EnCoRe: Towards a holistic approach to privacy

Nick Papanikolaou, Sadie Creese, Michael Goldsmith, Marco Casassa Mont, Siani Pearson

HP Laboratories
HPL-2010-83

Keyword(s):

privacy policies, policy hierarchy, policy refinement

Abstract:

Privacy requirements for IT systems and solutions arise from a variety of sources, including legislation, sector-specific regulation, organisational guidelines, social and user expectations. In this paper we present and discuss a holistic approach to the management of privacy - explored in the context of the EnCoRe project - which takes into account the need to deal with these different types of policies, at different levels of abstraction as well as risk assessment methods to assess them based on specific threats, needs and constraints. We discuss examples of privacy requirements and related policies coming from different sources. We then present how a 'privacy-aware risk assessment' approach (which leverages and extends traditional security-driven risk assessment approaches) can be used to analyse these policies, assess their compliance to requirements, identify gaps and mandate the adoption of specific controls. We explain its relevance and implications in an employee data case study, involving the management of privacy consent and revocation. This is work in progress, carried out in the context of the EnCoRe collaborative project [1].

External Posting Date: July 21, 2010 [Fulltext] Approved for External Publication
Internal Posting Date: July 21, 2010 [Fulltext]

ENCORE: TOWARDS A HOLISTIC APPROACH TO PRIVACY

Nick Papanikolaou, Sadie Creese, Michael Goldsmith
International Digital Laboratory, University of Warwick, U.K.
N.Papanikolaou@warwick.ac.uk, S.Creese@warwick.ac.uk, M.H.Goldsmith@warwick.ac.uk

Marco Casassa Mont, Siani Pearson
Systems Security Lab, Hewlett Packard Laboratories, Bristol, U.K.
marco_casassa-mont@hp.com, siani.pearson@hp.com

Keywords: privacy policies, policy hierarchy, policy refinement

Abstract: Privacy requirements for IT systems and solutions arise from a variety of sources, including legislation, sector-specific regulation, organisational guidelines, social and user expectations. In this paper we present and discuss a holistic approach to the management of privacy - explored in the context of the EnCoRe project - which takes into account the need to deal with these different types of policies, at different levels of abstraction as well as risk assessment methods to assess them based on specific threats, needs and constraints. We discuss examples of privacy requirements and related policies coming from different sources. We then present how a 'privacy-aware risk assessment' approach (which leverages and extends traditional security-driven risk assessment approaches) can be used to analyse these policies, assess their compliance to requirements, identify gaps and mandate the adoption of specific controls. We explain its relevance and implications in an employee data case study, involving the management of privacy consent and revocation. This is work in progress, carried out in the context of the EnCoRe collaborative project [1].

1 INTRODUCTION

As part of everyday business practice, enterprises manage and administer huge databases of personal data. This process involves meeting a wide range of privacy requirements, including data protection laws and the privacy preferences of individual users.

Privacy requirements emerge from several sources within and outside an enterprise, causing tensions and tradeoffs that need to be balanced. For instance, it is often of commercial interest to an enterprise to use personal data held in order to market its products; this is an example of a business need that needs to be balanced against relevant legislation, and the privacy preferences of individual customers. Privacy advocacy groups will apply pressures to enterprises to prevent collection of personal data, ensure that collected data is not misused, and/or ensure that customers have access and control over data held about them. Individual customers will often expect to have such control, as information held about them may have a direct influence on their livelihoods (e.g. in the case of creditworthiness ratings). So, there are tensions that arise between different countervailing interests, and any enterprise whose business relies on the collection and processing of personal data has to deal with pressures from different sources, beyond just what is stipulated by common law. These pressures, whether they result from the law or not, can have direct financial consequences making privacy a significant risk to data-subjects and enterprises handling personal data alike; when data is mishandled this can result in financial penalties, when there is even a perception that it is mishandled this could result in a fall in the share value for commercial enterprises.

Experience within the EnCoRe project ("Ensuring Consent and Revocation") [1] has shown that enterprises often adopt privacy practices as an afterthought, rather than including privacy considerations as an integral part of all business processes. Also, it can be the case that different departments in the very same enterprise, e.g. the legal and IT departments, view privacy requirements differently, potentially resulting in resolutions which do not make privacy a core requirement in system design. In an ideal world there would be a way of relating high-level privacy requirements (such as those found in privacy law, viz. the Data Protection Act in the UK) to the business and security needs of an enterprise, so as to check for potential conflicts, and enforce all applicable requirements in a consistent way in favour of preserving privacy.

