Privacy Protection Measures and Technologies in Business Organizations:

Aspects and Standards

George O.M. Yee
Aptus Research Solutions Inc., Canada & Carleton University, Canada

Information Science
REFERENCE

Managing Director: Lindsay Johnston
Senior Editorial Director: Heather Probst
Book Production Manager: Sean Woznicki
Development Manager: Joel Gamon
Development Editor: Myla Harty
Acquisitions Editor: Erika Gallagher
Typesetters: Milan Vracarich, Jr.
Print Coordinator: Jamie Snavely
Cover Design: Nick Newcomer, Greg Snader

Published in the United States of America by
Information Science Reference (an imprint of IGI Global)
701 E. Chocolate Avenue
Hershey PA 17033
Tel: 717-533-8845
Fax: 717-533-8661
E-mail: cust@igi-global.com
Web site: http://www.igi-global.com

Library of Congress Cataloging-in-Publication Data

Privacy protection measures and technologies in business organizations: aspects and standards / George O.M. Yee and Aptus Research Solutions Inc., editors.
p. cm.
Includes bibliographical references and index.
Summary: “This book is a collection of research on privacy protection technologies and their application in business organizations”--Provided by publisher.
ISBN 978-1-61350-501-4 (hbk.) -- ISBN 978-1-61350-502-1 (ebook) -- ISBN 978-1-61350-503-8 (print & perpetual access) 1. Business--Data processing--Security measures. 2. Computer security. 3. Data protection. 4. Privacy, Right of. I. Yee, George. II. Aptus Research Solutions.
HF5548.37.P755 2012
658.4’78--dc23

2011038433

British Cataloguing in Publication Data
A Cataloguing in Publication record for this book is available from the British Library.

Chapter 1
Privacy Enhancing Technologies for Information Control

Martin Gilje Jaatun
SINTEF ICT, Norway

Inger Anne Tøndel
SINTEF ICT, Norway

Karin Bernsmed
SINTEF ICT, Norway

Åsmund Ahlmann Nyre
SINTEF ICT, Norway

ABSTRACT

Privacy Enhancing Technologies (PETs) help to protect the personal information of users. This chapter will discuss challenges and opportunities of PETs in a business context, and present examples of currently available PETs. We will further study the Platform for Privacy Preferences (P3P), and discuss why it so far has failed to deliver on its promise. Finally, we provide our advice on further research on privacy preferences, and conclude with our conviction that businesses need to take a progressive stance on providing privacy to their customers.

INTRODUCTION

Privacy is a fuzzy concept with many definitions, one of which is "the right to be let alone" (Warren & Brandeis, 1890). This particular definition seems to have lost much of its validity, however, as we in modern society base so much of our existence on interaction with others over the internet. This is in particular true for us as consumers, as increasingly businesses assume that we will use the internet for both purchases and user support.

Privacy Enhancing Technologies (PETs) comprise a broad collection of tools and processes that help protect the privacy of end-users. PETs range from mechanisms that prevent disclosure of personal information, via ways of hiding location information, to methods for anonymous communication.

To the casual observer it might seem that most businesses have been more interested in privacy-

DOI: 10.4018/978-1-61350-501-4.ch001

invasive technologies than privacy-enhancing technologies, in that businesses have wanted to learn as much as possible about their (potential) customers, in order to deploy more targeted advertisements, sales projections and production planning. On the "dark side" of the business spectrum, this has led to the development of spyware that monitors individual computer users, but retailers have also looked into the possibility of exploiting RFID tags on merchandise to harvest information on their customers (CASPIAN, 2006).

To this day, most consumers seem oblivious to privacy concerns, and their behavior is usually not motivated by such concerns; to the contrary, most people seem overly free with their personal information, particularly in the context of online social networks such as Facebook (2011) and LinkedIn (2011). However, as advances in data mining techniques are progressing, the vast amounts of data available to businesses are bound to be recognized as a concern by consumers, and a backlash is imminent.

It has been said that online businesses instead of establishing a trust relationship with their customers, rather focus on avoiding *distrust* (Clarke, 2008), and that this is mainly achieved through how they treat the personal data of their customers – in many cases, the more personal data a business demands from a customer, the more the distrust increases. In this context, it would seem that organizations that facilitate the use of Privacy Enhancing Technologies in their interactions with their customers should have a business advantage. Many businesses tend to collect personal information about their customers as a matter of course, regardless of whether they actually need this information. This could potentially end up as a liability for a business, since in many jurisdictions, storing personal information requires informed consent, and forces the business to abide by privacy legislation.

The obligation to inform users of privacy practices is commonly resolved by using a comprehensive and high-level description of an organization's privacy policy (Guarda, 2009). In the digital world, privacy policies have become the main instrument for service providers to explain how users' personal data are collected, used, disclosed, and managed. Unfortunately, due to their complexity, difficult language and sheer length, users tend to neither read nor understand the policies prior to acceptance (Berendt, Günther, & Spiekermann, 2005). Vila et al. (2003) show that market forces are actually counter-productive to use of privacy policies when effort is required by users to verify a policy. This is an argument for more automated processes.

As the 20th century was drawing to a close, the Platform for Privacy Preferences (P3P) (W3C, 2006) emerged as an innovative privacy-enhancing concept, based on the transformation of textual privacy policies into machine-readable instructions for computers. A main motivation for the P3P project was to make it easier for users to understand privacy policies and make well-informed decisions on how to interact with services that collect personal data (Argyrakis, Gritzalis, & Kioulafas, 2003). Central to its vision was the privacy agents that allowed the user to specify what was acceptable and not, in terms of information sharing, and let the agent compare the user's privacy restrictions with the intentions of the web site that he was visiting. P3P is considered one of the most significant efforts to help web users control the sharing of their personal information. Later on in this chapter we will discuss its background, history and uptake, criticisms and technical obstacles, and explain the main reasons for its failure.

Strictly speaking, one may be reluctant to categorize approaches such as P3P as PETs, since they are more concerned with informing users about how their personal information will be (ab)used than actually protecting the users (often against themselves). It is important to consider P3P as a concept, however, as it primarily has relied on individual websites (i.e. businesses) for deployment, and the current verdict is that for the most part, this must be considered a dismal failure.

This chapter thus explores the failure of P3P based on a motivation to determine how future PET deployments for business organizations may avoid the same fate.

In this chapter we will explore Privacy Enhancing Technologies (PETs) that focus on active information gathering (Lioudakis, et al., 2007). We thus do not consider passive or semi-passive information gathering, which in turn means that mechanisms such as onion routing or other anonymisation efforts are beyond the scope of this chapter. For most users, privacy is really about *information control*, since we intuitively want to restrict who has access to our personal information. In practice, however, users find it difficult to live by the privacy rules they make for themselves, and PETs represent one way of balancing the scales in this respect.

Note also that this chapter takes a business perspective. The objective of the chapter is to present Privacy Enhancing Technologies that are suitable for deployment by businesses for the protection of their customers' privacy, and outline how PETs could be wielded as a competitive advantage by discerning businesses. Though we acknowledge that businesses are of various types and have varying needs for personal information, we do not go into these differences in this chapter. This also means that we do not put particular emphasis on, e.g., the special needs of social network providers. In most cases, we will refer to the business as a *service provider* and the customer as a *user*, but sometimes we will explicitly highlight the business/customer relationship.

The chapter is structured as follows. First, the chapter provides background information on the value or personal information and on the privacy regulation approaches that are in use today. Then, it describes the market opportunities of PETs, addressing both the potential problems that can be encountered as well as the opportunities that can be realized. An overview is provided of some existing PETs as well as research prototypes, followed by a discussion of P3P and the reasons for its failure. Based on the business needs and corresponding problems of P3P that are identified, we point at some important areas for future research.

BACKGROUND

In this section we have a look at some of the important characteristics of privacy and personal information as seen by different stakeholders; the businesses, the users and the society. First, we outline the main rationale for businesses to collect personal information from users. Next, we identify the motivation for users to share their personal information, before we give an overview of society's view on privacy through providing a summary of regulatory efforts to protect privacy.

The Value of Personal Information

Odlyzko (2003) argues that the main motivation for businesses to use personal information is to perform effective price discrimination. The idea is that the more the business knows about their consumers, the more accurately they can determine the price they are willing to pay. Hence, consumers willing to pay more are also charged more than other consumers for the exact same product, such that the charge on the entire customer base reaches its potential maximum. In an ideal world, everyone would be satisfied with price discrimination, since the idea is that you are charged the amount you are willing to pay for a service. However, it is a dangerous strategy that is likely to cause outrage if the discrimination is too blatant (Acquisti, 2002).

A variant of this motivation is the targeted offerings or advertisements that seem to interest businesses. Based on purchase history, ratings, demographic information and personal likings, potential customers may be offered to buy products that are likely to be interesting to them (Acquisti, 2002). Despite the obvious benefits for both users and service providers; e.g. users are to a greater extent only presented with advertisements for

products they are interested in, and advertisers experience greater impact of each presented advertisement; customers tend to disapprove of targeted advertising (CASPIAN, 2006). This may possibly be due to the discomfort users are experiencing when realizing the amount of information available on them. As we will see in the next section, people tend to not act rationally when it comes to privacy.

Most businesses will claim that their use of personal information is mainly motivated by improved customer experience. The ability to adapt to the various needs of consumers is fundamental to provide value-added features such as auto completion, location-based services, translations, etc, which in turn require collection, storage and use of personal information. One may of course argue that the underlying motivation is to increase profit, not to improve customer experience. However, as with conventional retailers, online businesses have a motivation for keeping their customers happy.

Using consumer information for these purposes is not new; it has been done for quite some years already. However, previously such information was used in a generalized manner mostly for market segments, geographic regions, or other particular user groups. A common example is the use of student discounts, or regional special offers or targeted advertising for readers of a magazine, issues that normally not constitute a threat to privacy. The difference today is that the information is not generalized in any way, instead it is used on an individual basis, not targeting groups of users but rather specific users. Evidently, this new approach to an old phenomenon will introduce added challenges.

Users' Attitude

Users are generally concerned about their online privacy and therefore initially reluctant to share personal information. However, research shows that this is not reflected through their actions, resulting in what is termed the *privacy paradox*; user behavior does not match their stated preferences (Berendt, et al., 2005; Jensen, Potts, & Jensen, 2005). There may be various reasons for the privacy paradox, one of them is that users do not care about privacy after all, another that users lack the knowledge and understanding needed (Flinn & Lumsden, 2005) to make privacy decisions that are according to their principles. Studies also suggest that users have a tendency of forgetting their privacy principles once interacting with services on the web (Spiekermann, Grossklags, & Berendt, 2001).

Privacy Regulations

Although privacy by many is regarded as a fundamental human right, it is not something you can expect as an automatic feature. In the following we will discuss how privacy is regulated through self-regulation, third parties and legislation.

