5-15-2012

PRIVACY-BY-DESIGN THROUGH SYSTEMATIC PRIVACY IMPACT ASSESSMENT - A DESIGN SCIENCE APPROACH

Marie Caroline Oetzel
Vienna University of Economics and Business

Sarah Spiekermann
Vienna University of Economics and Business

PRIVACY-BY-DESIGN THROUGH SYSTEMATIC PRIVACY IMPACT ASSESSMENT – A DESIGN SCIENCE APPROACH

Oetzel, Marie Caroline, Vienna University of Economics and Business, Institute for Management Information Systems, Augasse 2-6, 1090 Vienna, Austria, marie.oetzel@wu.ac.at

Spiekermann, Sarah, Vienna University of Economics and Business, Institute for Management Information Systems, Augasse 2-6, 1090 Vienna, Austria, sspieker@wu.ac.at

Abstract

A major problem for companies that develop and operate IT applications that process personal data of customers and employees is to ensure the protection of this data and to prevent privacy breaches. Failure to adequately address this problem can result in considerable reputational and financial damages for the company as well as for affected data subjects. We address this problem by proposing a methodology to systematically consider privacy issues in a step-by-step privacy impact assessment (so called 'PIA'). Existing PIA approaches lack easy applicability because they are either insufficiently structured or imprecise and lengthy. We argue that employing the PIA proposed in this article, companies will be enabled to realise a 'privacy-by-design' as it is now widely heralded by data protection authorities. In fact, the German Federal Office for Information Security (BSI) ratified the approach we present in this article for the technical field of RFID and published it as a guideline in November 2011. The contribution of the artefacts we created is twofold: First, we provide a formal problem representation structure for the analysis of privacy requirements. Second, we reduce the complexity of the privacy regulation landscape for practitioners who need to make privacy management decisions for their IT applications.

Keywords: Privacy impact assessment, privacy-by-design, security risk assessment, design science.

1 Introduction

Privacy maintenance and control is a social value deeply embedded in our societies. A global survey found that 88% of people are worried about who has access to their data; over 80% expect governments to regulate privacy and impose penalties on companies that don't use data responsibly (Fujitsu, 2010). At the same time, we witness an increasing occurrence of privacy breaches, including massive leakage of personal data to unauthorised parties. At fast pace, technical systems transition to allow for unprecedented levels of surveillance as they move from centralised data base systems to ubiquitous computing. These developments demand for new approaches to protect privacy.

One of these approaches is to require companies to conduct privacy impact assessments (PIAs) (see i.e. (EC, 2009)). Similar to established security risk assessments (ISO, 2008; NIST, 2002), the goal of a PIA is to make companies run through a systematic privacy risk assessment. In this way, they are supposed to identify organisational and technical privacy threats and choose controls that mitigate these. Typically, these assessments should be done early on in the development of an IT application, so that – following the principle of 'privacy-by-design'[1] – privacy enhancing techniques and measures can be pro-actively built into an application. Although, there are already some countries where PIA procedures are used (in particular for e-government services as well as highly sensitive areas, health and biometrics) (Wright et al., 2011), their adoption is very slow, especially in the private sector (Bennett and Bayley, 2007). This can be explained by the fact that until now PIAs are not mandatory. But even if PIAs were to become mandatory (as some scholars now argue (Wright, 2011)) a great challenge we see is that existing PIA methodologies lack easy applicability. They are cumbersome either because they are insufficiently structured or because they are imprecise and lengthy. As we will show below, none of them describe a step-by-step process a company could easily implement and integrate into its risk management processes.

To address these shortcomings of existing PIAs, we propose a set of new constructs and a novel methodology for systematically considering privacy issues in a step-by-step PIA. We extend prior work in this research area by introducing experiences and concepts from security risk assessments. We propose a new systematic methodology that aids practitioners in evaluating relevant privacy issues and realising privacy-by-design. The goal of our methodology is to complement existing risk management techniques and provide companies that develop and operate IT applications that process personal data with a more formal technique for analysing application-specific privacy requirements. To achieve this goal we adopted a design science research approach (Hevner et al., 2004; Gregor, 2006). Design science research involves the construction and evaluation of IT artefacts, constructs, models, methods, and instantiations, by which important organisational IT problems can be addressed. Our proposed set of artefacts includes constructs for representing and evaluating privacy requirements and a new methodology for systematically considering privacy issues in a step-by-step process. Together, they constitute a novel process for realising an effective PIA.

The remainder of the article proceeds as follows. The next section reviews the conceptual foundations of existing methods for PIAs. The third section outlines the PIA methodology we propose and defines our constructs, including the representation of privacy requirements in the form of privacy targets and qualitative evaluation techniques. The fourth section provides an evaluation of the utility of our proposed approach. In particular, we apply our PIA process model to the technology field of RFID where we tested it and established it through the German Federal Office for Information Security (BSI) as a guideline for the development of privacy-friendly RFID applications.

[1] http://privacybydesign.ca

2 Addressing privacy issues today: A review of the current knowledge base

The PIA methodology we present below (Section 3) is founded upon the critical reflection of already existing constructs and procedures. We reflect on how existing privacy compliance procedures, privacy principles and regulation as well as security risk assessments informed our PIA methodology.