In order to develop a suitable strategy for handling privacy requirements in an enterprise, it is necessary to perform risk assessments, so as to determine:

(i) which risks need to be mitigated,
(ii) at which points in the enterprise's business processes special privacy controls need to be introduced,
(iii) what mechanisms are needed to protect infrastructure and data.

But such assessments can only be complete if there is a good understanding of all the privacy requirements that exist, and this is difficult to achieve due to the tensions between different interests mentioned above, and the lack of a uniform representation.

Spiekermann and Cranor [10] have identified approaches to privacy enforcement in an enterprise as belonging to three broad classes: (i) *privacy by architecture*, in which privacy requirements are embodied by the very design of a company's IT infrastructure (e.g. through the use of anonymous credentials) and in some cases meaning that additional enforcement measures are unnecessary since there is no collection of personal data at all, (ii) *privacy by policy*, in which privacy requirements are met through audits and automated enforcement of policies, and (iii) *hybrid approaches*, which are various combinations of (i) and (ii). Not all policies can be enforced automatically, and certainly there exist several interrelated policies, which are applied at different levels of management.

Privacy by architecture corresponds to a scenario that may not always be realistic; there are actually many valid business reasons for data collection, in which case policies are needed to enforce privacy. Privacy by architecture can only be introduced before an enterprise's IT infrastructure is built, so that it influences the overall design and structure; a privacy-by-policy approach fits more naturally with an existing IT architecture. The latter approach does of course give rise to more direct threats and vulnerabilities, justifying the need for risk assessment and policy enforcement.

Consider a simple example of privacy by policy. To preserve the privacy of its customers, an enterprise may have in place a policy disallowing employees to access the customer database outside office hours (during office hours, other controls are used to mitigate misbehaviour of employees). Such a policy may need to be enforced not only by technical means (in this access control mechanisms), but also through regular checks/audits on employees who stay in the office late. Thus, to implement this privacy requirement, there will be two levels of policy: managers or special support staff will enforce one level, while the other will be enforced by a technical solution.

While we are interested in scenarios involving privacy by policy, we believe that all the different policies implemented in an enterprise need to be informed by, and balanced against, the outcomes of risk assessment. As we have seen there is no uniform representation of privacy requirements, whether they originate in legislation, regulation, business practice or IT needs, and so it is not obvious how an enterprise can integrate privacy considerations into its risk assessments and, more generally, into all the different business processes involving the collection, storage and dissemination of personal data.

In this paper we identify how a holistic approach to privacy may be devised, which takes into account all the above issues. In particular, we discuss how to redress the balance between the various tensions and countervailing interests that influence an enterprise's privacy practices by:

- having a uniform representation of privacy requirements from different sources, and
- using this representation to inform risk assessments.

In EnCoRe we are exploring this approach while specifically focusing on an important aspect of privacy: the management of users' preferences with regard to the handling of their personal data (their expressions of *consent* and *revocation*). The specific area of management of privacy policies, security constraints, and consent and revocation [2] is of particular interest because it is at the intersection of legislation, user requirements and management of privacy and security technical policies within and across organisations.

This paper is organised as follows: Section 2 describes our holistic approach to privacy manage-

ment. Section 3 discusses the issues involved in creating an uniform representation for privacy requirements. Section 4 gives examples of diverse privacy and privacy-related requirements in order to show that a uniform representation is possible. Section 5 proposes a privacy aware risk assessment process and we present our conclusions in Section 6.

2 A HOLISTIC APPROACH TO PRIVACY MANAGEMENT

We explore here the variety of sources which can result in privacy requirements, and observe the dependencies which exist between the privacy rules dictated by each. First, there is a set of international legal requirements, which are set out by international agreements and directives, such as the European Data Protection Directive or the EU Safe Harbour agreement. These requirements tend to be the most abstract, with the intention that they are open to interpretation and refinement in the law of individual countries. Not all countries have data protection legislation on a national level; examples of national privacy laws include the Data Protection Act (1998) in the UK, the HIPAA, GLBA, SB 1386, COPPA and various State Breach laws in the USA.

Regulation gives rise to another set of privacy requirements, particularly relevant to enterprises operating on an international level: this includes export and transborder flow restrictions on personal data that need to be enforced. Security requirements which are typically enforced at management levels, include, for instance, adherence to the Sarbanes-Oxley Act (SOX) for financial reporting and the PCI Data Security Standard (DSS). Similarly, business requirements include contractual obligations, information lifecycle policies and the enterprise's own internal guidelines.