Self-Regulation

Traditionally, industry players in the U.S. have been wary of new government rules and regulations, and this has also had consequences for American privacy legislation (or the lack of same). In the U.S., privacy is assumed to be handled through self-regulation. That is, the market mechanisms (buyer and supplier) should ensure a proper level of privacy protection for end-users. The guiding principles are set forth by the Federal Trade Commission (FTC) and labeled the *fair information principles* (Vila, et al., 2003). These principles comprise:

- **Notice/awareness:** Consumers should be given notice of an entity's information practices before any personal information is collected from them.
- **Choice/consent:** Giving consumers options as to how any personal data collected from them may be used.

- **Access/participation:** An individual has the right to view data collected about him/her and to contest the accuracy of that data.
- **Integrity/security:** Data must be accurate and protected against loss or unauthorized access.
- **Enforcement/redress:** Effective privacy protection requires some means to ensure that these principles are applied.

The self-regulation approach encourages the use of third-party providers to mediate and ensure privacy.

Third Parties

Privacy seals have emerged as means to convey trust between service providers and users. The idea is that a trusted third party can certify that a service provider follows a set of privacy requirements, which subsequently should induce trust in the service provider. For this setup to work, the user must first trust the third party certifier, and second acknowledge the requirements it has set forth. The first assumption may turn out to be problematic since most of these third parties are relatively new to users, and furthermore they always require a fee for issuing a seal. Clearly, there may thus be a conflict of interest when judging whether a site should be accepted or not, since rejected candidates do not generate revenue for the third party. The second assumption is perhaps more subtle, since the requirements set forth by the third party may be inappropriate, insufficient or ineffective in capturing what assurance the user actually needs. If the "privacy approved"-seal only verifies that the site has a privacy policy that contains all the required parts, a user may falsely assume that the actual policy has been approved. So, if the user has substantially different privacy requirements than the third party, the seal does not convey any particular trust. This assumption is often also misunderstood for other security evaluations, such as the Evaluation Assurance Levels (EALs) of the Common Criteria (Information technology -- Security techniques -- Evaluation criteria for IT security -- Part 1: Introduction and general model 2009). The assurance (expressed through seven EALs) is only concerned with verifying a match between the requirements for the product (as given by the developer) and the actual implemented product (as provided by the developer). So in essence, a Common Criteria EAL states how sure you can be that the product does what it says it does, and is not a ranking of security products (such that EAL4 products necessarily are better than EAL3 products).

The TRUSTe privacy seal (Benassi, 1999; TRUSTe) is one of the prominent examples of such privacy seal programs. TRUSTe uses its own set of requirements that are based on the Federal Trade Commission's Fair Information Principles. Although there are some specific requirements (e.g., Secure Socket Layer (SSL) required for transferring personal information, and presence of dispute resolution), most of them are more focused on the presence of certain aspects of the privacy policy and adherence to these aspects. Hence, the seal does not imply any qualitative assessment of the policy itself, but rather that the policy contains all the parts it should. A more thorough approach is the European Privacy Seal, or EuroPriSe (2010), which covers certification of all computerized systems and products. The European Privacy Seal indicates that the product (or web site) has been certified to adhere to European regulations on privacy protection (EU, 2002). Therefore, the requirements imposed on the web site operator are much more stringent than those of TRUSTe. The process to obtain a Seal resembles that of a Common Criteria evaluation for security products, and is also similar in choices of words. Assurance is only given in binary form, either the product is certified or it is not. Currently, there are only a handful of services that have received the EuroPriSe, presumably due to its extensive nature.

Legislation

Meeting government requirements is central to any service provider on the Internet, including with respect to privacy. The European Union member states have ratified the directive on protection of personal information (EU, 2002) and thereby committed themselves to laws on privacy protection that are at least as strict as the directive. By adopting and extending the fair information principals, the European privacy legislation is among the strictest. Particularly, the directive states detailed requirements on the third country legislation required for it to allow transfer of personal information to a company governed by that legislation. In fact, this initially excluded U.S. companies from exchanging personal information with European partners, but was solved later on by the introduction of the *safe harbor program* (US, 2011), a program designed to let U.S. companies voluntarily subject themselves to terms and conditions (not a law as such) laid down by the commission on how to handle personal information.

While it may seem that the legislative approach should require anyone to implement appropriate privacy measures, the fundamental problem is the lack of enforcement. There are few, if any, being prosecuted for not adhering to the statutory demands of the law. Although the law should protect users' privacy, it may seem that it is not capable of such. One reason may be disputes regarding which country's rules actually govern the service.

Whether to address the business motivation, user motivation, legislative demands or self-regulation, it is evident that all parties' could benefit from Privacy Enhancing Technology.

MARKET OPPORTUNITIES FOR PETS

In this section we discuss the main market opportunities of supporting PETs, and also the main reasons for not doing so. We focus on five different activity goals where privacy can be used as an opportunity: Building of reputation and trust, keeping customers happy, establishing a competitive advantage, reducing risk of privacy compromises, and making users share more information. For each of the activity goals we describe problems and opportunities, Problems constitute arguments for not investing in PETs and can e.g. describe uncertainties regarding the benefits that will be achieved or potential obstacles businesses should be aware off. Opportunities describe positive gains that can be achieved if managing to use PETs effectively.

Reputation and Trust

Businesses are expected to show corporate responsibility in a number of areas, including privacy. In addition businesses are dependent on customer trust, and privacy can be one important piece that needs to be addressed in order to gain more trust. Trust is however characterized by the long time it takes to earn it and the short time in which it can be broken. Loss of business and damage to brand has been suggested as one of the prominent risks of a privacy breach (Borking, 2009).

Trust is a very complex term that can take on very different meanings in various contexts. Here, we use trust to denote the level of confidence one part has that the other part will behave as expected. More specifically, how can a user be sure that the service provider actually lives up to its stated policy?

Problem – limited risk of getting caught: On the one hand, there is the risk of a privacy breach, and on the other, the risk of getting caught. Nehf (2007) points at two fundamental accountability problems: First, users seldom know about privacy breaches that occur, and second, if breaches are detected it may be near impossible to trace them to a particular source. In general, the majority of data collection occurs outside public view. Personal information is present in countless databases, and it is difficult for users to have any

overview of how this information is shared and sold, as this is often not detailed in the privacy policy (Gomez, Pinnick, & Soltani, 2009). When users experience harm because of data breaches, they are therefore in most cases not able to put the blame on one specific business.

Problem – consequence of privacy breach may increase: To benefit from PET investments, businesses will likely rely on an increased focus on privacy among their (potential) customers. Such an increased focus on privacy will, for most companies, be an opportunity but at the same time a bit risky. Odlyzko claims that when it comes to privacy, companies "tend to hide what they do, most likely because of the implicitly understood fear of consumer backlash, and potential government intervention." (Odlyzko, 2007) With increased focus on privacy comes a possible increase in the risk related to privacy breaches, as the consequences of a breach may increase. Promoting the use of PETs does not remove the risk of privacy breaches for own business. In the event that companies only seemingly promote trust, there may be a tremendous penalty for not following up.

Problem – costs vs. benefits: If businesses are to succeed in using privacy to build trust and reputation, they must expect to make investments both to improve their privacy (e.g. invest in PETs) but also to get the message through. The costs to do this are certain. The benefits are however not quite clear as the relation between privacy and trust is not yet fully understood. Industry often promotes the model of self-regulation for privacy. This model builds on the assumption that if consumers really care about privacy they will be offended by privacy invasive practices. As a consequence, companies having such practices will suffer a repudiation loss. When this has not happened so far, businesses have reason to believe that privacy is not that important anyway.

Opportunity – signals that there is nothing to hide: In terms of trust building, setting the stage for privacy can have a huge impact on consumers' trust in a business. By promoting privacy enhancing technology, a company is signaling that it has nothing to hide, which in turn may build consumer trust significantly faster than through positive experience. Since PETs are not normally supported by today's businesses, those that choose to support PETs have a better chance of being perceived as someone who is leading the way when it comes to privacy.

Opportunity – more privacy-aware future users: A lack of actual privacy concern is not the only possible explanation for the limited consumer demand for privacy. As also explained in the background section, several studies have shown that when it comes to privacy, users in generally do not act according to their principles (Ackerman, Cranor, & Reagle, 1999; Berendt, et al., 2005; Jensen, et al., 2005). The model of self-regulation, on the other hand, assumes that they do (Greenstadt & Smith, 2005). And although today's users' behavior related to privacy is best described by the privacy paradox, this may not be the case for future users. As users share more and more data, the risk of privacy breaches increases, and with it the potential media attention related to privacy. As users become more aware of the potential privacy problems of sharing their information with a variety of services, they may raise their demands for privacy. If the future brings an increasing awareness of privacy issues, businesses will benefit from being ready when the shift occurs.

Keeping Customers Happy

For any business it will be important to take good care of existing customers. An important part of this is to handle their customers' personal information in a responsible way.

Problem – future users may be less privacy concerned: As explained above, future users may demand better privacy. But there is also another possibility; users may simply become less privacy concerned. Currently, the popularity of blogging and social media services more

generally suggest that people are freely sharing a lot of information about themselves. When so much personal information is available on the Internet anyway, then why should companies invest in privacy protection? As pointed out by Odlyzko: "privacy technologies languish because practically no one bothers to use them, and often intimate information is given out for very little or no monetary reward."(Odlyzko, 2007) Privacy concerns are in many cases also associated with conspiracy theorists and those with "something to hide" (Wilton, 2009).

Problem – whether improved privacy will be appreciated: Businesses are unlikely to have knowledge of whether their customers will appreciate improved privacy and support for PETs. Existing customers have already some minimum level of trust in the company, and are willing to use its services with the current trust level. They have directly or indirectly accepted the current privacy policy, potentially also stating that the privacy policy can be changed at any time without notification. If positive changes in privacy management are to have any effect on the existing customers, they must be made aware of the changes. Though many customers will probably appreciate the improved practices, some may find it annoying and others may become suspicious – why are they introducing this now, have they had a privacy breach that we are not aware of...? In addition there is the risk that increased focus on privacy can make current (unfortunate) practices more easily understandable, and therefore easier to react to.

One of the main problems with privacy is that the users in fact have no means to verify or enforce any of the agreed policies. Hence the service provider has no incentives of making privacy understandable and improving their privacy, other than avoiding a publicly known privacy breach.

Opportunity – meet legislative requirements: Though users may not care about their privacy, legislation will in many cases put requirements on how personal information should be treated. These requirements need to be addressed independently of whether the customers care or not. PETs can help in addressing some of the legislative requirements (Szeto & Miri, 2007).

Opportunity – signal that the business cares for their customers' well-being: Businesses have no reason to simply assume their customers do not care about privacy because they do not experience strong demands for improved privacy. A large amount of users claim to be concerned about their privacy, though their current actions do not necessarily imply so. Convincing current users that their privacy is taken well care of may pay off, as it shows customers that the business cares about their well-being. Customers wishing to get services from a competitor can also become more aware of the privacy implications of doing so (since users then have to share information one more time).