2.1 Existing privacy compliance procedures

Privacy is a broadly defined concept that goes beyond data protection (Solove, 2006). It includes, for example, “the right to be let alone” (Warren and Brandeis, 1890), groups’ freedom to private speech and association (Raab and Wright, 2012), bodily privacy, etc. A first requirement for a PIA is therefore that it should be able to flexibly embrace the full spectrum of these concepts.

However, when it comes to current privacy compliance procedures (at least in Europe), typically these are reduced to data protection only. Data protection issues have a legal basis and are therefore subject to legal compliance checks conducted by national data protection authorities. Or, they are addressed by private auditing businesses, which offer privacy seals (i.e. (TRUSTe, 2011), (BBBOnLine, 2011), (EuroPriSe, 2011)). Yet, besides their limited focus on data protection, current compliance procedures face more challenges: First, they mostly take place at the end of the development of an application or even later when the application is already up and running. Thus, they review existing systems (Shroff, 2007) a change of which can only be fixed in a bolted-on and often costly fashion. With this they often fall short to respond to article 20 of the European Data Protection Directive, which demands “that these processing operations are examined prior to the start thereof.” (EC, 1995) Second, they are not done by engineers designing the system, but by auditors, lawyers or data protection officials who can just ‘checklist’ legal compliance, but hardly influence more ‘code-based’ and rigorous privacy controls. And third, current compliance checks lack a standard procedure, not only because national data protection laws vary, but also because so far data protection has not been perceived as part of companies’ quality controls that are ensured by standardised risk procedures.

As a consequence of this status quo and amounting public pressure for privacy protection, PIAs are now considered a key solution to offer a superior approach (EC, 2009). With their inherent risk management orientation they are integrating privacy considerations into the development of applications and thus enable privacy-by-design. Stewart (1996) describes a PIA as follows: "In large measure, PIAs are directed not simply towards issues of legal compliance but the policy choices involved in answering the questions ‘ought we to do this?”. Bennett and Bayley (2007) identified four common PIA requirements: (1) “conduct a prospective identification of privacy issues or risks before systems and programmes are put in place, or modified”, (2) “assess the impacts in terms broader than those of legal compliance”, (3) “be process rather than output oriented”, and (4) ”be systematic”.

Even though process orientation has been identified as a key element of a PIA, existing PIAs largely fall short of it. As Figure 1a demonstrates, even the UK PIA handbook (ICO, 2009) that is heralded as a global “best practice publication” on how to conduct a PIA (Wright et al., 2011) is far from anything that governance scholars would consider a valid process reference model. No input-output factors are described. Process steps are so generic (“information gathering”, “internal analysis”) that they leave persons responsible to conduct PIAs uninformed of what to do. No guidelines or conceptual tools support the risk assessment.

The first PIA with a short, but valid process model is contained in the PIA Framework for RFID (EC, 2011) that needed to be endorsed by the Art. 29 Working Party (which explicitly demanded a risk evaluation process!) and was signed by the European Commission in April 6th 2011. This framework requires European RFID application operators to describe their system landscape, identify privacy risks, then mitigate such risks through appropriate controls and finally document the analysis and residual risks in a PIA report. The procedure outlined in the PIA Framework for RFID has received a lot of regulator attention as it resulted from a US-European negotiation (Spiekermann, 2012). It is voiced as a "landmark for privacy-by-design" by Ontario's DPA Anne Cavoukian who invented the concept of privacy-by-design. Yet, again it has its own shortcomings: As the related field of security risk management demonstrates, risk assessments need to be very concrete and systematic. This is demonstrated in Figure 1b depicting the NIST risk assessment process (NIST, 2002). All steps in this process build on each other in a concrete and interlocked manner. Against this background we reviewed existing security risk processes to inform the creation of a new PIA.