There are various operational and technical policies that are machine readable and enforceable by policy management frameworks. These may be expressed using technical policy languages and policy frameworks such as XACML [6], EPAL, P3P [4], Privacy RBAC [7], the latter two being targeted specifically at privacy-related IT policies.

Hence there are many levels of policies an enterprise has to cope with. Ideally all these kinds of policies should be managed and enforced successfully, in such a way that their requirements and stipulations are unambiguous and mutually consistent. In practice this can be difficult. However we believe that by finding a uniform representation, we can bridge some of the disconnection between higher and lower levels of policies.

The privacy and privacy-related policies, which apply in any particular enterprise, cannot be divorced from the security considerations, which arise as a result of risk assessment. This thesis is central to our approach; thus, any attempt at designing an IT system architecture for managing personal data requires an understanding of privacy requirements at all levels that is aligned with an analysis of risks and suitable protection mechanisms.

In Figure 1 we represent the different approaches to privacy enforcement. We note that:

- both approaches to privacy are ultimately implemented by means of a technical mechanism for enforcement,
- and that there is a range of hybrid approaches (represented by the dashed line between privacy by policy and privacy by architecture).

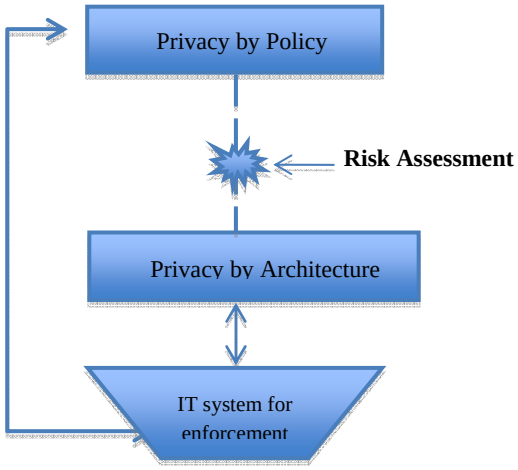

Figure 1: Different approaches to privacy enforcement.

In EnCoRe we aim to provide a solution that takes into account all the different perspectives; the approach will not aim to produce just a technical language for policies, divorced from the realistic needs of businesses and end-users. Rather, an assessment of risks will be made so that suitable privacy controls can be devised. Privacy enforcement will aim to be extensible and sufficiently general to handle a number of different enterprise scenarios.

A wide variety of knowledge and (technical, social and legal) expertise can be leveraged in EnCoRe to define approaches to privacy policy management at a legal level as well as at a technical implementation level. What is particularly desirable is to devise an intermediate representation of policies that is rich enough to embody all high-level requirements while being directly translatable to (potentially existing) low-level policies or access control languages such as XACML. Such a representation should not be tied

to a particular implementation language. This is one of our goals within EnCoRe.

3 TOWARDS A UNIFORM REPRESENTATION FOR PRIVACY REQUIREMENTS

Devising a uniform representation for privacy requirements involves investigating the tradeoffs between pragmatism and generality of policy representation approaches (so as to choose an approach that is neither overly pragmatic nor narrowly technical). It must also take into account all the levels of policy pertaining to personal data including legal, security and business angles.

Privacy policies typically contain stipulations about:

- For which purposes a data processor may collect personal data.
- Which types of personal data are considered sensitive, and hence are subject to additional restrictions.
- For how long collected personal data may be held.
- Whether and how personal data may be shared with third parties.
- Which actions a data processor must take in case of a privacy breach.

These reflect privacy principles that are common to the different levels of policies discussed in Section 2 above. A uniform representation should be able to express high level requirements such as those in national data protection legislation, as well as lower-level requirements which refer to how privacy is implemented in IT systems. Such a representation will consist of:

- a syntax for **conditions** that need to be checked,
- a syntax for **immediate actions** that should be performed if the conditions for a particular privacy rule are met,
- a syntax for **obligations** which the enterprise has if given conditions are met.

We believe that by defining these elements we can model most forms of privacy requirements, although we realise that these are context dependent, and will vary by type of policy: for instance transborder data flow policies will include conditions on the source and destination of data, while business policies will include conditions on the role of anyone accessing data, etc. In the next section we give examples of typical, but diverse, privacy requirements that an enterprise may be required to satisfy, and try to show that the elements of our representation as listed above can express them succinctly.

