we use what we know about them only in ways which preserve and enhance that dignity.

It's important to understand that this statement rules out the possibility of any program or machine “protecting a user's privacy.” The designers of a system can consider the dignity of individuals and ensure that the system doesn't degrade people's dignity—but it's the designers, not the system, who protect privacy in that case. And the system, left to itself, will never rethink its actions in light of new cultural norms or new sets of facts—so the system likely will eventually become a privacy hazard.

Note also that the Golden Rule doesn't distinguish between information which we collect with individuals' consent and information which comes into our possession in other ways. Our inclination to respect people's dignity is based only on the fact of our knowledge and not the mechanism by which we received the knowledge.

Burton Group's Golden Rule, of course, is too general to be a guide to the development of concrete privacy practices. But the Golden Rule motivates a set of common, shared privacy principles that can help guide the development of the following practices:

- **Governance:** Disrespect toward others is an affront to dignity; therefore, we establish formal and effective procedures for operating and continuously improving the privacy program, which ensures that we will treat others with respect.
- **Accountability for all parties:** Harming others is an affront to dignity; therefore, we hold all parties, including the data subject, accountable for their use and disclosure of personal information.
- **Transparency:** Deceiving others about our intentions or actions toward them is an affront to dignity; therefore, we allow people to see how we handle information about them.
- **Meaningful choice:** Robbing others of the ability to exercise their free will is an affront to dignity; therefore, we allow people to decide how we will use information about them. When presenting people with choices about how we will be allowed to use their information, we design easy-to-understand interfaces, which reduce the possibility of confusing people, and we avoid creating “Hobson's choice” situations in which people are forced to choose the lesser of a set of evils.
- **Minimal collection and disclosure:** Prying into the affairs of others is an affront to dignity; therefore, we ask for as little personal information as possible, and we disclose personal information only to those with a demonstrable need to access the information.
- **Constrained use:** Taking advantage of the generosity of others for our own gain or convenience is an affront to dignity; therefore, we use information about people only for purposes they have authorized in advance.
- **Data quality and accuracy:** Spreading gossip and crediting rumor are affronts to dignity; therefore, we strive to ensure that personal information is accurate, timely, and relevant.
- **Validated access:** Voyeurism is an affront to dignity; therefore, we acknowledge the information we collect and we allow people to access the information we have about them. We validate the identity of people asking for access to their own information because carelessly revealing information about people to strangers is inconsistent with dignity.

- **Security:** Endangering others for our own gain is an affront to dignity; therefore, we take active and effective steps to protect personal information in our possession against malicious or accidental access, misuse, or corruption.

Even these principles (which are compared with a wide variety of sectoral, national, and international privacy principles in “The Details” section of this overview) do not translate directly to practice. This difficulty in translating principles directly into practice isn't unique to privacy practices: Consider IT management practices. One set of generic practices would certainly not meet the needs of nor be appropriate for all enterprises in every sector and region. Organizations often use IT frameworks, such as Control Objectives for Information and related Technology (CobiT) or Information Technology Infrastructure Library (ITIL), as the basis for their own practices; these organizations' IT policy teams adapt these principles and standard practices to “meet the needs of the business,” which is to say that the teams contextualize these principles and practices. Developing privacy practices is a similar endeavor.

Principles Yield Practice

Putting principles into practice is where organizations do Shelley's work of imagining intensely and comprehensively, and putting themselves in the place of another and of many others. The privacy team does this work. Burton Group recommends that organizations use the following process for describing privacy practices:

- 1 Acknowledge the importance of privacy (agreeing to live by “the Golden Rule”).
- 2 Establish accountability for privacy by [funding](#) and [staffing a team](#).
- 3 Adopt a set of common shared principles (either those listed above or one of the frameworks listed in the “[A Comparison of Privacy Principles](#)” section of this overview).
- 4 Derive a set of written [privacy practice guidelines](#) from the principles.
- 5 Implement the practices.
- 6 Audit the practices periodically and make improvements where necessary.

The hardest work in this process is done in Step 4. To derive practices from principles, the privacy team must dedicate a lot of effort to *contextualizing the principles*.

The photographer Minor White advised his students that they should photograph an object not just to show “what it is” but also to show “what else it is.” Contextualizing privacy principles is a similar exercise. The privacy team must ask itself, for every data item it's tasked with protecting, not just “what it is” (that is, what the organization uses it for in the organization's own domain), but also “what else it is” (that is, what the data subject uses it for, what business partners use it for, what identity thieves would use it for, and so on) in every domain the organization is likely to encounter.

Doing this requires the privacy team to act as a kind of mirror to the organization. The team holds up a home address to the enterprise and says, “When you look at this address, you see something to print on an envelope containing a marketing brochure. But when a national data commissioner looks at it, he sees something he can collect a thousand-dollar fine for if he finds it lying on a backup tape in a Dumpster. And when a car thief looks at it, he sees a place he can go to steal the model of the luxury car your organization sells. And when an estranged husband looks at it, he sees a place he can go to attack his ex-wife. If you look in my mirror, you can see all these things, and you can think about them before you decide how you're going to handle this piece of data.”