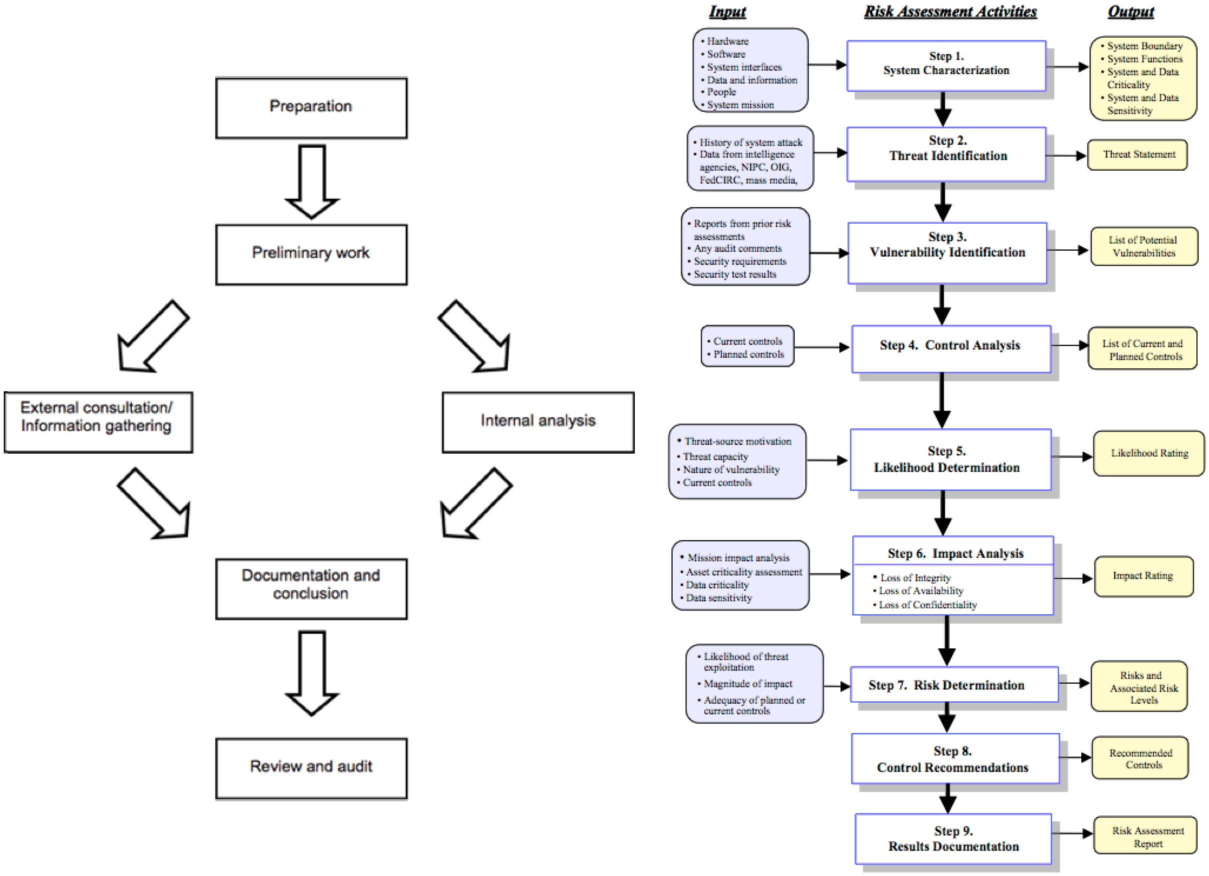
Preparation
Preliminary work
External consultation/ Information gathering
Internal analysis
Documentation and conclusion
Review and audit
Input
Risk Assessment Activities
Output
• Hardware
• Software
• System interfaces
• Data and information
• People
• System mission
Step 1. System Characterization
• System Boundary
• System Functions
• System and Data Criticality
• System and Data Sensitivity
• History of system attack
• Data from intelligence agencies, NIPC, OIG, FedCIRC, mass media,
Step 2. Threat Identification
Threat Statement
• Reports from prior risk assessments
• Any audit comments
• Security requirements
• Security test results
Step 3. Vulnerability Identification
List of Potential Vulnerabilities
• Current controls
• Planned controls
Step 4. Control Analysis
List of Current and Planned Controls
• Threat-source motivation
• Threat capacity
• Nature of vulnerability
• Current controls
Step 5. Likelihood Determination
Likelihood Rating
• Mission impact analysis
• Asset criticality assessment
• Data criticality
• Data sensitivity
Step 6. Impact Analysis
• Loss of Integrity
• Loss of Availability
• Loss of Confidentiality
Impact Rating
• Likelihood of threat exploitation
• Magnitude of impact
• Adequacy of planned or current controls
Step 7. Risk Determination
Risks and Associated Risk Levels
Step 8. Control Recommendations
Recommended Controls
Step 9. Results Documentation
Risk Assessment Report

Figure 1. a. UK PIA process overview (ICO, 2009) b. NIST process overview (NIST, 2002)

2.2 Security risk assessments

Standards and guidelines for information security management in organisations are already available for some time. The most prominent are the ISO/IEC 27000 series and NIST Special Publications 800 series. In Germany, the Federal Office for Information Security (BSI) provides industry with an IT Baseline Protection Catalog ("BSI IT-Grundschutz"). The latter not only complies with ISO/IEC 27000, but describes even stricter requirements. Similar to the NIST example, all of these standards contain an extensive number of interlocked steps.

Most important, all of these accepted standards offer guidelines that can be integrated into an organisation's risk management processes (see: (ISO, 2008), (NIST, 2002), (BSI, 2008)). Contrary to the PIA process (Figure 1a), the NIST risk assessment (Figure 1b) process not only depicts the needed steps in much more detail but it also defines input and output artefacts for each step. Both, the detailed steps as well as the required artefacts enable a company to realise the risk assessment in a coherent way, knowing exactly what to do. In addition, these standardised security risk procedures have defined interfaces with system development lifecycles. This implies that security issues are already considered early on during the development and implementation of IT applications. The approach decreases bolted-on security functionality and promotes security-by-design.