Competitive Advantage

If users take privacy into account when choosing services, privacy can be seen as a competitive factor where businesses with good privacy practices get a market advantage.

Problem – investments needed: Any business that decides to use PETs will need to spend money and resources on new technology and/or modification of existing systems. In addition, internal processes will probably also need to be updated. To be able to do this, businesses will need access to necessary expertise, not only on the technology itself, but also on legal and organizational issues. As PETs are not that commonly used today, there is a lack of experience and research that can give concrete guidance to businesses that consider implementing PETs. Businesses are unlikely to have the necessary competence in-house, and it may also be scarce among external consultants (Borking, 2009). It is also worth noting that many PETs, e.g. user agents such as AT&T Privacy Bird (introduced later in this chapter), will only be of

interest to users when they are supported by a relatively large number of providers (Wilton, 2009).

Problem – visibility of privacy: Currently we do not see a lot of competition or differentiation among organizations when it comes to privacy protection. Use of PETs does not result in social recognition. Usually privacy is mainly seen as a negative driver, where avoidance of privacy breaches is the most important. It is not common to view privacy as a positive driver that can result in market advantage (Borking, 2009). For privacy to become a competitive factor, users need to take privacy into account when selecting services. Currently this is not the case (ref. the privacy paradox). One reason may be the visibility of privacy and PETs, and specifically the availability of privacy information. Companies commonly use privacy policies in order to inform users of their privacy practices. These privacy policies are usually time consuming and difficult to read and understand. In addition they often can be changed without notice and are generally not believed by customers. As a consequence, they are not a good basis for users to make decisions as to whether to share information. It is also important to remember that when users interact with a service online, privacy is not in the forefront of the customers mind – "they're not buying privacy, they're buying some other thing" (Greenstadt & Smith, 2005). As data handling is tightly bundled with some good or service, people tend to forget privacy, and if they manage to keep privacy focus, there are often no good privacy choices available that will allow them to acquire the goods or services they want. As a result, consumers do not know about nor believe in good privacy practices, and businesses as a consequence have no incentive to create any (Greenstadt & Smith, 2005).

Problem – human factors that influence decision making: Businesses wanting to compete on privacy need to take into account how human factors influence decision making. Nehf (2007) describes the decision making goals of users as *accuracy of decisions*, *cognitive ease* and *emotional comfort*. People do not necessarily strive towards optimal solutions. Rather they seek satisfactory solutions that can be reached with a minimum of effort. As a result, some factors may be left out. Since the consequences of sharing information with a particular site are not clear to most users, privacy is a potential candidate for exclusion.

Since users strive for cognitive ease, they prefer alternatives that are easy to evaluate. As a consequence, service providers should not make it too complicated to choose their particular service. Forcing users to delve into privacy policies (e.g. by having them press an "I agree" button) in order to use a service is probably not a good idea, as this will increase the effort required to use the service. Unless the privacy policy is really superior, users will be likely to choose other easier available services instead.

If forced to make difficult trade-offs, users can experience negative emotions. People seek to minimize such discomfort in their selection of decision strategies. As a result important factors may be excluded from the decision base. It is difficult to put a price on privacy or compare its value towards other benefits. Users will also in many cases consider privacy as a protected value that should not be prized and traded away. As a result people may tend to avoid making privacy comparisons.

Nehf also describes other factors that influences decisions:

- *Inferences*: If an important attribute is not easy to evaluate, people may infer a value based on what they already know, e.g. that it is similar across brands.
- *Framing effects:* The way information is presented influences the decisions. In particular, users evaluate the cost of digging deeper into the subject against the cost of accepting it as is.
- *The availability heuristic*: People commonly over-respond to risks that they are

well aware of, and may underestimate risks that do not often come to their attention.

Inferences may result in difficulties of making privacy a competitive factor. Users may simply assume that the privacy policies of similar sites are also similar. Businesses wanting to stand out as strong on privacy will have to put effort into making this message come through. This can however result in increasing the cognitive effort of users and forces them to make more emotionally-laden comparisons. As a result the site may be seen as less appealing. The framing effect makes it possible to look privacy conscious without actually being that. The mere presence of a policy or a privacy seal may give users the impression that privacy is taken care of, and few users will study the details and discover the facts. Because of the availability heuristics users may underestimate privacy risks. Though privacy breaches gain some publicity, there is seldom much focus on what are the consequences experienced by users of the weak privacy practices of providers, and as users have little knowledge about what information is collected and with whom it is shared, it is difficult to couple the risks with actual businesses.

Opportunity – ability to support privacy concerned customers without bothering the other customers: Though there are problems with making privacy a competitive factor, PETs can actually reduce some of the problems. Many current PETs come in the form of user agents that user have to download and install themselves. Businesses in many cases just need to support these PETs, e.g. by making their privacy policy available in a proper machine-readable form. Thus, by publishing machine-readable policies businesses can support those users that are privacy concerned without troubling less interested users with privacy information that forces them to make difficult and emotionally laden trade-offs.

Opportunity – adapt to users' privacy preferences: Businesses also have the potential to offer more advanced privacy services to users. Consider the static nature of current online privacy policies that do not fit the otherwise dynamic world of the Internet. Typically, users are forced to either accept the privacy policy in full or reject it entirely; a situation that rarely is the case offline. This is particularly problematic since most web sites only provide a single privacy policy for potentially a myriad of different services. For service providers, the concern is that potential customers may refrain from using their services due to the stated privacy policy, without the provider even knowing about it. One solution to this problem is privacy policy negotiation, where the service provider and user negotiate and agree on a set of privacy rules governing the use of a particular service. For instance the service provider may provide only limited services to users that do not acknowledge a particular rule in the privacy policy. In the event that the particular rule is mandatory for delivering any service to the user, the service provider at least gains knowledge on the particular rule causing the problem. The irony is that by employing a negotiation strategy, the service provider may actually end up with more information than before. The obvious benefit of such an approach is that customers originally rejecting the privacy policy may now use the services offered by the provider, though with potentially reduced functionality. Additionally, service providers have the option to change their privacy policies (and implementation) based on the rules users are unwilling to accept.

Reduce Risk of Privacy Compromise

Businesses trusted with personal information always face the risk of privacy breaches. PETs can influence this risk, both positively and negatively.

Problem – uncertain costs of privacy breaches: There is currently a lack of empirical data on what privacy breaches cost and how likely they are. It is not possible based on current research and experience to say what damage should be expected in case of a major privacy breach. Research and experience often show conflicting results, and

trust is difficult to verify (Borking, 2009). As mentioned earlier, users experiencing problems that can be related to a privacy breach are often unable to trace the source of the problems. Thus, businesses have a quite good chance of not getting caught (Nehf, 2007). Still, regulations such as the Data Protection Directive in the European Union, and the Health Insurance Portability and Accountability Act (HIPAA) in the United States put requirements on how personal information should be treated. Often such legislation is general and quite abstract, and it can be difficult for organizations to understand what they actually require them to do. They also put no direct pressure to invest in PETs (Borking, 2009), but probably as important, "there is nothing to stop anyone from breaking these laws except the fear of getting caught." (Greenstadt & Smith, 2005) Since there up till now have been relatively few investigations and penalties, it is likely that the consequences of not complying with privacy laws are not considered that important (Borking, 2009).

Problem – consequence of privacy breach may increase: As also mentioned earlier, an increased focus on privacy does not eliminate the risk of experiencing privacy breaches. Though the likelihood of a privacy breach may be reduced by increasing organizational awareness of privacy, the consequences of a breach may increase as users have higher privacy expectations.

Opportunity – more holistic privacy management: As businesses collect more and more personal information about their customers, the risk of privacy breaches are likely to increase unless measures are taken to protect privacy. Investments in PETs will likely influence also the existing routines and technology related to handling of personal information. Thus it brings with it an opportunity to handle privacy more holistically. As will be seen in the next section on existing PETs, many PETs are concerned about making the privacy policy more understandable to users. A side-effect of such PETs will be that the privacy policy is also easier accessible to managers, developers and others dealing with personal information within the organization. Thus it can be an important part in building privacy awareness among own employees, and as a result reducing the likelihood of privacy breaches.

Users Will Share More?

From the customers' point of view, improved privacy can be achieved by sharing less information, with the result that many services cannot be used. Privacy can however also be improved by businesses taking better care of the personal information they are trusted with. Businesses that are considered trustworthy when it comes to privacy may therefore find that users are willing to trust them with more information, and more correct information, than other businesses.

Problem – privacy consciousness of users: We have already identified a number of problems with using privacy as a competitive factor; users may not care about privacy, may forget about it, and dislike privacy trade-offs because of lack of understanding of the issues and the emotional implications. In order for users to take privacy into account, they have to somehow get an idea of how different businesses handle privacy issues. A main challenge in this respect is the lack of visibility of privacy in today's Internet services.

Opportunity – quality of information: Personal information, and especially correct personal information, is of high value to businesses. Experiments have shown that when users get explanations on why certain information is collected they are more willing to share information (Kobsa & Teltzrow, 2005). Improved privacy is likely to further increase users' willingness to share information about themselves.

Summary

As can be seen from our discussion, with every problem there is an accompanying opportunity for businesses to capitalize on users' need of

privacy. Especially since privacy awareness and protection is rather scarce, one need not do very much in order to convince users that privacy is taken seriously and thereby get a competitive advantage. However, the fundamental problem of privacy in business is the uncertainty when it comes to predicting user behavior and acceptance. It is therefore extremely difficult to calculate the benefit (for a business) of implementing privacy enhancing measures to ensure a reasonable return on investment. That being said, there are several measures to be taken that are relatively inexpensive and at least are assumed to comfort the most privacy conscious users.

A SURVEY OF PETS SUITABLE FOR BUSINESS ORGANIZATIONS

The foundation for most PETs dealing with information control is the privacy policies of service providers. Privacy policies are the most common way of explaining to users how data is collected, used, stored and shared. Privacy policies are however usually difficult to understand for the users (Cranor, Guduru, & Arjula, 2006). A common approach among existing PETs is, in varying ways, to make the content of privacy policies more accessible to users. To do this, the form in which policies are presented is important.

In this section we start describing the different shapes PETs may take. Then we give an introduction to P3P, before we provide an overview of a number of different PETs that are available either as products or as research prototypes – many of which build on P3P.

Different Shapes of PETs

PETs for information control differ in their sources of information, the actions they take and in where they are implemented. Below we give an overview of the main options available.

Sources of Privacy Information

Privacy related information can be made available to users and presented in various ways. Below we present the main alternatives. It should be noted that they can be used in combination; a business can for instance create a machine-readable policy that is based on the textual policy, and make both available to their users. In addition third parties can create additional sources of privacy information, e.g. through communities.

The most common way of presenting privacy policies is to publish them in textual form on the web site. Several surveys however show that few users actually bother to read such policies (Jensen, et al., 2005). There have been various surveys of the content of privacy policies. One such example is the study by Pollach (2007) of 50 policies belonging to popular web sites from various business areas. This study found that the policies were of varying length (the longest was more than ten times longer than the shortest one), and that the language of the policies is used actively to highlight positive aspects of privacy practices, and put more negative aspects in the background. Actually, the word "may" was found to be fourth most frequently used non-grammatical word! In addition, many, if not most, policies lack information that users care about.