The process of contextualizing principles is complicated. It requires thinking about lots of different kinds of data in lots of domains and going through the “what *else* is this” analysis over and over and over again for each datum in each domain. The work of contextualization can't be automated; humans have to do it, because only humans understand the full range of implications each context has. And it's collaborative; having more minds equals getting better results in the contextualization process. This is where the privacy team will gain maximum benefit from reaching out to partner organizations in the enterprise. Including people from lots of parts of the organization in the contextualization process will have the added benefit of making a lot of people deeply familiar with why privacy really matters.

It's through contextualization that principles yield practice. After evaluating the various domains within which the enterprise works, after consulting the legal department to understand jurisdictional obligations, and after speaking with the ombudsman and users, a privacy team possesses an inkling of understanding into the privacy contexts with which they must grapple. Viewing these contexts through the lens of its privacy principles, a privacy team can begin to work with its partners in the enterprise to build sound privacy practices that are contextually aware and meaningful.

Like a physics exam in college, the process of transforming principles into practices requires the privacy team to show its work. Privacy principles grant “moral authority” to a privacy team, but that moral authority erodes if it's not clear that the principles have been thoughtfully applied in order to create the practices. A privacy team can and should market the principles, both internally and externally. The team can overcome any claims that its derived privacy practices are arbitrary or ill-fitting by demonstrating that the privacy practices are a logical extension of privacy principles in view of the relevant domains and contexts. Furthermore, especially for large, multinational enterprises, “showing the work” of transforming principles into practice serves as a training tool for regional privacy teams that must adapt corporate practice to a more restrictive, focused regional context. The tools and methods a privacy team uses to transform principles into practices are discussed in “The Details” section of this overview.

The Details

Building an effective privacy program requires specific organizational structures and specific activities.

Characteristics of an Effective Privacy Program

In interviews with organizations of various sizes and in various sectors, Burton Group has observed that effective privacy programs share the following attributes:

- Formal governance structure
- Written policies and practices
- Line-item funding for privacy and the privacy team
- Formal procedures for receiving and resolving inquiries and complaints
- A designated point of contact for privacy issues

These attributes do not ensure success in and of themselves, but including them in a privacy program can greatly increase an organization's chances of successfully addressing privacy concerns.

Privacy Governance Structure

In order for an enterprise privacy program to be effective, a formal privacy governance structure must be established. The governance structure requires funding for privacy and the privacy team as well as executive sponsorship. Without both, no matter how well intended, the privacy program can only be moderately successful in a limited set of areas and activities.

In any governance structure, executive sponsorship is important; but in the case of privacy governance, executive sponsorship is essential. As privacy policies and procedures affect most aspects of the enterprise, the privacy governance structure needs the support of the CEO. Due to the small size and resources of the privacy team, the chief information officer (CIO), chief information security officer (CISO), and general counsel also must be active supporters of the governance structure and the privacy team.

Regardless of the size of the enterprise, information about the state of privacy and privacy programs should flow up to an overall privacy governance team. This governance team should include executive representation from the CIO, CISO, chief privacy officer (CPO), general counsel, and privacy ombudsman (or whoever is responsible for representing individuals' views of privacy issues). The CPO (or equivalent) takes the lead in reporting the state of privacy initiatives to the governance team. To gather program state information, the CPO relies not only on the privacy team but also on various partners of the privacy team. In large organizations, the CPO often must gather information from line-of-business (LOB) and regional privacy champions or teams. How the privacy team is structured is handled in more detail in "[The Few, the Proud, the Privacy Team](#)" section of this overview.

Documented Principles, Policies, and Practices

A privacy program without a documented set of principles, policies, and practices is very unlikely to be effective. Before jumping straight into details such as information protection procedures, privacy teams need to document the enterprise's privacy principles. Often, these principles stem from a set of jurisdictional or sectoral domain privacy principles, such as those developed by the Organisation for Economic Co-operation and Development (OECD) or the American Institute of Certified Public Accountants (AICPA). This early guidance from the privacy team will shape all future work; in a sense, these principles serve as a mission statement for privacy in the enterprise. These principles lead, through a process of contextualization, to privacy practices and policies.

Once a set of principles is in place, the privacy team can begin, often with the help of other teams, to construct privacy policies and practices. One of the first of these practices must be a formal procedure for receiving and resolving privacy-related inquiries and complaints. Both internal and external sources of privacy inquiries exist, and complaint and privacy policies and practices must handle both sources. As seen in the “[A Comparison of Privacy Principles](#)” section of this overview, procedures for redress are common to privacy principles. Often these principles are external facing, focusing on the needs of users, customers, and citizens, but similar needs are present internally. Internal, especially whistleblower, procedures are needed as well.

Privacy compliance policies and procedures are by no means the only policies and practices that an effective privacy team must create. Working with information security (and potentially records management), privacy teams must build information-handling policies. This topic is covered in greater detail in the “[Policy Management](#)” section of this overview.