Nevertheless, researchers describe the following problems that are inherent to existing security risk assessments and that we consequently need to keep in mind for our proposed methodology: focus on process and not on content and its quality (Siponen, 2006), focus on generic security requirements and thus disregard of company-specific requirements (Siponen and Willison, 2009), validation based on common practice and not on profound research methods (Siponen and Willison, 2009).

Seen that privacy and security are interrelated but still distinct and not synonymous (Oetzel and Krumay, 2011), it is not surprising that both ISO/IEC 27002 and BSI IT-Grundschutz, include privacy protection. Yet, despite this claim the ISO standard leaves privacy policies and measures unspecified. The BSI IT-Grundschutz does apply the security risk analysis to privacy principles. Yet, in doing so it reduces privacy protection to the concepts of anonymity, pseudonymity, unobservability and unlinkability (BSI, 2008). Thus, BSI (2008) does demonstrate how their security risk assessment method works for privacy, but it fails to embrace the wider spectrum of privacy principles as they are imbedded in the European Data Protection Directive or in the OECD privacy guidelines (i.e. data subjects' right to transparency, collection limitation, etc.). Privacy principles that should be embraced by PIAs are reflected in the next section.

2.3 Privacy principles and data protection regulation

At the outset of any privacy analysis or PIA should be the question of what it actually is that needs to be protected (Rost, 2011). The oldest description of information privacy principles is the OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (OECD, 1980) that was adopted in 1980. These guidelines differentiate between principles of national and international application. Principles of national application are: collection limitation, data quality, purpose specification, use limitation, security safeguards, openness and individual participation. Principles of international application are free flow and legitimate restrictions. The goal of these guidelines was to harmonise national laws and to motivate member states to include a certain degree of privacy protection into their national laws. The European Data Protection Directive hence mirrors many of these principles (Greenleaf, 2011): data quality, legitimate processing of personal data, legitimate processing of personal sensitive data, the data subject's right to be informed, the data subject's right of access to data, to correct and erase data, the data subject's right to object, confidentiality and security of processing and notification. The principles show that the European Data Protection Directive explicitly differentiates between personal data and sensitive personal data. What is important to consider is that both of these major privacy frameworks focus only on information privacy issues. As was noted above, however, privacy is actually a wider concept. Therefore, if PIAs are to embrace the whole privacy arena, then protection goals will probably need to go beyond current legal frameworks and agreements. And indeed this is what thought leaders currently argue for when they propose that PIAs should extend beyond their current limited focus on individual privacy and embrace privacy as a social and political construct (Raab and Wright, 2012).

3 PIA methodology and constructs

Against the background of the reviewed knowledge base we can now present the PIA methodology we propose (Figure 2). For the reasons outlined above, the methodology is initially founded on (BSI, 2008). Each step produces a result artefact (depicted in the grey boxes on the right).

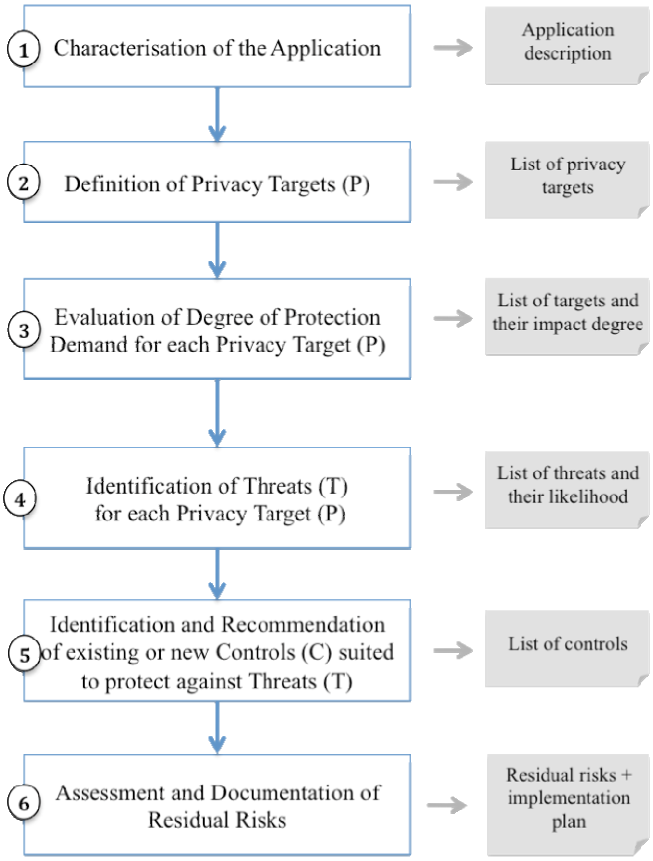
1 Characterisation of the Application
Application description
2 Definition of Privacy Targets (P)
List of privacy targets
3 Evaluation of Degree of Protection Demand for each Privacy Target (P)
List of targets and their impact degree
4 Identification of Threats (T) for each Privacy Target (P)
List of threats and their likelihood
5 Identification and Recommendation of existing or new Controls (C) suited to protect against Threats (T)
List of controls
6 Assessment and Documentation of Residual Risks
Residual risks + implementation plan