The varying quality and accessibility of current privacy policies, and the fact that users do not bother to read them calls for other ways to present privacy policies in order for the message to come through. Several initiatives have worked on defining privacy icons (Hansen, 2009) so that users can get important privacy information by just taking a quick look at the icons – if familiar with the icons' uses.

Privacy policies can also be provided in machine-readable form so that PETs can use them in providing services to users. The most common example of such machine-readable policies is P3P.

Privacy information can also be provided without involving the service provider. One option is to

have third parties offering information to users on the privacy policies of popular service providers. The offering of privacy seals can be considered a special type of such service where some third party verifies that basic privacy issues are taken care of. Another commonly suggested approach is to have community services, e.g., where users can rate privacy policies of a business, and provide their own experiences when it comes to privacy.

PET Actions Taken

PETs related to information control can take on different task in order to provide support for users. First, PETs can present privacy information to users in a more easily understandable form than what is commonly available. As an example, PETs can interpret machine-readable policies and use them to extract parts that the PET's users are likely to care about. Then this information can be presented to the user. PETs can also retrieve information from privacy communities, and present information to users that is relevant for the site the user is currently visiting. However, PETs have the potential to move beyond simple information presentation. A very common type of functionality is matching of user preferences towards privacy policies of service providers. With such an approach users are expected to explain their privacy preferences to the agent, which in turn warns users in cases where the preferences and the privacy policy do not match. It will also be possible for PETs to block user behavior, if this is desirable.

A more advanced type of functionality that also can be supported by PETs is that of negotiating privacy terms associated with a service. Today, the common approach is for service providers to state their privacy terms, and users accept by using the service. It is however possible to envision that the privacy terms are negotiable, and that PETs can negotiate automatically on behalf of users based on knowledge of users' privacy preferences.

Some PETs also keep track of which information has been shared with the different service providers. Thus users are better able to exercise rights of access to, and eventually correction or deletion, of information concerning them.

PET Implementation

PETs also vary in who install or implement them into their systems. Service providers can implement PETs that e.g. make privacy policies more understandable for users. Users can install PETs, e.g. in form of browser plug-ins. Such user agents are typically running during all web interactions, and can assist with real-time evaluation of possible information disclosure as it is being performed by the user. To support agents, businesses do not necessarily have to make changes to their systems, other than for instance providing machine-readable policies if the PETs rely on their availability. PETs can also require third-party support, e.g. in the form of community services. But there are also more complex PETs available that take the form of middleware that needs to be supported by both users and service providers.

The Platform for Privacy Preferences (P3P)

P3P (W3C, 2006) is a standard that enables service providers to communicate the privacy policies of web sites to their clients. P3P provides both a standardized format for privacy policies and a protocol that enables web browsers to read and process the privacy policies automatically. P3P was developed by the World Wide Web Consortium (W3C) and was the first privacy policy language to be standardized by the W3C. The latest specification for P3P is version 1.1, which was finalized in 2006.

P3P is an XML-based language that allows service providers to express privacy policies in a machine understandable way. By using P3P, service providers can specify rules that include the

type of data, the type of use, the user of the data, the purpose of the use and how long the data will be retained. User agents can then fetch P3P policies, interpret them and present them to the end-user in order to assist the end-users in determining whether a particular service provider's published privacy policy matches the user's individual privacy preferences. The user agent can also warn the user in the case of a mismatch between the privacy policy and the user's preferences. Such user agents can be built into web browsers, browser plug-ins or proxy servers. They can also be implemented as a part of electronic wallets, form-fillers or any other user data management tool. Users therefore do not need to read the privacy policies at every site they visit or for every service they access.

P3P-Enabling a Website

To P3P-enable a web site, the service provider translates the privacy policy for the site into one or more P3P policy files. The policy files are then made available to the user, either by storing them in a well-known location (/w3c/p3p.xml), or by embedding information in the HTML response (directly in the header or by using link tags) (Cranor & Lessig, 2002). The policy file will then be collected by the client when visiting the web site, and read and processed by the client computer.

The P3P Specification

The P3P data scheme consists of a set of data elements, sets and structures. Data elements are individual pieces of data, for example first name or telephone number. Data sets are used to combine data elements into groups; For example, the home postal address data set contains data elements for street address, city, state, postal code and country. Data structures can be used to create templates which can be reused with multiple data sets (Cranor & Lessig, 2002).

The P3P specification defines two types of P3P policies; full policies and compact policies. The full policy is written in a XML format based on the P3P vocabulary and data scheme. The compact policy is used to describe the privacy practices related to the use of cookies. According to the P3P specification, a compact policy is required to have a corresponding full P3P privacy policy.

A full P3P policy is written as a sequence of STATEMENT elements, which includes several sub-elements:

- PURPOSE: The purpose for information collection. The P3P specification contains 12 pre-defined values, for example "current" (to complete and support the activity for which the data was provided) and "pseudo-analysis" (to infer e.g. habits and interests of individuals without identifying specific individuals).
- RECIPIENT: The intended user of the information. The P3P specification contains 6 pre-defined values, for example "ours" (ourselves) and "same" (legal entities following our practices).
- RETENTION: the duration that the collected information will be kept. The P3P specification contains 5 pre-defined values, for example "stated-purpose" (discard information at the earliest time possible) and "indefinitely".
- DATA-GROUP: lists the individual data items that will be collected for the stated purpose

A simple example of a full P3P policy is the privacy policy for the P3P handbook website (http://p3pbook.com/). The human-readable privacy policy looks like this:

This is the web site for the book Web Privacy with P3P by Lorrie Faith Cranor. We do not currently collect any information from visitors to this site except the information contained in standard web server logs (your IP address, referrer, information about your web browser, information about your

HTTP requests, etc.). The information in these logs will be used only by us and the server administrators for website and system administration, and for improving this site. It will not be disclosed unless required by law. We may retain these log files indefinitely. Please direct questions about this privacy policy to privacy@p3pbook.com.

The corresponding machine-readable P3P policy file looks like this:

```
<POLICIES>
<POLICY discuri=”http://p3pbook.com/privacy.
    html” name=”policy”>
<ENTITY>
<DATA-GROUP>
<DATA ref=”#business.contact-info.online.
    email”>privacy@p3pbook.com </DATA>
<DATA ref=”#business.contact-info.online.
    uri”>http://p3pbook.com/
</DATA>
<DATA ref=”#business.name”>Web Privacy With
    P3P</DATA>
</DATA-GROUP>
</ENTITY>
<ACCESS>
<nonident/>
</ACCESS>
<STATEMENT>
<CONSEQUENCE>
Our Web server collects access logs containing
    this information.
</CONSEQUENCE>
<PURPOSE>
<admin/>
<current/>
<develop/>
</PURPOSE>
<RECIPIENT>
<ours/>
</RECIPIENT>
<RETENTION>
<indefinitely/>
</RETENTION>
<DATA-GROUP>
<DATA ref=”#dynamic.clickstream”/>
<DATA ref=”#dynamic.http”/>
</DATA-GROUP>
</STATEMENT>
</POLICY>
</POLICIES>
```

A compact P3P policy contains a set of tokens that represents the following elements from the P3P vocabulary: ACCESS, CATEGORIES, DISPUTES, NON-INDENTIFIABLE, PURPOSE, RECIPIENT, REMEDIES and RETENTION. The vocabulary is designed to reduce the number of bytes transferred within a HTTP response header. An example of a P3P compact policy is

CP=”NOI DSP COR NID ADM DEV PSA OUR IND UNI PUR COM NAV INT STA”

where for example “PSA” and “OUR” represent the <pseudo-analysis/> and <ours/>” value in the PURPOSE and RECIPIENT element respectively (i.e. in this simple example these two tokens mean that the site will use the collected data only for themselves, in order to infer, e.g., habits and interests of individuals without identifying specific individuals).

APPEL: User Preferences

The creators of P3P also designed the P3P preference language (APPEL), which is a standard for encoding the user’s privacy preferences in a machine-readable way. Similarly to P3P, an APPEL preference file therefore consists of a set of machine-readable rules. The primary goal when developing APPEL was to simplify the sharing and installation of user preferences in terms of machine-readable rule-sets. By storing privacy preferences in a file, the preferences can be shared amongst users with similar privacy principles and easily transferred between different devices. The APPEL rule set is composed of patterns that can be matched against P3P policies, and specifies action that will be taken is case the policy does not match

the preferences. The syntax is very similar to the P3P policy syntax. APPEL is further described in (Cranor, Langheinrich, & Marchiori, 2002).

P3P User Agents

A P3P user agent is capable of matching privacy policies to stated user preferences autonomously. Such user agents can be implemented in web browsers, electronic wallet, ISP proxies etc. The basic concept is as follows: Whenever a user makes a service request, the user agent will retrieve the privacy policy of the service and compare it to the user preferences. Depending on the capabilities of the agent and languages used for specification, the agent may enter a negotiating phase, initiate mitigating measures, block access to the service or simply issue a warning that the policy does not match the user's preferences. Microsoft's Internet Explorer 6 (IE6) was one of the first P3P user agents available on the market. It contains built-in support for fetching and displaying full P3P policies in a human-readable format, and to match user preferences with P3P compact policies (concerning the use of cookies). Examples of stand-alone P3P user agents will be presented later in this chapter; the most well-known being the AT&T Privacy Bird (Cranor, Arjula, & Guduru, 2002).

The user agent was considered a crucial part of the P3P project. The P3P specification therefore includes guiding principles, which are specific recommendations for P3P user agent software design. The guiding principles include advice on how the users should be informed, and how the software should be configured by default.

P3P Policy and Preferences Tools

The goal of the P3P project was twofold. It allows service providers to formulate and present their privacy policies in standardized, easily-located, machine-readable manner. It also provides the users with a mechanism to understand how their personal data will be collected and used. However, the syntax of both P3P and APPEL makes it very difficult for human beings to correctly specify their intentions, both regarding privacy policies and privacy preferences. Especially, writing a P3P policy can be a tedious procedure, just look at the simple example provided above. Several tools have therefore been developed to facilitate this process. An example is the IBM P3P policy editor (Bergmann, Rost, & Pettersson, 2006) that was developed to assist the service providers in translating their privacy policy into the P3P format.

Most of the existing P3P user agents contain support for preference generation. For example, the AT&T Privacy Bird contains a graphical user interface that allows the user to specify preferences based on a subset of the P3P vocabulary (Cranor, Arjula, et al., 2002).

Examples of PETs

Below we briefly describe a number of existing PET solutions and research prototypes for improved information control. We give an overview of their main idea and solution, and explain what information source they rely on, what actions they take, and where they are implemented. An overview of the PETs is given in Table 1.