The Few, the Proud, the Privacy Team

Among the organizations Burton Group interviewed for this overview, privacy teams are always small, with an average size of 1 person and a maximum size of less than 10 people. The observed ratio of privacy team members to number of employees ranges from around 1:6,000 to 1:40,000. Such ratios illustrate one of the primary reasons why an enterprise must have a formal privacy governance structure with strong executive support. Furthermore, as privacy teams must rely on partnerships with other organizations, executive support is even more critical.

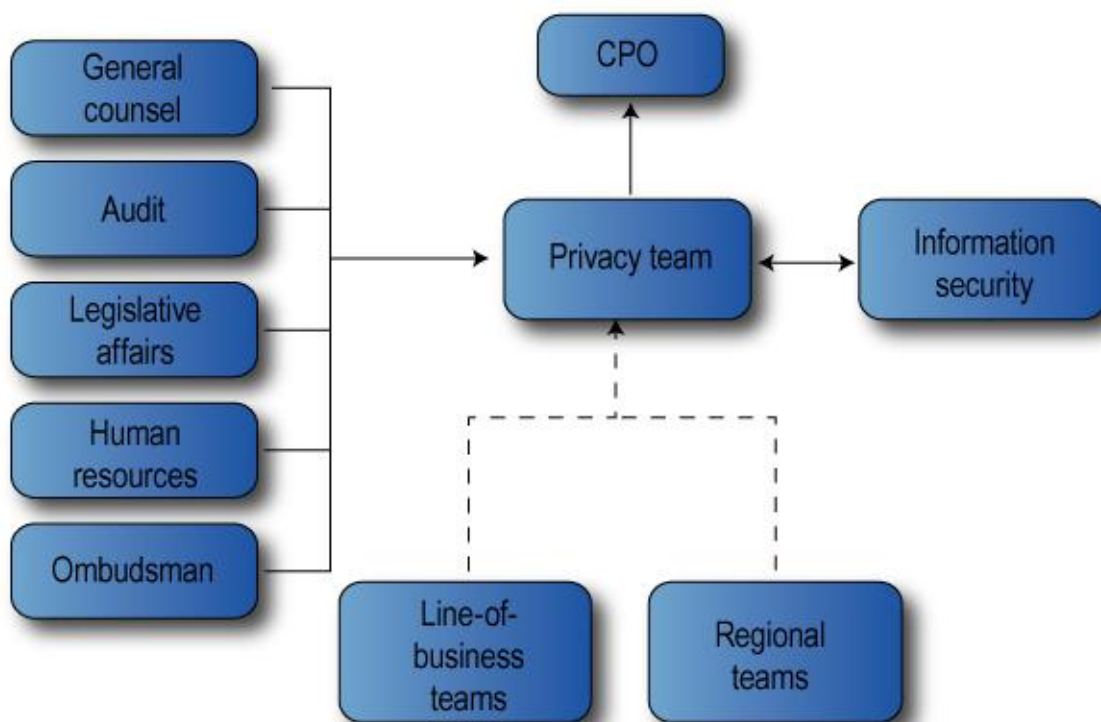


Figure 2: *Common Privacy Team Partnerships*

Figure 2 shows that the privacy team relies on many partners to be effective. A privacy team's closest working partners are likely to be the information security department and the office of the corporate counsel (or, in educational or government organizations, the general counsel). These teams must work together to build policies and procedures that are fundamental to the privacy efforts of the enterprise. But these aren't the only partnerships a privacy team has; privacy teams commonly must:

- Advise the general counsel's office of the potential impact of privacy laws
- Work with the general counsel's office to build policy
- Alert the legislative affairs department about proposed legislation and regulations that the enterprise needs to track or attempt to influence
- Receive results from audit teams and suggest future audits and audit requirements
- Work with HR (or its equivalent; e.g., admissions or student affairs, in the case of a university) to resolve issues arising from privacy concerns

In short, an effective privacy team must be skilled in maintaining working relationships with numerous teams in the enterprise, addressing tactical and strategic issues, and setting and sticking to a series of risk-based priorities.

As privacy teams are typically so small, they cannot be aware of the privacy-related activities in each LOB and each region of the enterprise. Especially in large, multinational enterprises, the privacy team relies on in-country, in-LOB staff members whose duties include some privacy work. These staff members may have dotted-line reporting responsibilities to the CPO, but it is rare that they directly report to the CPO. As an alternative to a dotted-line reporting structure, some enterprises, typically those with a matrixed structure, use privacy councils to ensure that the privacy concerns of the enterprise reach every LOB and region. Regardless of whether the enterprise is matrixed or not, the CPO is responsible for relaying privacy information between the overall privacy governance team and regional or front-line privacy champions.

Often a CPO-like position is not created in the enterprise until after a breach or accidental disclosure has occurred. Burton Group found in the course of its research that often the CPO is promoted from within in the midst of handling the breach. It is therefore not surprising that, during interviews with Burton Group, several privacy professionals made statements like, “There was really nothing here when I took the job. I had to build everything.” Bootstrapping a privacy program is hard enough; attempting to do it while handling a privacy event is significantly more difficult.