4 EXAMPLES OF DIFFERING PRIVACY REQUIREMENTS

The following examples show the value of being able to explicitly and uniformly represent the concepts and constraints involved in different types of policies as a way to reason about them. We believe that a conceptual model should provide a way to consistently represent all these concepts across different domains without the constraints induced by any specific "technical" language. To ensure continuity of the mapping between different layers, these requirements and policies need eventually to be mapped into enforceable technical policies, for example in languages such as XACML. This is where most of conceptual gaps can be identified as well as limitations of current technical approaches to policy languages. In the case of technical policies, we need to take into account a variety of details, for example where PII data and data subjects' preferences are stored, how to express constraints in a way that can be automatically enforced, how to deal with consent and revocation.

An example of a simple privacy-aware access control policy could be expressed as follows (note the use of conditions and actions):

```
Target: Personal Data D
if (Data Requestor wants to ac-
cess personal data D for Pur-
pose P)
and (data subject has given
consent for this data)
then Allow Access
else Deny Access
```

[Example 1]

For transborder data flow, rules may also be represented in the same form, for instance:

```
if (all source countries are
members of EEA and all target
countries are members of EEA)
then (no problems with trans-
border data flow)
```

[Example 2]

This type of rule is found neither in an access control policy nor in an obligation policy, but in what is known as a 'compliance policy'.

Notice and notifications require checking for "triggering" conditions and the context. Again, an **if...then** rule could be used to capture these concepts . For example:

```
if (<country legal entity re-
sides in> is member of [Bel-
```

```
gium, Portugal])
then (provide notification)
```
[Example 3]

This is more like an obligation policy, but note that it is not triggered by access control [2]. Another example would be that if there were a data breach then it would be necessary to notify the legal authorities and end users. This is an obligation policy, of a type that is triggered by an event. The key point here is that it is possible to identify some common patterns and concepts across these types of policies along with intermediate representations (e.g. rules) that are independent of underlying technical policies but which may nevertheless be fairly directly mapped onto these.

A similar analysis of policies can be made from a business and security perspective. Business policies, for example, relate to the treatment of information throughout its lifecycle and include: availability and recovery time policies, change control policies, binding contractual arrangements with third parties, service level agreements (SLAs) and IT governance policies. Also in this category are internal guidelines (that can map onto access control policies, obligation policies and/or compliance policies), and contractual obligations, which could relate to clauses included in contracts with clients, or to information contained within SLAs, etc.

Security requirements and related policies often originate in information security standards dictating methodologies and common security practices. These include: PCI DSS, Standard of Good Practice for Information Security, OCTAVE & CORAS (these are risk management methodologies), ISO 27001/2 (an international standard outlining best practices), BS 10012:2009 (British Standard outlining best practices); DoD MIL-STD-1629A (US Department of Defense risk management methodology). Usually these security requirements dictate constraints on who can do what on which protected resource, given a specific context. In the context of an employee data scenario, where the HR department of an enterprise has access to individual employee profiles, there might be access control policies with rules such as:

```
Target: EmployeeProfile P
if (Data Requestor is Role=HR)
then (Allow access to P)
```
[Example 4]

At a conceptual level we notice similarities about how to represent these constraints across different domains. In the specific case of management of personal data, privacy and security concepts can be conceptually bundled in a uniform representation. For example, both privacy and security constraints could be represented in the same **if...then...else** rule model:

```
Target: Personal Data X
If (Data Requestor is User
U/Role R in Context C)
and (Data Requestor wants to ac-
cess personal data D for Purpose
P)
and (data subject has given con-
sent for this data)
then (Allow access to X)
else (Deny access)
```
[Example 5]

5 PRIVACY-AWARE RISK ASSESSMENT

Once privacy requirements from all different sources and levels have been uniformly represented, they can be accounted for consistently during risk assessment. We propose here a form of risk assessment that takes into account privacy considerations, which are usually broader in scope than security considerations. Our approach differs from privacy impact assessments (PIAs) [12], in that often PIAs do not consider security requirements and are focused on particular components or system functionalities, whereas what we propose applies to an enterprise's privacy practices as a whole.

Risk assessment is used to quantify the potential negative impact resulting from security threats to and vulnerabilities within an enterprise information system. Such assessments determine which risks need to be protected against; the outcome of risk assessment influences the policies used in a system and the protection measures/controls that are implemented. Considering what could go wrong is important for understanding what needs to be done to effectively manage and protect personal information. However, the level of risk is determined in part by the perceived impact of a data loss, which is subjective and will be different depending upon whether it is the data-subjects perspective or the perspective of the enterprise handling the personal data. Our proposed risk assessment approach will require enterprises to consider the impact from the data-subjects' view. We are, in effect, taking the general ethos of PIAs but applying it within an integrated security and privacy risk assessment approach.