P3P User Agents: AT&T Privacy Bird, Integrated Privacy View

The AT&T Privacy Bird (Cranor, Arjula, et al., 2002) is an early P3P user agent, which was implemented as a browser helper object for Internet Explorer. A graphical window allows the user to set up his/her privacy preferences, based on a subset of the P3P specification vocabulary. The AT&T Privacy Bird then automatically retrieves privacy policies from service providers and compares these with the user's specified privacy preferences. The user is warned of any mismatches, and can also access a summary of the P3P policy.

Table 1. Overview of existing PETs

PET	Information source	Actions taken	Implementation
AT&T Privacy Bird	P3P policy	Match and present	Browser plug-in
Integrated Privacy View	P3P extension, edit HTML	Match	Browser add-on
Privacy and Identity Management for Europe (PRIME)	Service provider	Manage credentials, control interactions, present, match, negotiate	Needs to be supported by both users and service providers
Collaborative privacy	Community	Match and present	Browser plug-in
Web of trust	Community	Present	Browser add-on
PIPWatch toolbar	Community	Match with legislation	Web browser add-on
Privacy panel	Service provider	Provide access to privacy information and functionality	Service provider; web site and underlying system
Privacy nutrition label	P3P policy	Present	Not specified
Privacy finder	P3P policy	Present	Online search engine

Another example of a user agent that utilizes P3P policies is the Integrated Privacy View (IPV) (Levy & Gutwin, 2005). Where standard P3P operates on a page level, IPV considers each input field separately. This is possible through an extension of P3P that makes it possible to link policy statements directly to HTML elements. IPV can thus present how an individual piece of information is handled, and whether or not this matches the user's preferences.

The potential of Privacy Bird has been shown in several user studies (Cranor, et al., 2006). In practice however, its usefulness has turned out to be very limited. A major reason is that the agent is dependent on the availability of P3P privacy policies, and that few service providers have published such policies. IPV similarly relies on the availability of P3P policies, and in addition service providers need to make changes to their HTML code.

Privacy Management: Privacy and Identity Management for Europe (PRIME)

Privacy agents such as Privacy Bird only deal with a limited part of privacy management. The PRIME project is an example of a more comprehensive privacy solution that has the potential of offering better privacy, but at a higher cost. The PRIME architecture (Andersson, et al., 2005; Camenisch, et al., 2005; Pettersson, et al., 2005) brings different PETs together, in order to protect and improve Internet users' privacy. Their goals include

- to achieve user informed consent and control related to all disclosure of personal data,
- to allow privacy negotiations in order to come to agreements with service providers on how personal data should be handled,
- to provide solutions for data minimization so that only the necessary information is collected and that it is deleted when no longer needed,
- to develop solutions for user controlled identity management,
- to allow a spectrum of anonymity services, ranging from the use of anonymous communication channels to the use of identity proofs issued by third parties, and
- to ensure accountability for anonymous users.

Note that the three latter goals strictly speaking fall outside the scope we are considering in this chapter.

PRIME services involve three main parties: the users, the service providers and a certification authority. Users and service providers share essentially the same architecture. Both hold a database with certificates and declarations of the party, default policies for handling of data, and logs of previous interactions with other parties. Service providers in addition hold data collected from other parties. Access to this database is controlled by the access control component, identity control component and the graphical user interface. Service providers also have an obligation management component. PRIME offers the possibility to include components that assure policy compliance.

For service providers, support of PRIME requires modifications to their systems. It is however possible to start supporting parts of the PRIME solution and then to gradually include more and more privacy support. One particularly interesting feature of PRIME is the possibility to assess the trustworthiness of service providers. This feature makes it easier for smaller companies to gain trust more quickly than what is common today.

Communities: Collaborative Privacy Management, Web of Trust, PIPWatch Toolbar

Because of the lack of availability of P3P policies, it may be – at least in the near future – that PETs are more likely to be successful if they rely less on service providers, instead of putting more requirements on them, as is done with PRIME. Kolter et al. (2010) uses similar arguments as a basis for suggesting collaborative privacy management. Their tool consists of four parts; a privacy preference generator, a privacy agent, a data disclosure log and a browser plug-in. The basic concept is similar to that of e.g. Privacy Bird; users specify preferences that are matched with the published P3P policies of service providers. The major difference is that the tool is not dependent on support from the service providers. Rather than convincing the service providers to implement and update their privacy policies, the tool relies on support from an online privacy community that will contribute to privacy related information on service provides by using a Wiki-like Web front-end. Examples of information that can be provided include general information on the service provider, the amount of data that is required, information on who the service provider shares data with, explanations of the textual privacy policy, historical policies, information on whether or not the service providers adhere to their policy and individual experiences with this service provider when it comes to privacy.

A simpler but already up and running approach is the Web of trust (http://www.mywot.com/) add-on where users can share information relevant for trust through a community. Users can rate websites based on their own experiences, and can in turn get information on the trustworthiness, vendor reliability privacy and child safety of a web site. For privacy the add-on can provide information on whether or not the site has a privacy policy and the privacy implications of the practice stated in this policy.

Another community based browser add-on is the PIPWatch Toolbar (Clement, et al., 2008) that helps the users to interpret whether the privacy policies of the websites they visit comply with the Canadian private-sector privacy legislation. The PIPWatch toolbar is not based on any privacy policy language. Instead users are expected to contribute with information on the website by using functionality embedded in the toolbar. This includes filling out basic information on the website and to send email to privacy officers, asking them to fill out a questionnaire. The responses will be stored at a central server and used by the toolbar to evaluate to what degree the site fulfils the users' privacy expectations and the privacy legislation.

The community based PETs do not put requirements on service providers, and can therefore be

considered of little direct relevance for businesses. Businesses should however care about what is registered about them in various privacy communities. Privacy communities are in many ways a reaction to the lack of useful privacy information available from today's businesses. As businesses do not provide what users need to make privacy decisions, users have to provide it themselves, but then businesses have no guarantee that what is published about them is indeed correct.

Policy Presentation: Privacy Nutrition Label, Privacy Panel

Some PETs focus mainly in policy presentation, in order to make privacy information more available and understandable to users. We have already mentioned the concept of privacy icons. Another example is the Privacy Nutrition Label (Kelley, et al., 2009), which is a design approach to improve the visual presentation of privacy policies to end users. The design was inspired by nutrition facts panel on food products in the United States, hence the name. The Privacy Nutrition Label is a refinement of its predecessors; the Expendable Grid (Reeder, et al., 2008), which implemented the entire P3P specification and thus turned out be very hard to read and comprehend; and the Simplified Label, which the authors considered to be *too* simple (Kelley, et al., 2009). The Privacy Nutrition Label uses a combination of symbols and color codes to illustrate how the user's personal data will be treated. The data is organized according to what type of information it represents, how it will be used by the service provider itself and whether it will be shared with 3rd party service providers. The terminology used in the Privacy Nutrition Label is derived from the P3P specification, but simplified in order to fit into a one-page summary. The design has been evaluated by a laboratory study that compared how end-users perceived the difference between ordinary textual-based privacy policies and privacy policies presented as Privacy Nutrition Labels. The results indicate that privacy policy information is easier to find and more enjoyable to read when presented as a Privacy Nutrition Label (Kelley, et al., 2009).

Where the Privacy Nutrition Label is mainly concerned with visualizing the content of the policy, the privacy panel proposed by Schunter and Waidner (2007) is more concerned with how users can easily access important privacy information on a website. They suggest a panel that consists of four icons: one for access to the privacy policy, one for access to the stored data about oneself, one to block identity and another to delete identity. Schunter and Waidner propose that this panel can be standardized and used at all participating web sites, such that the users will have a common way to control their data. The main contribution of Schunter and Waidner's paper is not the design of the quite simple user interface, but rather the formalization of the underlying protocol mechanisms that handles how privacy controls are implemented and used both locally and across multiple organizations.

Policy presentation is relevant for businesses that want to make their privacy policy more available. The Privacy Nutrition Label is based on P3P and can be used by e.g. user agents to provide presentations of various site policies, but businesses can also use this visualization on their own web site. The privacy panel is however something that must be included on the website, and supported in the underlying systems.

Search-Based Approaches: Privacy Finder

Search-based approaches can be especially useful when users are considering whether or not to conduct business with a particular provider. An example of such an approach is the privacy-enhanced online search engine Privacy Finder (http://www.privacyfinder.org/) developed and operated by the CyLab Usable Privacy and Security Laboratory (CUPS) at Carnegie Mellon University. Privacy Finder orders search results according to their

P3P privacy policies. A "privacy meter" next to the search result indicates whether a P3P policy exists, and to what degree the policy corresponds with a list of preset privacy preferences. Clicking on the privacy meter will open up more detailed a privacy report on the site.

Like for the other PETs that are based on P3P policies, businesses that want to be evaluated on privacy by Privacy Finder need to make their policy available in P3P. The incentives for doing so will depend on the popularity of such search engines, and the action taken if a P3P policy is not available (e.g., if the site is then given the lowest score).

CASE STUDY: THE RISE AND FALL OF P3P

The work on P3P started in the 1990ies, and represents a great step forward in the quest of automating the process of communicating data management practices. However, its history has been controversial and the standard has been subject to numerous disputes. In this section we discuss the background of P3P and point out some of its technical limitations and obstacles, before touching upon the P3P uptake and summarizing the main reasons for its failure.

Background and History

P3P was developed over several years. Early work started in 1995 when three associates at the Center for Democracy and Technology (CDT) discussed the idea of using the Platform for Internet Content Selection (PICS) (W3C, 1997) for protection of Internet user privacy. The idea was then further explored by the Internet Privacy Working Group (IPWG) who developed a draft privacy vocabulary (a standard set of terms for web sites to describe their privacy practices), which eventually was named the Platform for Privacy Preferences (P3P) (Cranor & Lessig, 2002; W3C, 2006). P3P was officially launched as a W3C project in May1997. The development process was initially projected in an 18-month period but turned out to continue through the next three years. The details of the specification were changed a number of times during this period. Initially automatic data transfer mechanisms were proposed as a part of the P3P specification, but were later removed. Negotiation of privacy policy was also considered to be a part of the P3P protocol and vocabulary, but never made it into the final specification. The P3P 1.0 specification was eventually released as a "candidate recommendation" in late 2000 (Cranor & Lessig, 2002). The relatively long development process was explained by the P3P developers to be caused by the "deliberative and thoughtful process" behind P3P (Cranor, Schwartz 1999), which amongst other things included revision of the specification based on feedback from the European Commission (Cranor & Lessig, 2002).

Criticisms

The P3P specification has been heavily questioned, already from the beginning. Large parts of the critique against P3P as a privacy protection mechanism is based its limited scope. P3P was a U.S. initiative and has its roots in the U.S. Federal Trade Commission (FTC) privacy model (Hochheiser, 2002). As described earlier in this chapter, the FTC privacy model consists of five parts: Notice/awareness, choice/consent, access/participation, integrity/security and enforcement/redress.