Point of Contact

An effective privacy program must have at least one dedicated, publicized point of contact for handling privacy inquiries and complaints. Depending on the size and structure of the organization and its privacy team, this point of contact may or may not be the CPO. The point of contact must address both major streams of privacy issues: internal and external. To gather external privacy inquiries, most enterprises use a shared e-mail address that the public can send inquiries to. Similarly, a shared phone number and voicemail account as well as a postal address are often used. In some cases, the enterprise has an established office of the privacy ombudsman. The ombudsman can serve not only as the point of contact for privacy issues, but also as the privacy advocate for customers and partners. The same organizational structures can be applied for internal privacy concerns, though in some cases, instead of using a shared account, the CPO or designee's e-mail and phone number are distributed as internal contact points.

Budget

The small size of privacy teams makes their job difficult; without resources, “difficult” becomes “impossible.” Privacy teams will typically try to pass on as much cost as they can to LOB, IT, and legal departments. That being said, these teams need a budget to carry out their own projects, including:

- Benchmarking and evaluations
- Education and awareness initiatives
- Hiring external consultants and legal experts

A privacy team without its own dedicated budget is unlikely to be effective. The budget need not support a large staff, but it needs to support the team's essential activities.

Overview of the Activities of an Effective Privacy Program

As different sectors and jurisdictional impart privacy requirements, there is no one prescriptive list of activities that a privacy program must execute to be successful. In the course of researching this overview, Burton Group identified the following as a common set of privacy program activities. It is important to note that due to differences in team maturity, industry, and priorities, not all privacy teams Burton Group interviewed were conducting all of these activities to the same extent. The common activities of effective privacy programs include:

- Data inventories and classification
- Risk assessment
- Privacy impact assessments
- Education and awareness
- Privacy audits
- Regulatory awareness

Data Inventory and Classification

While privacy is about people rather than about data, improper use of data about people is a common source of privacy incidents. Not knowing that private data exists, or not knowing where private data is, makes incidents more likely. A privacy team without a data inventory will be destined to wander in a fog, chasing the specter of potential breaches while arriving late to the scene of actual breaches. A data inventory is an invaluable set of information for privacy teams that enables privacy teams to prioritize their resources and efforts, provide risk assessments, evaluate existing policies, and formulate new policies. Building a data inventory takes time and requires effort. This effort can be partially reduced by using tools to automate the data discovery and classification process; for more information on these tools, see the *Security and Risk Management Strategies* overview “[Information Classification: The Most Important Security Thing You're \(Still\) Not Doing](#).” The process of building a data inventory can highlight political issues in the enterprise concerning privacy. It can also serve as input to risk assessments and privacy impact assessments that will define the privacy team's roadmap.

What, Who, and How to Ask

The data inventory must provide enough information to help the team evaluate this risk associated with all of the enterprise's information repositories that contain privacy-sensitive information. Regardless of the purpose, format, and location of privacy-sensitive information repositories, the data inventory must contain them all. Unlike with a census, statistical modeling is not sufficient; the privacy team has to knock on everyone's door and record repositories found.

At the minimum, a data inventory should include the elements illustrated in Table 1.

Element	Purpose
The nature of a repository of privacy-related information	Provides context and describes the purpose of the repository.
The notional owner of the repository	A starting point for further investigation into the repository if needed.
The volume of information in this repository	How much data is actually in the repository?
The format of the information	Is this a paper or electronic repository? Is it structured or unstructured?
The use of the information	How is the information being used?

Type (or types) of privacy-related information in the repository	What kinds of information are in the repository? (E.g., physical or e-mail addresses? SSNs? Health information? Salary information?)
--	--

Table 1: *Typical Components of a Data Inventory*

The overall goal of the data inventory is to help the privacy team build a set of risk-based priorities. To do this, the team needs to understand what information is in the enterprise, how much of it there is, and what requirements it is subject to. The nature of the repository and the use of its information provide the context of the repository; this information can be augmented by the owner of the repository. The volume of information is one of the key elements in evaluating privacy-related risk. Knowing the format of information in a repository will help teams decide what skill sets they'll need (or need to find) to work with the repository. Because different types of information have different privacy requirements, a data inventory needs to gather details about which types are in which repositories. It is highly likely that a single repository will have multiple types of private information and thus will be subject to different sets of requirements.

Beyond the six items described in Table 1, a complete data inventory will also include sources of the information, details about how it is collected and disseminated, and opt-in and consent tracking data as well. Enumerating sources of information and dissemination points aids a privacy team in understanding the flow of information into and out of the enterprise. Onto this data, the team can map jurisdictional and regional privacy requirements and regulations. By using opt-in and consent tracking information, the privacy team, in conjunction with audit teams, can build benchmarks as well as evaluate the efficacy of its privacy controls.

There are two approaches to gathering data inventory information: interviews and surveys. Using the interview method, the privacy team interviews repository owners about the privacy-sensitive information those owners have under their control. This is a time-consuming approach for the privacy team, but it is the most accurate way to gather the information. It requires the cooperation of application and repository owners and their management. Because the data inventory is concerned with all the personal information in the possession of the enterprise, and not just electronic information, building the inventory often requires the assistance of people far removed from IT in the LOB.