Figure 2. PIA methodology overview

3.1 Step 1 – Characterisation of the application

The goal of this first step is to clearly identify the scope as well as the boundaries and thus the assets (resources and information) that need to be protected and thus considered throughout the impact assessment. The comprehensiveness and the level of detail of the characterisation have an influence on the successful execution of a PIA. Risks may not be identified due to missing information in the application characterisation. The characterisation should therefore contain: application and system components, roles, generic business processes, detailed use cases, data flow diagrams (of internal and external data flows) as well as categories of processed data. This information can either be derived from requirements and design documents when the application is still in the initiation, design or development phase. If the application is already operational, relevant information needs to be collected from its production environment. Thus, information gathering is not restricted to a specific phase but it can be conducted throughout the impact assessment process. As a result the application characterisation gains more and more detail so that the risk evaluation can be based on a sound foundation.

3.2 Step 2 – Definition of privacy targets

The purpose of a risk assessment is to understand what is at risk. Existing security risk assessments take the application characterisation from step 1 as a basis and then identify the assets described therein as security targets that need to be protected. The above-described security risk assessment standards offer lists of security targets that can be used in the assessment of an application. Such a guideline does not yet exist for the consideration of privacy issues. Many catalogues of potential privacy targets exist (Rost, 2011). Consequently one of our constructs is a list of systematically derived privacy targets from the relevant legal frameworks. Taking EU legislation and thus the European Data Protection Directive (EC, 1995) as a starting point, results in 8 privacy targets and 16 more concrete sub targets that can be specified as depicted in Table 1.

Privacy target code and name		Sub targets	
P1	Safeguard of quality of personal data	P1.1	Ensuring fair and lawful processing through transparency
		P1.2	Providing purpose specification and limitation
		P1.3	Ensuring data avoidance and minimisation
		P1.4	Ensuring quality of data
		P1.5	Ensuring limited duration of data storage
P2	Legitimacy of processing personal data	P2.1	Legitimacy of processing personal data
P3	Legitimacy of processing sensitive personal data	P3.1	Legitimacy of processing sensitive personal data
P4	Compliance with the data subject's right to be informed	P4.1	Providing adequate information in cases of direct collection of data from the data subject
		P4.2	Providing adequate information where the data has not been obtained directly from the data subject
P5	Compliance with the data subject's right to access, correct and erase data	P5.1	Facilitating the provision of information about processed data and purpose
		P5.2	Facilitating the rectification, erasure or blocking of data
		P5.3	Facilitating the notification to third parties about rectification, erasure and blocking of data
P6	Compliance with the data subject's right to object	P6.1	Facilitating the objection to the processing of personal data, direct marketing activities and disclosure of data to third parties
		P7.1	Facilitating the objection to being subject to decisions that are solely based on automated processing of data
P7	Safeguard of confidentiality and security of processing	P7.1	Safeguarding confidentiality and security of processing
P8	Compliance with notification requirements	P8.1	Compliance with notification requirements

Table 1. Generic privacy targets and concrete sub targets

In order to address the second of the above-mentioned (Subsection 2.2) problems of existing security risk assessments, at the outset of this second step, each privacy target needs to be described against the background of the respective industry or company context. Seen the wide spectrum of the privacy concept as well as national laws or industry-specific regulations more targets can and should be added.

3.3 Step 3 – Evaluation of degree of protection demand for each privacy target

The leading question of this third step is "What would happen if …?". The goal is to identify the degree of protection demand that is feasible for each privacy target. Damage scenarios that consider the impact that would result from potential privacy breaches are used. This evaluation of privacy targets differs from the evaluation of security targets in that we focus less on the loss of assets, but consider more of the 'soft' implications that privacy breaches have. As we are dealing with applications that process personal information of customers or employees, it is not sufficient to only consider the operator's perspective and reflect on reputational or financial damages that can ensue for an operator, but it is essential to also consider the data subject's perspective and to reflect on damages that can ensue for the data subject (e.g. reputation, financial well being, personal freedom). Table 2 describes these two perspectives and adds three degrees of protection demand (low, medium and high) that can help to systematically evaluate the degree of protection demand for a privacy target.

Protection demand	Criteria for the assessment of protection demand				
	Application operator perspective		Data subject perspective		
	Impact on reputation and brand value	Financial loss	Social standing, reputation	Financial well being	Personal freedom
Low – 1	The impact of any loss or damage is **limited** and calculable.				
Medium – 2	The impact of any loss or damage is **considerable**.				
High – 3	The impact of any loss or damage is **devastating**.				

Table 2. Protection demand categories and perspectives

For each privacy target, the evaluations of the five criteria are then combined using the maximum principle resulting in an overall evaluation. In a later state (step 5) of the assessment, this evaluation helps to choose privacy controls that are corresponding in strength and vigour.

3.4 Step 4 – Identification of threats for each privacy target

Based on the defined privacy targets, this step aims at systematically deducing threats for each of the privacy targets. These threats can either be generic in terms of the privacy target, technology-, application- or context-specific. After identifying the threats, it is necessary to consider the likelihood of their occurrence. We differentiate only two levels: likely and not likely. Not all threats may be equally probable. Some may not materialise at all from a specific operator's perspective or in a specific operating context. Only those threats will later be mitigated that are likely to occur.