As can be seen, the P3P specification only covers the two first parts: notice/awareness (fetching and displaying electronic policy files) and choice/consent (specifying user preferences and matching them with privacy policies); it does not attempt to cover any of the other three parts. Probably the most fundamental problem is that P3P does not comprise enforcement. P3P in itself does not provide any mechanism for making sure service providers act according to their stated policies.

Hence, the technology involves a risk that the user may be mislead into believing that if a privacy policy matches the stated privacy preferences, it will be safe to proceed. Critics such as Coyle (Coyle, 2000) therefore argue that P3P fail to address privacy since it will never be able to provide general privacy protection for Internet users. This view is shared by Catlett (Catlett, 2000) who states that "P3P is not going to protect privacy, and the public shouldn't continue to be told it will".

The above criticism mainly applies to the U.S. (Hochheiser, 2002). Many countries have laws that protect the individual's privacy. As an example, the EU privacy directive (2002) puts strong restrictions on the use of personal data. To comply with the privacy principles stated in the EU privacy directive, organizations must, amongst other things, inform their users about the use of their personal data. For example, organizations need to get informed consent from data subjects before processing personal data (the principle of consent). Moreover, the directory states that data subjects shall be able to check and influence the processing of their data (the principle of data subject control). In addition, organizations are not allowed to process data in a way that is incompatible with the stated purpose of collection. The EU privacy directive is therefore considerably stricter that the FTC privacy model, which is more ambiguous in several respects. Non U.S. Internet users, such as EU citizens, may therefore find P3P useful as a complement to already existing legal requirements that limits the collection of personal information (Hochheiser, 2002).

Other critics are based on the political context surrounding P3P. In the middle of the 90ths several workshops and reports from the FTC promoted self-regulation as the preferred approach to privacy protection (Hochheiser, 2002). However, many privacy activists rejected the idea of self-regulation and privacy. Catlett compared privacy in the U.S. context with copyright in the music industry to motivate why privacy preferences and "promises" (i.e. policies) is a model ill-suited to protect consumer privacy (Catlett, 2000). He stated that clearly a law is a necessity to protect the privacy of American consumers. Also Coyle (1999) argued that since the U.S does not have data protection laws, a possible breach of the privacy agreement stated in a P3P policy will not have any negative consequences whatsoever for the company in question. A common argument articulated by privacy activists was therefore that the large corporation support of P3P is an attempt to avoid privacy legislation. Similarly privacy activists also maintained that the slow development pace of the standard was a mean to delay or avoid regulatory actions (Hochheiser, Catlett).

Technical Limitations and Obstacles

There has also been heavy criticism regarding the technical limitations of P3P. Hochheiser (2002) identified the limited scope (only web sites), lack of limitations on personal data collection and the restricted vocabulary as the main areas of technical criticism against the P3P specification. Particularly the precision of terms is said to be sub-optimal, which for example makes it difficult for service providers to express conformance to privacy legislation. While extensions may be developed to answer several of these shortcomings, this would conflict with the P3P intention of being simple and easy to use.

On the other hand, Catlett argued that the P3P vocabulary was unnecessary complex, since "the core of consumers' desires for privacy are simple and easily stated, but unpalatable for marketers: consumers don't want their personal information sold, shared, or reused for secondary purposes" (Catlett, 2000). His view was shared by Coyle, who stated that collection of personal data is only necessary when purchasing a product online (Coyle, 2000).

Service providers also experienced practical difficulties when attempting to translate to textual policies to the P3P format, something that contributed to the large number of erroneous policies that

can be found online. Moreover, P3P lacks support for stating that accepting the current policy also implies accepting any future versions of the policy.

Also APPEL has its limitations, as pointed out by e.g. Agrawal et al. (2003). They state that simple preferences are surprisingly hard to express in APPEL and that writing the preferences is an error prone procedure. In addition, APPEL suffers from some serious problems that arise from the fundamental interactions between P3P and APPEL. In particular, they point out that with APPEL users can only specify what is unacceptable, not what is acceptable. Another problem with APPEL is that its limited expressiveness makes it difficult to translate real human intentions to a set of machine-readable preferences, even with support from a user agent with a graphical interface.

P3P Uptake

To achieve its goal, the P3P concept relied on support from both the client and the server side. On the client side, P3P support had to be built into the browsers, as core functionality or as browser add-ons. Additionally, the users, who access P3P enabled web sites using their browsers, had to understand and accept the technology in order to install and configure the P3P preferences settings. On the server side, the service providers had to encode their privacy policies in the P3P format, maintain them, and make them available to the P3P user agent.

Even though there have been occasional reports of increasing rates of P3P adoption amongst service providers, especially for e-commerce sites (Egelman et al. (2006) showed that out of a sample of e-commerce websites, 21% contained a P3P policy), it seems like adoption of the specification is doomed to fail. A study performed in 2005 showed that although 28% of the 75 most popular domains (the most clicked on domains from America Online search results) have P3P policies, only 9% of the sites of a random selection were P3P enabled (Cranor, et al., 2008). The study also revealed that a majority of the P3P policies contained syntax errors. After what seems to be the peak in 2005, the number of P3P enabled sites has been steadily decreasing. A simple sample performed at the time of writing revealed that large service providers, such as Google.com and Apple.com, currently do not provide P3P privacy policies at all. The same is true for social networking sites such as Facebook.com and LinkedIn.com.

After the launch of P3P in early 2000, Microsoft has been an avid supporter of the technology, to the extent that Internet Explorer version 6 (IE6) and its successors (IE7 and IE8) provide the ability to display P3P privacy policies and to compare the policies with the user's privacy preferences settings. However, the support is quite limited, in that only a subset of the specification has been implemented in the conformance checking; namely the compact P3P policy, which covers the use of cookies. IE will not alert the user if the web site violates the privacy preferences regarding any other personal information, such as user-provided data. IE uses the P3P compact policy, which is transmitted in the HTTP headers, to make cookie blocking decisions. If the cookie policy of the service provider does not match the user's preferences, IE will display an eye covered by a do-not-enter sign in the browser frame.

The release of the P3P-enabled IE6 by Microsoft in 2001 was probably the most important factor for the large number of websites that adopted P3P in the following years. Service providers quickly found out that supporting P3P was necessary for their web sites to function properly when viewed using this browser (Cranor & Lessig, 2002). IE6 employed a fairly strict privacy policy regarding cookies. As previously discussed, IE looked for P3P compact policies; however, in addition to warning the user, IE6 also automatically blocked third-party cookies if a compact policy could not be found. This lead to a situation where sites that were not P3P enabled, but employed targeted advertising, page counters, etc. (that rely on 3rd part cookies), did not work in IE6, and its subsequent

versions (P3P user agents based on compact policies were also implemented as a part of Mozilla Firefox and Netscape web browsers in early 2000, however, the functionality has since then been removed).

At its birth, P3P was praised by the Internet community and widely believed to be the keystone to resolving large privacy issues on the web. In practice, the specification has never been widely adopted, something that by some was predicted already early in the project (Catlett, 2000). In 2006, the work on P3P was officially suspended by the specification group due to the "insufficient support from current browser implementers" (W3C, 2007). In an interview from 2009 one of the W3C spokespersons on privacy stated the main reason for the P3P failure was that they "did not manage to convince the browsers" (Out-Law, 2009). Even though P3P policies were implemented and made available by the service providers, transparency was never really achieved because the client side did not utilize all the privacy information (i.e. P3P policies) that was actually out there.

Reasons for Failure

There are several reasons why P3P never became a success in the market. In hindsight, we are here able to point out some of the most obvious ones.

P3P's dependence on service providers to declare their privacy policy using the P3P policy language is one of the factors explaining why it never reached critical mass of adoption (the stage where adoption becomes self-reinforcing). Some have asserted that P3P suffered from the "chicken and egg" problem (Schwartz, 2009). The problem is that with few P3P compliant web sites, the user demand for P3P user agents is low, which in turn reduces the usage of P3P-declared policies. And if no one is using the P3P policies, why should service providers bother providing them? Reaching the critical mass is what turns this around to a positive reinforcement, rather than a negative one.

It has been said that many of the service providers who were early P3P adopters supported the technology to show their customers that they respected their privacy and to demonstrate a voluntary step to address privacy concerns (Cranor & Lessig, 2002). On the contrary, service providers currently have little or no incentive to implement or facilitate the use of machine-readable privacy policies such as P3P. From the service providers' perspective, information gathering is extremely important; however, there is a trade-off between obtaining as much information about the users a possible, while still preserving their trust. If service providers are to implement privacy enhancing measures, there will have to be a very concrete and measurable upside - fuzzy concepts such as "community spirit" and "doing the right thing" are unlikely to carry the day. In our opinion, for a user agent to become useful in practice, convincing the service providers to adopt the concept is the key to success. P3P never represented a competitive advantage for the service providers in this respect.

Even though the P3P vocabulary may have its limitations, its specification was by many considered as far too complex. There was never any user interface that could use all the metadata in the specification; most of the tools implemented only a subset of the specification. P3P was therefore never fully implemented, as the creators had hoped (Schwartz, 2009).

Microsoft did it best to push the usage of P3P by forcing web site developers to create P3P compact policies to be able to use third-party cookies. However, the result was not as intended. Many developers chose a "quick fix" by simply looking for P3P policies and compact policies on random Web sites and copying them onto their own sites (Cranor, 2002). Even though Microsoft's effort contributed to the increased number of P3P-enabled sites, the result was not as expected. A P3P policy is considered a contract and must therefore be consistent with the web site's human readable policy. The large number of random (and often

erroneous) P3P compact policies that could be found online did not contribute to its credibility.

As discussed in the previous section, the P3P working group never managed to convince the big browser developers to support the technology (with the exception of Microsoft and its cookie handling in IE6). One of the key concepts of P3P was transparency, i.e. its ability to let the users say what is acceptable and what is not, and be informed of to what degree a particular service provider treats his personal information in ways that he does not like, without any knowledge of the underlying technology. This objective could never be achieved without support from the browser developers. Even though several browser add-ons were developed and distributed (the most well-known being the AT&T Privacy Bird), the lack of built-in P3P support in the browsers meant that the majority of the users were never aware of the technology and its possible implication on their personal data management.

The fact that P3P in its basic form cannot ensure that a service provider conforms to its own privacy policy, contributes to the hesitancy to adopt the technology. Especially browser developers are likely to be reluctant of enabling support for the technology because a P3P policy represents a privacy promise that may not be kept. Supporting P3P may therefore fool the users into believing that their browser will be responsible for handling their private information in the way stated in the users' privacy preferences.

The end-users also represented a huge challenge in the P3P project. The deployment of P3P assumed that the users regarded the privacy policies as trustworthy, and that they would act in accordance with their privacy preferences. Unfortunately, as discussed earlier in this chapter; even though most users claim to be concerned about their privacy, this has little effect on their actual behavior. This implies that most people do not take the time to read through the privacy policy, and those who do probably do not understand it. We also believe that most users do not understand the implications of personal data sharing; they look for short-time benefits rather than long-time consequences. It is also questionable whether users pay attention to, or tinker with, the default settings for any software that is already working. It is therefore highly uncertain whether users will spend time to actively specify their privacy preferences, or even to browse through the P3P summaries presented by the client software. Finally, most Internet users are not technically savvy. It is therefore a challenge to make people play an active role in the protection of their own privacy.