A privacy team using the survey method distributes some form of questionnaire throughout the enterprise to gather data inventory information. Although from the perspective of the privacy team, this is the least labor-intensive approach, it may end up being more time consuming than the interview method because the privacy team must wait for the information to be returned. Further, the privacy team must follow up with stragglers. To build a data inventory using the survey method in a timely manner requires strong executive sponsorship and management support to drive the survey process. The survey approach relies on a strong element of trust. The privacy team has to trust that respondents will not omit any repositories of privacy information; it is a lot easier for a respondent to forget about a file cabinet or server locked in the closet when filling out a form than when a representative from the privacy team is in the room asking questions.

Burton Group recommends the following approach in constructing and maintaining a data inventory. First, the privacy team should understand and document the context and domain requirements before starting a new data inventory or updating an existing one. For example, an enterprise may operate in multiple domains, subject to multiple legal regimes. A data inventory project may only focus on one type of data, such as patient health information, financial records, or personal identifiers. Understanding the context helps to focus the effort. Second, the privacy team should consult its legal department before starting, so as to avoid problems later. The legal team may advise the privacy team that if the team discovers a certain set of information, that team will be responsible for using specific procedures to handle that information. The goal here is to avoid having the privacy team discover a stockpile of privacy-related nuclear waste armed only with a set of chopsticks to dispose of it. Third, the privacy team should use the interview method to construct the initial data inventory. This effort can be augmented through the use of data classification tools in some cases (though it's critical to remember that these tools discover only information in electronic form). Updates to the data inventory may be implemented using the survey method. If the enterprise has incorporated some form of [privacy impact assessment](#) into its project management process and software development lifecycle (SDLC), then those assessments can be used to update the data inventory as well.

Although the amount of information contained in a data inventory may seem small, gathering this information can take an inordinate amount of time. For the small to medium-size organizations Burton Group interviewed in the course of researching this report, it usually took a year or more to build the first iteration of a data inventory. The privacy team spent some of this time ferreting out all of the repositories of privacy-related information, and some of this time justifying why it was asking for the information in the first place. Furthermore, in these cases, the privacy team was solely responsible for building the data inventory.

Who Owns the Effort and the Result

The data inventory is institutional knowledge about where sensitive personal information lives in the enterprise. This seemingly implies that the CIO ought to be responsible for maintaining the data inventory. And given that the “I” in “CIO” doesn't stand for “information technology” but rather for “information,” the CIO is thus responsible for the enterprise's information *in all forms* as well as information about that information—such as the data inventory. Sadly, that statement is woefully optimistic to the point of being utopian.

Among the enterprises that Burton Group interviewed while researching this overview, only a few CIOs owned the data inventory. In these few cases, the members of the privacy team (and not the CIOs) were also not held responsible for building the data inventory. In most cases, over time, the ownership and maintenance of the data inventory transitioned from the privacy team to the CIO. Considering that most privacy teams are formed as a result of a breach, the privacy team not only had to contend with the repercussions of the breach (and the associated internal politics) but also had to conduct an enterprise-wide data inventory. Given the reality that privacy teams are often very small, the combination of all these factors make building that first data inventory extremely difficult.

The data inventory is a highly valuable and sensitive enterprise asset; it can be valuable to organizations other than just the privacy team. It warrants executive-level support and management. An enterprise privacy team is not the appropriate team to maintain such an asset. The team must be involved in formulating the questions that the data inventory needs to answer, and it can help elicit those answers, but it should not be solely responsible for the ongoing maintenance of the data inventory. Maintenance should be the responsibility of the CIO.

Risk Assessment

In the hands of the privacy team, the data inventory can become a powerful risk assessment tool, and this risk assessment can help set priorities for privacy initiatives as well as information security programs. Although all the elements of a data inventory are useful in analyzing privacy-related risk, evaluating just the volume and type of private information in a repository is sufficient for a “thumb in the air” estimation of risk and thus can help set priorities for the privacy team. For example:

- The warehouse full of tens of thousands of paper patient records deserves more attention than the billing systems with thousands of patients.
- The donation tracking system containing the donation history of thousands of alumni and donors warrants more attention than the athletic ineligibility spreadsheet listing a dozen students.
- The system that tracks passport use for millions of citizens requires more attention than the system that tracks passport applications.
- The system that stores SSNs, account numbers, and credit card data for a bank's customers requires more attention than the mailing list the marketing department maintains to recruit new customers.

Information about data type and volume, while useful by itself, is even more powerful when considered with the other elements of the data inventory. Understanding where the private information lives (e.g., on a thumb drive, in a protected data center, in a filing cabinet down the hall) and how that information is used (printed out and interoffice-mailed to each branch, published to a website that doesn't require authentication, stored under lock and key for regulatory purposes and then destroyed after seven years) enrich risk assessments. The reality is that IT, records management, and privacy teams are all strained for resources. While multiple privacy-related risks may exist in the enterprise, not every one of them can be addressed at the same time. Using the data inventory, a privacy team can set and justify its priorities; these priorities scope the effort for not only the privacy team but also for IT security, records management, and other constituents.