3.5 Step 5 – Identification and recommendation of existing or new controls suited to protect against threats

The crucial step in the assessment process is to identify controls that can help to minimise, mitigate or eliminate the identified threats. Controls can either be of a technical or non-technical nature. Technical controls are directly incorporated into a system, e.g. access control and authentication mechanisms; pseudonymisation, anonymisation and encryption methods. Non-technical controls, on the other hand, are management and operational controls as well as accountability measures, e.g. policies or operational procedures and information measures taken vis-à-vis data subjects. Furthermore, controls can be categorised as being either preventive or detective. Preventive controls inhibit violation attempts, while detective controls warn operators about violations or attempted violations. Keeping in mind that PIAs have the goal to foster privacy-by-design, there should be a focus on identifying and recommending preventive controls.

For each threat, controls need to be identified and for each control, three levels of effectiveness (low, medium and high) need to be defined. These levels need to be taken into account when recommending a control, because the control's level of effectiveness should match the beforehand-identified (in step 3) degree of protection demand. E.g. high protection demands combined with likely threats should be mitigated with highly effective controls.

3.6 Step 6 – Assessment and documentation of residual risks

The list of recommended controls that results from step 5 needs to be considered during the now following risk mitigation phase. Recommended controls need to be evaluated, e.g. concerning feasibility and effectiveness, a cost-benefit analysis can be conducted, which then results into a list of prioritised controls. The result is a control implementation plan, from which residual risks are derived. Residual risks remain for example if an implemented control reduces the magnitude of the impact of a threat but does not eliminate the threat completely due to technical or business reasons.

4 Utility evaluation of the proposed artefacts

An important aspect of design science research is the evaluation of the proposed artefacts. To demonstrate the utility of our artefacts, we follow Hevner et al. (2004), who suggested five evaluation methods, two of which are appropriate for the context we have studied. The first is the observational approach, which is exemplified by case study and interviewing methods. In subsection 4.1, we report the results of several workshops with IT industry experts who assessed the use of the proposed artefacts. We additionally use the second evaluation method, the descriptive approach, by employing the informed argument method using information from the knowledge base of our research domain to build arguments for the utility of our proposed artefacts. In subsection 4.2, we again use the descriptive approach by employing the scenario method and report the results of three scenarios and their resulting PIAs. With the completion of these evaluation methods, we also address the above-mentioned (Subsection 2.2) problems of existing security risk assessments. We use profound research methods to evaluate content and quality of our proposed artefacts and not only rely on common practice.

4.1 Workshops with industry experts

Conducting interviews, in our case in the form of workshops, is one of the most important gathering tools in qualitative research (Myers and Newman, 2007). Each workshop was structured as follows: (1) we explained our PIA methodology and the constructs, (2) the experts conducted a PIA based on an exemplary scenario we had prepared in advance, and thus evaluated the utility of our proposed artefacts in a context of use, (3) the experts were asked to suggest potential improvements that might be appropriate to our proposed artefacts. We ensured that all three parts of the workshop were completed; however related topics of discussion were permitted in order to increase the richness of the information captured. The workshops took on average 4 to 6 hours. The information was captured in the form of result protocols.

We worked with a set of industry experts with participants from 5 distinct stakeholder groups: (1) general risk manager, (2) IT department manager, (3) technology innovations manager with a strong background in technical security management, (4) member of a governmental institution, which is focused on information security, with a strong background in theoretical risk management and mathematics and (5) academic researchers. We worked with a total of 7 experts.

In general, all participants found value in our proposed artefacts for conducting a systematic and step-by-step PIA. The following three key dimensions about the utility of our proposed artefacts consistently emerged in the participants' opinions:

Documentation comprehensiveness: The comprehensive documentation of an application characterisation as required in step 1 is considered to be necessary in terms of a successful impact assessment but also expensive in terms of the amount of time and labour that needs to be invested, at least at the beginning. Participants agree that such comprehensive documentation is not readily available in a typical company where the main interest generally lies in a running application and not in a well-documented one. Furthermore, some participants were concerned about publishing such detailed company internal information to external parties like data protection authorities or customers. As the concept of a PIA indeed does recommend publishing a PIA report, which can be one of the resulting artefacts of step 6, it does not specify the level of detail that should be published. Thus, it seems to be acceptable to create a PIA report for the public that does not contain all internal information especially no confidential information, but enough information so that an external party can comprehend the decisions that were taken throughout the PIA.

Complexity reduction: All of the participants highly valued the given privacy targets, especially the 16 concrete sub targets (see Table 1). The privacy targets systematically structure the confusing (because of the legal language) and extensive landscape of privacy requirements in a way that practitioners feel confident to work with. Nevertheless, it remained the problem of the 'correct' interpretation of some of the targets. Especially the target P1.1 'Ensuring fair and lawful processing through transparency' resulted in discussions on how to interpret 'transparency' and what might be the measures to ensure transparency. Interestingly, this problem of 'correct' interpretation was not considered to be insurmountable; discussions always lead to a certain interpretation and participants agreed to consult legal personnel if necessary.