Some common risks that need to be addressed include lost or stolen media, over-sharing of personal information, good intentions but misused data, weaknesses of a third party with whom data has been shared, hackers, fraud, and social engineering. Failures may result in financial or legal penalties,

but ultimately may impact brand and reputation. Risk assessments usually take into account only security aspects. In the previous section (Example 4) an access control policy rule is shown; this may apply to a database of personal data. Taking into account privacy requirements would require adapting this policy rule significantly, into the form of Example 5. The difference is that the latter rule ensures privacy requirements (in this particular case, end-user consent) are satisfied.

6 CONCLUSION

We have discussed in this paper issues to do with the description, management and enforcement of policies in organisations. Specifically we highlighted the gap existing from a high-level approach to policies driven by risk and privacy impact assessment and low-level technical policies. We have explained and demonstrated the diversity of sources from which privacy requirements and constraints originate, and shown that they exhibit a relatively common structure, with a core set of conditions and actions. Our contribution has been to show that risk assessment and an understanding of privacy requirements go hand in hand, and that security and privacy risks need to be considered in tandem. Developing the integrated risk assessment methodology is a topic for future work, as is a formal conceptual model. As part of the EnCoRe project we are developing and refining this approach, and we expect that it will influence future developments at the intersection of security and privacy.

ACKNOWLEDGEMENTS

The position outlined in this paper is the result of many fruitful interactions within the EnCoRe project (see www.encore-project.info). We thank the project sponsors – TSB, EPSRC and ESRC.

REFERENCES

[1] EnCoRe. http://www.encore-project.info/

[2] Marco Casassa Mont (2006). *On the Need to Explicitly Manage Privacy Obligation Policies as Part of Good Data Handling Practices.* Proceedings ofW3C Workshop on Languages for Privacy Policy Negotiation and Semantics-Driven Enforcement, 17-18 October 2006, Ispra, Italy.

[3] Marco Casassa Mont, Siani Pearson, Gina Kounga, Yun Shen, and Pete Bramhall (2009). *On the Management of Consent and Revocation in Enterprises: Setting the Context.*Technical Report HPL-2009-49, HP Labs, Bristol.

[4] L. Cranor, B. Dobbs, S. Egelman, G. Hogben, J. Humphrey, M. Langheinrich, M. Marchiori, M. Presler-Marshall, J. M. Reagle, M. Schunter, D. A. Stampley, and R. Wenning (2006). *The Platform for Privacy Preferences 1.1 (P3P1.1) Specification.* World Wide Web Consortium Note NOTEP3P11-20061113.

[5] Marco Casassa Mont, Robert Thyne, Privacy Policy Enforcement in Enterprises with Identity Management Solutions (2006). Proceedings of PST 2006.

[6] OASIS eXtensible Access Control Markup Language (XACML). Standard available from http://www.oasisopen.org/committees/tc_home.php?wg_abbrev=xacml

[7] Qun Ni, Alberto Trombetta, Elisa Bertino, and Jorge Lobo (2007). Privacy-aware role based access control. In *Proceedings of the 12th ACM Symposium on Access Control Models and Technologies* (Sophia Antipolis, France, 20-22 June 2007). ACM, New York, pp. 41-50.

[8] Rodolfo Ferrini, Elisa Bertino (2009). A Comprehensive Approach for Solving Policy Heterogeneity. In *ICEIS 2009 -Proceedings of the 11th International Conference on Enterprise Information Systems* (Milan, Italy, 6-10 May 2009), pp. 63-68.

[9] Ioannis Agrafiotis, Sadie Creese, Michael Goldsmith, and Nikolaos Papanikolaou (2009). Reaching for Informed Revocation: Shutting Off the Tap on Personal Data. *Proceedings of Fifth International Summer School on Privacy and Identity Management for Life* (Nice, France, 7th – 11th September 2009).

[10] Sarah Spiekermann, Lorrie Faith Cranor (2009). Engineering Privacy. *IEEE Transactions on Software Engineering.* **35**(1), pp. 67-82.

[11] Peter P. Swire, Sol Bermann (2007). *Information Privacy: Official Reference for the Certified Privacy Professional (CIPP),* IAPP.

[12] A. Warren, R. Bayley, A. Charlesworth, C. Bennett, R. Clarke, C. Oppenheim (2008). Privacy Impact Assessments: international experience as a basis for UK guidance. Computer Law and Security Report, **24** (3), pp. 233-242.