Policy Management

The primary goals of a privacy team should be setting privacy-related policies for the enterprise and measuring adherence to those policies. These policies will:

- Publicize the privacy principles of the enterprise
- Raise privacy awareness
- Protect enterprise brand value
- Reduce the risk of disclosure
- Reduce the amount of personally identifiable information (PII) and sensitive health information (SHI) collected

It is important to note that privacy policies will not and cannot reduce the risk of a breach. Preventing a malicious outsider from gaining access, both physical and logical, is the realm of IT and physical security.

Privacy teams must work with information security teams to build and maintain policies that affect the entire lifecycle of information for all formats and types of information in the enterprise. This lifecycle includes:

- Information collection and creation
- Use
- Revision and correction
- Storage
- Sharing and secondary use
- Destruction

Each of the above points in the lifecycle has a rich set of sub-elements. For example, storage policies should also include policies about backup copies of the information, both on- and off-site if applicable.

Sharing and secondary use are gremlins in a privacy team's life. Policies related to sharing of information should include guidance on the printing, transferring, and transporting of information. Examples of questions that need to be answered by sharing-related privacy policies include:

- What information are employees not allowed to store locally on their systems?
- If employees are traveling abroad, what information should be removed from their systems and/or external storage devices if it is present?
- What set of information, if any, should not be accessible via remote access from employees' home computers?
- What are permitted uses of universal serial bus (USB) and other external storage media, and what are the exceptions?
- What obligations does the organization impose on partners with whom information is shared?

To cover as many sharing scenarios as possible, the privacy team must collect use case information from program and project leads as well as users (that is, the real consumers and users of the information) and data subjects.

Secondary use of information is harder to discover, manage, and regulate than primary use. Secondary use of information can include testing new IT systems that use personal information, marketing campaigns, and aggregation and data mining applications. Recognizing that privacy teams will not be able to predict all of the future secondary uses of information, privacy policies must be revisited periodically to ensure that they stay current with secondary uses.

As mentioned above in this section, privacy policies must cover all types and formats of potentially private information in the enterprise. This includes information:

- On paper
- In electronic formats—both structured and unstructured
- On laptops
- On home computers using remote access
- On portable storage devices, including USB drives, iPods, CD-ROMs, and DVDs
- In backup systems—both on and off site
- In storage area networks
- In departmental SharePoint servers

There are hundreds, if not thousands, of primary use, secondary use, type, and format permutations for privacy-related information. Furthermore, each jurisdictional and sectoral domain imposes its own requirements. Add societal norms and the starting assumptions of customers, partners, and regulators, and it is easy to see why privacy teams can easily get overwhelmed when trying to start building policy. To avoid this, privacy teams need to let their risk assessment of the data inventory be their guide. Privacy teams are almost universally small, and all rely on other groups within the enterprise for help. In this resource-bound situation, a privacy team cannot possibly build policies that address every permutation: every repository, every use, and every secondary use. A risk-based approach keeps privacy teams focused on the most critical tasks while outlining future goals.

A final note is needed about privacy policies and working with information security. Privacy teams must form a strong bond with information and IT security teams in the enterprise. It is only through a strong partnership that privacy policies can be enacted in IT. Furthermore, most privacy teams are composed of lawyers who often do not speak in technological terms; IT must work with the privacy team to translate privacy requirements into meaningful IT policies and procedures.

Privacy Impact Assessments

As new (or major updates to existing) projects and programs are being designed and planned, privacy teams have an opportunity to head off potential privacy problems by conducting privacy impact assessments (PIAs). A PIA examines a project or system to determine what, if any, privacy implications the system may have. By conducting a PIA, a privacy team will learn, at a high level:

- What, if any, privacy-related data is being collected?
- Where is this data coming from and how is it being collected?
- Why is this data being collected?
- How and by whom will this data be used?
- How and with whom will this data be shared?
- How long will this data be stored?

PIAs are commonly conducted by agencies of government; the questions that PIAs raise are broadly applicable and can serve all sectors and domains.

A PIA should not be used as an audit, in and of itself, but as a design tool. By raising privacy-related questions during the design of a project, before systems have been implemented and before data has been collected, privacy teams can work with project managers, process designers, and architects to build privacy-sensitive systems. It is not just new projects and systems that can (and should) benefit from PIAs, but also existing ones. A PIA should be conducted when a system is being upgraded or when a new phase of a project is being designed. To that end, some enterprises Burton Group interviewed for this overview have begun to integrate a PIA phase into their SDLC processes.

In enterprises that respect privacy and the dignity of their customers and citizens, underlying the questions in a PIA is the desire for minimization. This desire raises questions such as:

- Can the application collect less information?
- Does the system have to store that data in an identifiable form for so long?

The desire for minimization may read like a privacy activist's hymnal, but this desire actually expresses the risk- and cost-reduction needs of the enterprise. Collecting less information can mean lower bandwidth costs. Storing less information means lower storage and backup costs. Not collecting a piece of private information means one less piece of information that could be lost to breach or accidental disclosure—and thus, one less risk and one less cost.