Qualitative evaluation: The workshops lead to the conclusion that in both steps (3 and 4) where an evaluation is needed, a qualitative approach is the most feasible for practitioners. First there is the evaluation of the degree of protection demand for each of the privacy targets. As already indicated above, the evaluation of the impact that results from a privacy breach is different from evaluating the impact of security breaches (e.g. a loss of availability might be easily quantifiable in terms of business losses). The two perspectives (operator and data subject; see Table 2) we proposed were considered to be very helpful to evaluate the 'soft' factors that are typical for impacts of privacy breaches. Second there is the evaluation of the likelihood of each identified threat. Participants heavily discussed whether it should not be possible to assign a quantitative probability to each threat, but the nature of most of the threats did not make this a feasible approach. Thus, for the time being we settled on the simple differentiation between likely and unlikely.

4.2 Three scenarios and exemplary PIAs

The scenario method of the descriptive approach is described as follows: "Construct detailed scenarios around the artefact to demonstrate its utility." (Hevner et al., 2004) We translated the above described methodology and constructs, which are generally applicable, to a 'PIA guideline for RFID applications' (BSI, 2011) thus targeting specific IT applications, namely those who use RFID technology. The herein defined privacy targets and protection demand categories were taken as is but a additional list of 60 threats and 27 controls was compiled, which are feasible in the context of RFID applications. Additionally, we constructed three comprehensive scenarios: (1) a retail scenario involving an RFID-enabled loyalty card and tagged products, (2) a public transport scenario using RFID-enabled tickets and pay-per-use models, and (3) an automotive scenario involving an RFID controlled assembly and an RFID-enabled employee access card (BSI, 2011). All three scenarios were developed with the help of industry partners in order to describe business and use cases that are realistic. We then conducted exemplary PIAs for all of the three scenarios using our described methodology and constructs and documented the six steps in detail.

The accomplishment of the exemplary PIAs leads to several key findings. These key findings have consequences for the design of the RFID applications and for related business processes. Thus, in all three scenarios, the findings of a systematic PIA lead to system design, function and process adaptations in terms of privacy-by-design.

Retail scenario: In short, this scenario is composed of an RFID-enabled loyalty card, tagged products, RFID-enabled shop-floor applications (e.g. smart-trolley, smart-shelf, self-checkout) and added-value services (especially for expensive goods). The consideration of P1.1 and its related threats results in controls that require a lot of effort to extensively inform customers about the involved technology, the customer data that is collected and how this data is processed, so that customers can make informed decisions. Furthermore, it is required to separate logistical from customer data (P1.2 and P1.3), to implement fine-grained access rights and to regularly update assigned access rights (P1.2 and P7.1), to implement deletion rules that guarantee that customer data that is no longer needed for the specified purpose is deleted or anonymised (P1.5 and P1.3), to offer personalised as well as pseudonymised loyalty cards to customers (P1.3), to kill all product tags during checkout and to not kill a product tag if a customer utters the explicit wish to use the (expensive) product later on in conjunction with added-value services (P1.2 and P7.1).

Public transport scenario: In short, this scenario is composed of a ticket (e.g. a dedicated RFID-enabled card, a multi-application card or an NFC mobile device), entitlements that can be loaded onto this ticket, vehicles and gates of the public transport system that automatically read these tickets. The consideration of P1.3 and its related threats results in controls that require the public transport operator to not only offer personalised and pseudonymised tickets but also anonymised tickets, so that customers can still use public transport in an anonymous way. Furthermore, it is required to implement and regularly update fine-grained access rights to the collected data (e.g. in order to prevent the disclosure of customers' movement profiles) (P1.2 and P7.1), to implement deletion rules that guarantee that customer data that is no longer needed for the specified purpose is deleted (P1.5 and P1.3), to provide extensive information material so that customers can make informed decisions (P1.1), to ensure that customer data is correct and up-to-date so that monthly or best-price bills are correct (P1.4).

Automotive scenario: This scenario is twofold; in-house and outbound logistics of a car manufacturer are managed with the help of RFID-technology (e.g. car bodies and security-relevant/upscale modules are tagged) and access control to the facilities is managed with the help of an RFID-enabled employee card. The consideration of P1.3 and its related threats results in controls that require the car manufacturer and the car dealer to kill all tags before handing the car over to the customer. If some of the tags (e.g. on security-relevant modules) are intended to remain activated for defect management and recall purposes, these tags need to be cryptographically secured. In the case of the employee card, not only personal but also sensitive personal data is processed because the human resource management system is involved in the personalisation of the cards and in the assignment of access rights to individual employees. This leads to strong requirements for the implementation of access rights to the processed data (P1.2 and P7.1) and the assurance of the legitimacy of processing of this data (P2.1 and P3.1).