Current Status

P3P is not dead. Creating new user agents or new languages based on P3P is not a waste of time. In fact, in a time when social networking and personal data sharing has become an integrated part of the daily life, tool support to help the users control their privacy is more important than ever before. The concept of machine-readable policies, and preferences, is still considered a very promising approach. We should learn from the P3P creation, deployment and promoting process when building future privacy enhancing technology.

FUTURE RESEARCH DIRECTIONS

The adoption of PETs is currently low. One possible way of increasing adoption is of to push PETs through legislation, forcing businesses to support and provide such technology. Another option is to build PETs that are relevant and beneficial enough to be implemented of own free will. Such new and more relevant PETs need to:

- Constitute a competitive advantage for service providers.
- Be easy and pleasant to use for end users.
- Provide functionality that adequately represents the complexity of the privacy area.

As explained above, P3P fails in all these respects. Service providers that supported P3P early on did not gain a competitive advantage, as the technology was only useful to users when a large portion of the service providers supported the technology, and when proper user agents were available that built on P3P. For end users, the agents that are based on P3P have suffered from usability problems, e.g. when it comes to presentation of policies in an understandable way, and in the support users are given in their preference specification. P3P has also gained critique for shortcomings of the vocabulary.

If service providers are to implement PETs, they will need a clear and positive answer to the question: "What's in it for them?" As previously explained, the motivation for offering and supporting PETs can be to build reputation and trust, to keep their customers happy, to gain a competitive advantage or to reduce the risks associated with privacy compromise. PETs may even result in an increased willingness of users to share personal information. However, the investments that need to be made, both technically and for marketing purposes, and the uncertainty of the effects of these investments are likely to result in businesses investing their money elsewhere. More research is thus needed on the likely effects of using PETs so that businesses can have an improved decision basis.

Current PETs are, like P3P, commonly dependent on a certain level of adoption to be considered useful for end-users. Thus, businesses have little incentives to be first out on PET adoption. New PETs need to solve these problems, so that businesses that are pioneers in implementing PETs, can gain from this investment even though others may not follow them.

If PETs are to be successful, they need to be perceived as useful by the end-users. Users also must be made aware of them. This is challenging as today's Internet users do not spend a lot of time and effort on managing their privacy. Research is needed to improve our knowledge of what privacy support that users actually want. More research is also needed on how to present privacy information in a user-understandable way. But maybe even more important, we need PETs that are useful to users despite users' lack of privacy understanding, and lack of willingness to spend time on privacy. PETs should make the right privacy information and advice available at the right point in time. The Integrated Privacy View solution can be used as a starting point in this respect, as IPV allows users to get easy access to the privacy implications of providing specific information pieces at the point in time when this information is shared. Using PETs should also not be time-consuming and require users to take part in complex setup processes, potentially forcing them to make difficult and emotion-laden decisions regarding what they allow or disallow when it comes to personal information. The privacy paradox aside, users are most frequently put off by lengthy configuration processes for specifying preferences. One way of tackling this challenge may be through machine learning approaches that track user behavior, helping users to behave consistently with respect to privacy (Tøndel, Nyre, & Bernsmed, 2011). An additional challenge here is today's pervasive computing practices, where users access the web through a multitude of devices, many of which are mobile. This also implies that a future "preference generator" will need to take context information such as time and place into account.

Privacy is complex and personal. Individuals may have quite different opinions on what is acceptable and what is not, and what is acceptable may additionally vary between situations. Solutions that take an over-simplistic view of privacy, e.g. by having users specify a set of simple rules that always apply, are likely to be less useful to users, producing advices that do not match the situation the user is facing. Instead, privacy solutions are needed that take into account the complexity of privacy decisions. Especially solutions should acknowledge that privacy decisions are in many ways cost-benefit tradeoffs. Tools

that only consider the costs will be less relevant than tools that also consider the benefits involved. Currently PETs in general have only focused on the costs of sharing information, while businesses and end-users are more focused on achieving benefits. These benefits should also be taken into account by PETs.

More research is also needed on the role that third-parties should play in privacy technology. A limitation of the privacy policy approach (e.g. P3P) is the lack of control that companies actually follow what is written in their policies. Stricter privacy legislation is an option, but one could also consider improved privacy seal programs and third-party control, possibly integrated with PET tools.

Participating in social networks while at the same time preserving privacy is currently a challenging task. Neither Facebook nor LinkedIn currently support P3P, and it is not clear how this could be successfully implemented. Privacy in social networks could be a research topic of its own.

In the quest for relevant and useful PETs it will be important to look beyond the common focus on improving presentation and availability of privacy information. Real advantages are likely to be experienced when privacy support is built into the systems, and not only addressed by adding privacy policies. As an example of such a future privacy service, we describe some ideas related to service adaptation (Nyre, Bernsmed, Bøe, & Pedersen, 2011). This type of PET can be useful in cases where collecting personal information is not central to the service to be provided, but still considered by the provider as a nice option. Assuming that a user employs a browser or third-party plug-in that has access to the user's privacy preferences, these can be communicated to the service provider during session initiation. Based on these preferences, the service provider can refrain from asking for information which would violate the user's privacy preferences. This would improve the user experience, saving the user from eschewing a site which is too noisy, and the service provider would not lose a potential customer because of information that the provider doesn't really need. For other users with more relaxed preferences (or even with no privacy preferences) the service provider can collect more information, possibly in return for more services. For this to work, a universal, standardized format for privacy preference specification is needed. It may even be possible to establish a community portal of privacy preferences, where a simple ID (or URI similar to those provided by http://tinyurl.com) can be used to identify a complex policy. The user (or client) then only needs to communicate the ID to the provider, which can retrieve the full policy from the online repository. Tools for generating privacy preferences would upload complete policies to the repository, but only new unique policies will create a new ID.

To sum up the research challenges, there is a need for improved knowledge and understanding of the likely effects for businesses of supporting or offering PETs, but we also need to develop improved privacy technology. Researchers should learn from the failure of P3P and develop PETs that are better able to meet the requirements of both end-users and businesses, and can thus be used as a business advantage.

CONCLUSION

This chapter has discussed Privacy Enhancing Technologies in a business context, highlighting both challenges and opportunities. One insight that emerges from this work is that the most important part businesses can play is the facilitation of PETs deployed by others; i.e., even if businesses are not pushing PETs at their customers, they should provide foundations and interfaces that cater to the PETs that their customers decide to use. This includes making machine-readable privacy policies available.

Although P3P has failed to deliver on its promise, we maintain that it is important for busi-

nesses to be proactive when it comes to PETs. The alternative would be to surrender the playing field to the community-based approaches, and although these can provide a useful service to consumers, the businesses themselves will have few opportunities to influence them with respect to incomplete or erroneous information.

REFERENCES

W3C. (1997). *Platform for internet content selection* (PICS). Retrieved March, 2011, from http://www.w3.org/PICS

W3C. (2006). *Platform for privacy preferences*. Retrieved April, 2011, from http://www.w3.org/P3P/

W3C. (2007). *The W3C privacy page.* Retrieved March, 2011, from http://www.w3.org/Privacy/

Ackerman, M. S., Cranor, L. F., & Reagle, J. (1999). *Privacy in e-commerce: Examining user scenarios and privacy preferences*. Paper presented at the 1st ACM Conference on Electronic Commerce.

Acquisti, A. (2002). *Protecting privacy with economics: Economic incentives for preventive technologies in ubiquitous computing environments*. Paper presented at the Workshop on Socially-informed Design of Privacy-enhancing Solutions in Ubiquitous Computing.

Agrawal, R., et al. (2003). *An XPath-based preference language for P3P*. Paper presented at the 12th International Conference on World Wide Web.

Andersson, C., et al. (2005). *Trust in PRIME*. Paper presented at the Fifth IEEE International Symposium on Signal Processing and Information Technology.

Argyrakis, J., Gritzalis, S., & Kioulafas, C. (2003). Privacy enhancing technologies: A review. In Traunmüller, R. (Ed.), *Electronic government* (*Vol. 2739*, pp. 282–287). Springer. doi:10.1007/10929179_51

Benassi, P. (1999). TRUSTe: An online privacy seal program. *Communications of the ACM*, *42*(2), 56–59. doi:10.1145/293411.293461

Berendt, B., Günther, O., & Spiekermann, S. (2005). Privacy in e-commerce: Stated preferences vs. actual behavior. *Communications of the ACM*, *48*(4), 101–106. doi:10.1145/1053291.1053295

Bergmann, M., Rost, M., & Pettersson, J. S. (2006). Exploring the feasibility of a spatial user interface paradigm for privacy-enhancing technology. In A. G. Nilsson, R. Gustas, W. Wojtkowski, W. G. Wojtkowski, S. Wrycza & J. Zupancic (Eds.), *Bridging the gap between academia and industry* (pp. 437--448).

Borking, J. (2009). *Organizational motives for adopting privacy enhancing technologies (PETs)*. Paper presented at the D 7.3 PRISE Conference Proceedings: Towards Privacy Enhancing Security Technologies - The Next Steps.

Camenisch, J., et al. (2005). *Privacy and identity management for everyone*. Paper presented at the 2005 Workshop on Digital Identity Management.

CASPIAN. (2006). *Metro future store - Special report overview*. Retrieved March, 2011, from http://www.spychips.com/metro/overview.html

Catlett, J. (2000). *Open letter to P3P developers & replies*. Paper presented at the Tenth Conference on Computers, Freedom and Privacy: Challenging the Assumptions

Clarke, R. (2008). Business cases for privacy-enhancing technologies. In Subramanian, R. (Ed.), *Computer security, privacy and politics: Current issues, challenges and solutions* (pp. 135–155). Hershey, PA: IGI Global. doi:10.4018/978-1-59904-804-8.ch007

Clement, A., et al. (2008). *The PIPWatch toolbar: Combining PIPEDA, PETs and market forces through social navigation to enhance privacy protection and compliance*. Paper presented at the Technology and Society, 2008. ISTAS 2008. IEEE International Symposium on, http://dx.doi.org/10.1109/68.593394

Coyle, K. (1999). *P3P: Pretty poor privacy? - A social analysis of the platform for privacy preferences (P3P).* Retrieved April, 2011, from http://www.kcoyle.net/p3p.html

Coyle, K. (2000). *A response to "P3P and privacy: An update for the privacy community" by the Center for Democracy and Technology*. Retrieved April 2011 from http://www.kcoyle.net/response.html

Cranor, L. F. (2002). Help! IE6 is blocking my cookies. Retrieved April 2011 from http://www.oreillynet.com/pub/a/javascript /2002/10/04/p3p.html

Cranor, L. F. (2003). *"I didn't buy it for myself" privacy and ecommerce personalization*. Paper presented at the 2003 ACM Workshop on Privacy in the Electronic Society.