Privacy Audits

Within the enterprise, privacy is a governance matter; as such, privacy needs a quality-control feedback loop for its policies: Implement policies, observe their efficacy, and revise as necessary. Privacy audits are a key part of the feedback loop as they provide the privacy team with an opportunity to observe the efficacy of its policies. As with policy management, privacy teams should take a risk-based approach in setting privacy audit priorities. A small privacy team cannot, in a reasonable amount of time, audit all of the privacy-sensitive systems and repositories in the enterprise; they must focus attention on the riskiest enterprise systems. In most cases, privacy teams do not conduct audits on their own but instead work with internal audit teams to incorporate privacy elements into larger internal audit activities.

Privacy teams also play an indispensable role in preparing the enterprise for an external privacy audit. Burton Group has observed an increasing desire among enterprise privacy teams to ensure that enterprise partners have solid privacy practices. We've seen evidence of this in privacy teams working with procurement and contract management teams to add privacy language into the terms and conditions of enterprise contracts. Further, privacy teams have begun to request privacy audits of partners. Privacy teams can serve as translators helping to bridge communication gaps between external auditors and various enterprise teams.

Education and Awareness

One of the best methods of improving enterprise privacy performance is education. It is not sufficient to include privacy-related instructions and guidance in an employee handbook, though Burton Group does recommend this practice. Successful privacy teams build training materials and deliver training sessions to a variety of types of people, including:

- Executive management
- IT staff
- Internal audit staff
- Administrative staff
- New hires
- Summer interns
- Guest lecturers and researchers
- Professional staff
- On-site contractors and consultants

- Partners

To reinforce training sessions and materials and to serve as a one-stop shop for privacy information, privacy teams need to maintain an enterprise-wide repository of privacy principles and values, policies, practices, and procedures.

One effect of effective privacy education is simply a greater awareness of privacy issues. The awareness that education builds thrives on reinforcement; privacy teams deliver this reinforcement through marketing and PR. Greater privacy awareness can serve to reduce the risk of breach or accidental disclosure.

A Comparison of Privacy Principles

The fact that privacy is contextual prevents attempts to create a meaningful definition of privacy. As mentioned in the “[A Single Strict Definition of Privacy May Not Be Desirable](#)” section of this overview, a single strict definition of privacy is probably neither achievable nor desirable. What is achievable and desirable is a shared set of privacy principles that can guide enterprises and help shape contextual enterprise privacy policies, practices, and culture.

Building a Set of Contextual Privacy Principles

Enterprises should not start from scratch in creating their privacy principles, but instead should use existing principles as the nucleus of their principles. Industries, government agencies, nations, multi-national organizations, and other domains have published privacy principles from which enterprises can draw. Within limits, the cores of these sets of principles are similar. That being said, some sets of privacy principles contain domain-specific principles and/or domain-specific renditions of shared principles. Enterprises should follow this pattern:

- 1 Base privacy principles on common shared principles
- 2 Interpret principles in light of the domains relevant to the organization; add domain specificity to the language of common shared principles as needed
- 3 Add domain-specific principles only when absolutely necessary

Normalized Privacy Principles

In order to illustrate the commonalities of some of the many publicly available sets of privacy principles, Table 2 maps Burton Group's common shared privacy principles (enumerated in the “Analysis” section of this overview) to corresponding principles from a variety of sets of privacy principles. The sets of privacy principles included are:

- The American Institute of Certified Public Accountants (AICPA)—Generally Accepted Privacy Principles⁶
- Organisation for Economic Co-operation and Development (OECD)—Guidelines on the Protection of Privacy and Transborder Flows of Personal Data⁷
- European Union (EU) Directive 95/46/EC on the protection of individuals with regard to the processing of personal data and on the free movement of such data; also known as “the EU Data Directive”⁸
- Canada's Personal Information Protection and Electronic Documents Act (PIPEDA)⁹
- U.S. Safe Harbor¹⁰
- Asia-Pacific Economic Cooperation (APEC)—Privacy Framework¹¹
- Network Advertising Initiative (NAI) Self-Regulatory Code of Conduct¹²
- U.S. Department of Homeland Security (DHS)—The Fair Information Practice Principles: Framework for Privacy Policy at the Department of Homeland Security¹³

- U.S. Department of Health and Human Services (HHS)—The Nationwide Privacy and Security Framework for Electronic Exchange of Individually Identifiable Health Information¹⁴

Principle/ source	OECD	AICPA	EU Directive	PIPEDA	U.S. Safe Harbor	APEC	NAI	DHS
Governance		Management		Accountability				
Accountability for all parties	Accountability	Monitoring and enforcement	Data subject's right to object Judicial remedies, liability, and sanctions Transfer of personal data to third countries	Accountability Challenging compliance	Enforcement Onward transfer	Accountability		Accountability and auditing
Transparency	Openness	Notice	Information to be given to the data subject	Openness Identifying purposes	Notice	Notice	Notice Transparency	Transparency Purpose specification
Meaningful choice	Collection limitation	Choice and consent	Criteria for making data processing legitimate	Consent	Choice	Choice	Choice Transparency	Individual participation
Minimal collection and disclosure	Collection limitation	Collection	Special categories of processing	Limiting collection Limited use, disclosure, and retention		Collection limitation		Data minimization
Constrained use	Use limitation Purpose specification	Use and retention	Special categories of processing	Limited use, disclosure, and retention		Use of personal information	Use limitations	Use limitation
Data quality and accuracy	Data quality	Quality	Principles relating to data quality	Accuracy	Data integrity	Integrity of personal information	Reliable sources	Data quality and integrity
Validated access	Individual participation	Access	Data subject's right of access to data	Individual access	Access	Access and correction	Access	Individual participation
Security	Security safeguards	Security for privacy	Confidentiality and security of processing	Safeguards	Security	Security safeguards	Security	Security Data quality and integrity