5 Conclusions, limitations and future work

Following the design science research paradigm, the major theoretical contribution of this research is the development of a new set of artefacts designed to help practitioners and researchers to understand the relevant privacy regulation landscape better and to analyse and assess privacy issues in a systematic step-by-step process. Specifically, the artefacts provide tools for representing privacy requirements in the form of privacy targets, evaluating the degree of protection demand of these targets and to systematically derive threats and adequate controls. We build on prior risk assessment experiences and research especially in the security risk assessment area. We evaluated the proposed artefacts using two methods that are asked for in theoretical design science research. First, we used qualitative workshops with IT industry experts to assess the use of the proposed artefacts by practitioners. Second, we used scenario construction to demonstrate the applicability of our artefacts in the context of three realistic RFID application scenarios.

In this research we focused on the development of the methodology and the supporting constructs and we did not consider the integration of our step-by-step methodology into the existing risk management process landscape of a company. Although, we explicitly chose to base our methodology on existing security risk assessments to facilitate such an integration, we did not yet examine if this is actually the case.

To continue to provide useful tools to practitioners we already implemented an instantiation of our artefacts in the form of a web application and plan to do case studies to further evaluate the utility of our proposed artefacts.

References

BBBOnLine (2011). BBBOnLine – BBB Accredited Business Seal. http://www.bbb.org/online/, accessed December 5, 2011.

Bennett, C. and Bayley, R. (2007). Privacy Impact Assessments: International Study of their Application and Effects, Loughborough University, UK.

BSI (2008). Risk Analysis on the Basis of IT-Grundschutz, BSI Standard 100-3.

BSI (2011). Privacy Impact Assessment Guideline for RFID Applications.

EC (1995). Directive 95/46/EC on the protection of individuals with regard to the processing of personal data and on the free movement of such data. Official Journal of the European Communities, L 281, 31-50.

EC (2009). Commission recommendation on the implementation of privacy and data protection principles in applications supported by radio-frequency identification.

EC (2011). Privacy and Data Protection Impact Assessment Framework for RFID Applications.

EuroPriSe (2011). EuroPriSe – European Privacy Seal. https://www.european-privacy-seal.eu/, accessed December 5, 2011.

Fujitsu (2010). Personal data in the cloud: A global survey of consumer attitudes. Fujitsu, Japan.

Greenleaf, G. (2011). Global data privacy in a networked world. Research Handbook of the Internet. Cheltenham, Edward Elgar.

Gregor, S. (2006). The Nature of Theory in Information Systems. MIS Quarterly, 30 (3), 611-642.

Hevner, A. R., March, S. T., Park, J., Ram, S. (2004). Design Science in Information Systems Research. MIS Quarterly, 28 (1), 75-105.

ISO (2008). ISO/IEC 27005 Information technology – Security techniques – Information security risk management.

Myers, M. D. and Newman, M. (2007). The Qualitative Interview in IS Research: Examining the Craft. Information and Organization, 17 (1), 2-26.

NIST (2002). Risk Management Guide for Information Technology Systems, NIST Special Publication 800-30.

Oetzel, M. C. and Krumay, B. (2011). Differentiating Privacy and Security: A Content Analysis of B2C Websites. AMCIS 2011 Proceedings, Paper 211.

OECD (1980). Guidelines on the protection of privacy and transborder flows of personal data.

Raab, C. and Wright, D. (2012). Surveillance: Extending the Limits of Privacy Impact Assessments. Privacy Impact Assessment: Engaging Stakeholders in Protecting Privacy. Wright, D. and De Hert, P.. Dodrecht, Springer Verlag.

Rost, M. (2011). Datenschutz in 3D. DUD - Datenschutz und Datensicherheit, 5, 351-354.

Shroff, M. (2007). Privacy Impact Assessment Handbook. Office of the Privacy Commissioner, Auckland, New Zealand.

Siponen, M. (2006). Information Security Standards – Focus on the Existence of Process, Not Its Content. CACM, 49 (8), 97-100.

Siponen, M. and Willison, R. (2009). Information security management standards: Problems and solutions. Information & Management, 46, 267-270.

Solove, D. J. (2006). A Taxonomy of Privacy. University of Pennsylvania Law Review, 154 (3), 477-560.

Spiekermann, S. (2012). The RFID PIA - Developed by Industry, Agreed by Regulators. Privacy Impact Assessment: Engaging Stakeholders in Protecting Privacy. Wright, D. and De Hert, P.. Dodrecht, Springer Verlag.

Stewart, B. (1996). Privacy Impact Assessments. Privacy Law and Policy Reporter, 3 (4), Article 39.

TRUSTe (2011). TRUSTe privacy seal. http://www.truste.com/, accessed December 5, 2011.

UK Information Commissioners Office (ICO) (2009). Privacy Impact Assessment Handbook. London.

Warren, S. D. and Brandeis, L. D. (1890). The Right to Privacy. Harvard Law Review, 4 (5), 193-220.

Wright, D. (2011). Should Privacy Impact Assessments Be Mandatory? CACM, 54 (8), 121-131.

Wright, D., Wadhwa, K., De Hert, P., Kloza, D. (2011). A Privacy Impact Assessment Framework for data protection and privacy rights. Deliverable D1 of the EU PIAF Project. Brussels.