Cranor, L. F. (2008). P3P deployment on websites. *Electronic Commerce Research and Applications*, *7*(3), 274–293. doi:10.1016/j.elerap.2008.04.003

Cranor, L. F., Arjula, M., & Guduru, P. (2002). *Use of a P3P user agent by early adopters*. Paper presented at the 2002 ACM Workshop on Privacy in the Electronic Society.

Cranor, L. F., Guduru, P., & Arjula, M. (2006). User interfaces for privacy agents. *ACM Transactions on Computer-Human Interaction*, *13*(2), 135–178. doi:10.1145/1165734.1165735

Cranor, L. F., Langheinrich, M., & Marchiori, M. (2002). *A P3P preference exchange language 1.0 (APPEL1.0).* (Working Draft): W3C.

Cranor, L. F., & Lessig, L. (2002). *Web privacy with* (p. 3P). O'Reilly & Associates, Inc.

Egelman, S., Cranor, L. F., & Chowdhury, A. (2006). *An analysis of P3P-enabled web sites among top-20 search results*. Paper presented at the 8th International Conference on Electronic Commerce: The New E-commerce: Innovations for Conquering Current Barriers, Obstacles and Limitations to Conducting Successful Business on the Internet.

EU. (2002). *Directive on privacy and electronic communications*. 2002/58/EC

EuroPriSe. (2010). *Welcome! - EuroPriSe - European privacy seal*. Retrieved March, 2011, from https://www.european-privacy-seal.eu/

Facebook. (2011). *Welcome to Facebook - Log in, sign up or learn more*. Retrieved April, 2011, from http://www.facebook.com

Flinn, S., & Lumsden, J. (2005). *User perceptions of privacy and security on the Web*. Paper presented at the Third Annual Conference on Privacy, Security and Trust (PST 2005)

FTC. (2010). *Protecting consumer privacy in an era of rapid change: A proposed framework for businesses and policymakers*. Federal Trade Commission. Retrieved April 2011 from http://www.ftc.gov/os/2010/12/101201 privacyreport.pdf

Gomez, J., Pinnick, T., & Soltani, A. (2009). *KnowPrivacy*. Retrieved April 2011 from http://www.knowprivacy.org/report/ KnowPrivacy_Final_Report.pdf

Greenstadt, R., & Smith, M. D. (2005). *Protecting personal information: Obstacles and directions*. Paper presented at the Fourth Workshop on the Economics of Information Security (WEIS05)

Guarda, P., & Zannone, N. (2009). Towards the development of privacy-aware systems. *Information and Software Technology*, *51*(2), 337–350. doi:10.1016/j.infsof.2008.04.004

Hansen, M. (2009). *Putting privacy pictograms into practice - A European perspective*. Paper presented at the GI Jahrestagung.

Hochheiser, H. (2002). The platform for privacy preference as a social protocol: An examination within the U.S. policy context. *ACM Transactions on Internet Technology*, *2*(4), 276–306. doi:10.1145/604596.604598

IPTF. (2010). *Commercial data privacy and innovation in the internet economy: A dynamic policy framework*: The Department of Commerce, Internet Policy Task Force. Retrieved April 2011 from http://www.ntia.doc.gov/reports/2010/IPTF_Privacy_GreenPaper_12162010.pdf

ISO/IEC 15408-1. (2009). *Information technology -- Security techniques -- Evaluation criteria for IT security -- Part 1: Introduction and general model*. Retrieved April 2011, from http://www.commoncriteriaportal.org/files/ ccfiles/CCPART1V3.1R3.pdf, Retrieved

Jensen, C., Potts, C., & Jensen, C. (2005). Privacy practices of internet users: Self-reports versus observed behavior. *International Journal of Human-Computer Studies*, *63*(1-2), 203–227. doi:10.1016/j.ijhcs.2005.04.019

Kelley, P. G., et al. (2009). *A nutrition label for privacy*. Paper presented at the 5th Symposium on Usable Privacy and Security.

Kobsa, A., & Teltzrow, M. (2005). *Contextualized communication of privacy practices and personalization benefits: Impacts on users' data sharing and purchase behavior. Privacy Enhancing Technologies* (*Vol. 3424*, pp. 329–343). Berlin, Germany: Springer.

Kolter, J., Kernchen, T., & Pernul, G. (2010). Collaborative privacy management. *Computers & Security*, *29*(5), 580–591..doi:10.1016/j.cose.2009.12.007

Levy, S. E., & Gutwin, C. (2005). *Improving understanding of website privacy policies with fine-grained policy anchors*. Paper presented at the 14th International Conference on World Wide Web.

LinkedIn. (2011). *Relationships matter*. LinkedIn. Retrieved April, 2011, from http://www.linkedin.com/

Lioudakis, G. V. (2007). A middleware architecture for privacy protection. *Computer Networks*, *51*(16), 4679–4696. doi:10.1016/j.comnet.2007.06.010

Nehf, J. P. (2007). Shopping for privacy on the Internet. *The Journal of Consumer Affairs*, *41*(2), 351–375.

NTIA. (2010). Information privacy and innovation in the internet economy. *Federal Register*. Retrieved April 2011 from http://www.federalregister.gov/articles/ 2010/12/21/2010-31971/information-privacy-and-innovation-in-the-internet-economy

Nyre, Å. A., Bernsmed, K., Bøe, S., & Pedersen, S. (2011). A server-side approach to privacy policy matching. Paper presented at the Sixth International Conference on Availability, Reliability and Security.

Odlyzko, A. (2003). *Privacy, economics, and price discrimination on the Internet*. Paper presented at the 5th International Conference on Electronic Commerce.

Odlyzko, A. (2007). *Privacy and the clandestine evolution of e-commerce*. Paper presented at the Proceedings of the Ninth International Conference on Electronic Commerce Out-Law. (2009). *Privacy policy tool failed because of browser rejection, says W3C lawyer*. Retrieved March, 2011, from http://www.out-law.com/page-10437

Pettersson, J. S., et al. (2005). *Making PRIME usable*. Paper presented at the Symposium on Usable Privacy and Security.

Pollach, I. (2007). What's wrong with online privacy policies? *Communications of the ACM*, *50*(9), 103–108. doi:10.1145/1284621.1284627

Reeder, R. W., et al. (2008). *A user study of the expandable grid applied to P3P privacy policy visualization*. Paper presented at the 7th ACM Workshop on Privacy in the Electronic Society.

Schunter, M., & Waidner, M. (2007). *Simplified privacy controls for aggregated services: Suspend and resume of personal data*. Paper presented at the 7th International Conference on Privacy Enhancing Technologies.

Schwartz, A. (2009). *Looking back at P3P: Lessons for the future*. Retrieved April 2011 from http://www.cdt.org/files/pdfs/P3P_Retro_Final_0.pdf

Spiekermann, S., Grossklags, J., & Berendt, B. (2001). *E-privacy in 2nd generation e-commerce: Privacy preferences versus actual behavior*. Paper presented at the 3rd ACM Conference on Electronic Commerce.

Szeto, M., & Miri, A. (2007). *Analysis of the use of privacy-enhancing technologies to achieve PIPEDA compliance in a B2C e-business model*. Paper presented at the Eighth World Congress on the Management of eBusiness, 2007. WCMeB 2007.

Tøndel, I. A., Nyre, Å. A., & Bernsmed, K. (2011). *Learning Privacy Preferences*. Paper presented at the Sixth International Conference on Availability, Reliability and Security.

TRUSTe. (n.d.). Privacy seals & services. Online Trust & Safety from TRUSTe. Retrieved February 17, 2011, from http://www.truste.org

US. (2011). Welcome to the U.S.-EU & U.S.-Swiss safe harbor frameworks. Retrieved April, 2011, from http://www.export.gov/safeharbor/

Vila, T., Greenstadt, R., & Molnar, D. (2003). *Why we can't be bothered to read privacy policies - Models of privacy economics as a lemons market*. Paper presented at the 5th International Conference on Electronic Commerce.

Warren, S. D., & Brandeis, L. D. (1890). The right to privacy. *Harvard Law Review*, *4*(5), 193–220. doi:10.2307/1321160

Wilton, R. (2009). What's happened to PETs? *Information Security Technical Report*, *14*(3), 146–153. doi:10.1016/j.istr.2009.10.010

ADDITIONAL READING

Clarke provides a largely non-technical treatise on the business motivations for deploying Privacy Enhancing Technologies (Clarke, 2008). A European perspective on such motivations is provided by Borking (2009), although it should be noted that we have found no independent confirmation of Borking's claim that an organization which is innovative with respect to implementation of advanced Identity and Access Management (IAM), will also be more likely to deploy PETs.

The P3P web site (http://www.w3.org/P3P/) provides a wealth of information on P3P, and also provides links to other P3P documentation. The book "Web Privacy with P3P" (Cranor & Lessig, 2002) explains the working of the P3P protocol in depth. The book is written by Lorrie Faith Cranor who was a co-author of the P3P and APPEL specifications, and who served as chair of the P3P specification working group at the W3C. Lorrie Faith Cranor and colleagues have thus been working with P3P from its inception, and provide an insider view of the technology (Cranor, 2002, 2003; Cranor, Arjula, et al., 2002; Cranor, et al., 2008; Cranor, et al., 2006; Cranor & Lessig, 2002). A more critical view of P3P is offered by Hochheiser (2002). Argyrakis et al. (2003) presents a thorough survey of PETs, which in spite of its age is still surprisingly relevant. For an up-to-date view of US policy work on electronic privacy, see the reports from the U.S. Federal Trade Commission and Department of Commerce (FTC, 2010; IPTF, 2010; NTIA, 2010).

KEY TERMS AND DEFINITIONS

P3P: Platform for Privacy Preferences.

Personal Information: Any information that can be stored, and associated with an identifiable person. Personal data includes, e.g., a person's full name, e-mail address, IP address, credit card number, digital identity, membership in groups, relationships to other people, financial data and purchasing history, and so on. Sometimes, the term "personal identifiable information (PII)" is used instead of personal information.

PET: Privacy Enhancing Technology.

Privacy Paradox: The fact that although most users claim to be concerned with privacy, their actions are usually anything but privacy-preserving.

Privacy Policy: A policy published by a service provider stating how the service provider will use personal information collected from users of the service.

Privacy: The right to be left alone. Can sometimes be achieved through confidentiality, but not always.

Reputation: The opinion or social evaluation of a group toward an entity such as a person, business or organization.

Risk: Probability of (unwanted) incident multiplied by the cost occurred when the incident manifests itself. Cost does not necessarily have to be a monetary value, but in order to simplify calculations, non-monetary costs (such as loss of reputation) are usually assigned a monetary value.

Security: Mechanisms that protect a system from its environment. Composed of the three parts Confidentiality, Integrity and Availability.

Trust: Expecting another party to behave as they say they will.