Table 2: Privacy Principles Compared

Conclusion

Because privacy is fundamentally contextual, a privacy team's most important work is to contextualize common privacy principles into useable, meaningful privacy practices in the contexts relevant to the organization. Doing this requires collaboration with partners throughout the enterprise, strong executive support, a budget, and a formal privacy governance structure. Success in this effort builds an organization that respects the dignity of customers, employees, citizens, and partners; forms stronger relationships; and reduces risks to its finances and reputation.

Notes

- ¹ Ludwig Edelstein. *The Hippocratic Oath: Text, Translation, and Interpretation*. Baltimore, Maryland: Johns Hopkins Press, 1943.
- ² Edward Hanna. "The Sacrament of Penance." *The Catholic Encyclopedia*. Vol. 11. New York: Robert Appleton Company, 1911. 27 Jan 2009. <http://www.newadvent.org/cathen/11618c.htm>.
- ³ Former Republican National Committee chairman Mike Duncan as reported by *The Washington Post*. 5 Jan 2009. http://www.washingtonpost.com/wp-dyn/content/article/2009/01/05/AR2009010502771_2.html.
- ⁴ Monica Chew, Dirk Balfanz, and Ben Laurie. "(Under)mining Privacy in Social Networks." *Google*. 2 May 2008. <http://w2spconf.com/2008/papers/s3p2.pdf>.
- ⁵ John Locke. Second Treatise of Government §27, at 19(1980)(1690).
- ⁶ "Generally Accepted Privacy Principles" *The American Institute of Certified Public Accountants (AICPA)*. <http://infotech.aicpa.org/Resources/Privacy/Generally+Accepted+Privacy+Principles/>.
- ⁷ "OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data." *Organisation for Economic Co-operation and Development (OECD)*. http://www.oecd.org/document/18/0,3343,en_2649_34255_1815186_1_1_1_1,00.html.
- ⁸ "Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data." *Official Journal of the European Communities*. http://ec.europa.eu/justice_home/fsj/privacy/docs/95-46-ce/dir1995-46_part1_en.pdf.
- ⁹ "The Personal Information Protection and Electronic Documents Act." *Office of the Privacy Commissioner of Canada*. 13 Apr 2000. http://www.privcom.gc.ca/legislation/02_06_01_e.asp.
- ¹⁰ U.S. Safe Harbor. <http://www.export.gov/safeharbor/>.
- ¹¹ "APEC Privacy Framework." *Asia-Pacific Economic Cooperation*. 2005. http://www.apec.org/apec/news_media/fact_sheets/apec_privacy_framework_MedialibDownload.v1.html?url=/etc/medialib/apec_media_library/downloads/taskforce/ecsg/pubs/2005.Par.0001.File.v1.1.
- ¹² "2008 NAI Principles: The Network Advertising Initiative's Self-Regulatory Code of Conduct." *Network Advertising Initiative*. 2008. http://www.networkadvertising.org/networks/2008%20NAI%20Principles_final%20for%20Website.pdf.
- ¹³ Privacy Policy Guidance Memorandum. *U.S. Department of Homeland Security*. 29 Dec 2008. http://www.dhs.gov/xlibrary/assets/privacy/privacy_policyguide_2008-01.pdf.
- ¹⁴ "Nationwide Privacy and Security Framework for Electronic Exchange of Individually Identifiable Health Information." *Office of the National Coordinator for Health Information Technology, U.S. Department of Health and Human Services*. 15 Dec 2008. http://www.hhs.gov/healthit/documents/NationwidePS_Framework.pdf.

Author Bio

Ian Glazer

Senior Analyst

Emphasis: provisioning, privacy, identity audit, controls management

Background: Ian Glazer is a senior analyst for Burton Group's Identity and Privacy Strategies service. He covers identity audit, user provisioning, controls management, and privacy. Prior to joining Burton Group, Ian was senior director, of program management at Approva Corporation, director of identity strategy at Trusted Network Technologies, and senior product manager at IBM where he was a top-ranked product manager on the IBM Tivoli Identity Manager team, heading provisioning offerings for small and medium businesses. Ian is a strong advocate for industry standards and efficacy. He was a contributor to OASIS Provisioning Services Technical Committee and is a co-inventor of the patent pending Web Services Federated Provisioning. Ian is a frequent speaker and panelist at identity leadership events and is an active blogger identity management and security issues.

Copyright 2009 Burton Group. ISSN 1048-4620. All rights reserved. All product, technology and service names are trademarks or service marks of their respected owners. See Terms of Use and publishing information at <http://www.burtongroup.com/AboutUs/TermsOfUse.aspx>