Towards an Integrated Approach to the Management, Specification and Enforcement of Privacy Policies

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

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

ABSTRACT
We make the case for an integrated approach to privacy management within organisations. Current approaches to privacy management are either too high-level, enforcing privacy of personal data using legal compliance, risk and impact assessments, or too low-level, focusing only on the technical implementation of access controls to personal data held by an enterprise. High-level approaches tend to address privacy as an afterthought in ordinary business practice, and involve *ad hoc* enforcement practices; low-level approaches often leave out important legal and business considerations. As part of the EnCoRe project we are developing a methodology which tries to bridge the gap between privacy risk and impact assessment with the technical management of privacy policies.

We offer our thoughts on how a conceptual model might be devised as a means of expressing policy requirements as well as users' privacy preferences and offer a way to bridge the gap described above.

Keywords
Access control, policy, privacy, preferences, consent, revocation

1. POSITION STATEMENT
Enterprises manage and administer huge databases of personal data which are collected as part of normal business practice. This process is complex and involves meeting a wide range of requirements, including the need to satisfy data protection laws and privacy.

There is not yet a unified view of the different approaches to policies existing in an enterprise.. In general there are two extreme approaches to management and enforcement of privacy policies. There is firstly a *pragmatic approach*, mostly driven by risk assessment and risk management and tailored to current business practices. It involves identifying suitable high level policies and points to act on, but then typically requires the deployment of pragmatic control points, which are very dependent on the specific scenario/environment. In particular, the control points enforcing policies are often hardcoded within applications and services in an ad-hoc way, and so cannot be easily reused in different scenarios and organisational contexts. This makes it very expensive to reuse with other applications. However, this seems to be the norm in business practice today.

On the other hand, researchers often focus instead on a purely *technical approach* and narrowly propose yet another language or formal model for security, access control or obligations without taking into account legal, business and operational requirements. Hence, related policy languages might be too generic or detached from real requirements. What happens is that often these languages and models are of interest to the research community but seldom widely adopted in real environments.

We believe that there is a major gap between the two approaches and that there is a unique opportunity to combine aspects of each and provide mechanisms to bridge and provide continuity. Risk assessment and privacy impact assessment can help to identify threats and explore mitigations by means of suitable controls and related high-level policies. We believe that a conceptual model of these high-level policies enables reasoning about them and supports their refinment and mapping into low-level technical policies for practical enforcement in an IT information system. In the EnCoRe ("Ensuring Consent and Revocation") project[1], 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*).

Based on this position, our general approach is as follows:

1. Policies regarding the handling of personal data may be represented at different levels of abstractions within an enterprise, and so a unified, conceptual representation which allows us to compare and integrate them is desirable;
2. Current policy management approaches, tools and representations are suited only to particular classes of policies within this hierarchy, and so we aim to define approaches which bridge levels from legal requirements all the way down to technically implementable privacy and security policy;
3. When handling policies we want to take into account privacy preferences expressed by data subjects (end-users) and enforce them, hence enabling a user-centric perspective to privacy.
4. We envisage the need for a formal access control model embodying policy and preference concepts which enables reasoning from an abstract level (including legal, social and business aspects) to a technical, implementable level.

The different levels of privacy policies and their key characteristics are discussed in sections 2 and 3. Section 4 analyses the pros and cons of current approaches to privacy management and proposes a hybrid approach to the development and enforcement of privacy policies, which takes into account: legal, security and business requirements; the outcomes of risk and privacy impact assessment; what is feasible technologically.

[1] See www.encore-project.info and [2].

2. A HIERARCHY OF POLICY LEVELS

As discussed above, organisations need to cope with a variety of policies and constraints at different levels of abstraction, dictated by legal, social, business and individuals. This includes security and privacy requirements as well as data subjects' (end-users') preferences.

At the highest level of the hierarchy, there is a set of requirements which are set out by international agreements and directives, such as the European Data Protection Directive or the EU Safe Harbour agreement. Further, many countries have national data protection legislation, such as the Data Protection Act 1998 in the UK, or the HIPAA, GLBA, SB 1386, COPPA and various State Breach laws in US. With regards to regulation in particular, there are export and trans-border flow restrictions on personal data that need to be enforced. Privacy laws and regulations constitute the topmost layers of policy hierarchy regarding personal data with which an enterprise must comply. Such policies are often expressed in natural language as is typically the case with related data subjects' preferences.

At this high level of abstraction, security requirements may include adherence to the Sarbanes-Oxley Act (SOX) for financial reporting, or the PCI Data Security Standard (DSS). These may be refined to a set of policies at a lower level. Similarly, business requirements include contractual obligations, information lifecycle policies and the enterprise's own internal guidelines. All of the above influence how personal data is collected, stored and administered.

At the lowest level there are various operational, technical policies that are machine readable and enforceable by policy management frameworks. This includes XACML [5], EPAL, P3P [3], P-RBAC [6], and other technical policy languages.

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 introducing a conceptual model, we can bridge some of the disconnects between higher and lower levels of policies.

3. TOWARDS A CONCEPTUAL MODEL: CHARACTERISTIC ELEMENTS OF POLICIES IN THE HIERARCHY

Access control and privacy policies related to the protection of personal data 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 different levels of policies in the hierarchy presented in Section 2.

What is desirable is to have a uniform conceptual representation of the policies defined in the different layers. We consider here some of the distinctive features of the different types of policies, and for some we identify a general format; this is precisely what is needed for a conceptual representation. In future work we hope to find the structures that are common to all the different types of policies and to characterise them in a formal manner.

An example of a high-level policy is the set of data protection principles set out in the Data Protection Act. These principles (paraphrased versions of which follow) require that personal data shall:

- Be processed fairly and lawfully and shall not be processed unless certain conditions are met;
- Be obtained for a specified and lawful purpose and shall not be processed in any manner incompatible with that purpose;
- Be adequate, relevant and not excessive for those purposes;
- Be accurate and, where necessary, kept up to date;
- Not be kept for longer than is necessary for that purpose;
- Be processed in accordance with the data subject's rights;
- Be kept secure from unauthorised or unlawful processing and protected against accidental loss, destruction or damage by using the appropriate technical and organisational measures;
- Not be transferred to a country or territory outside the European Economic Area, unless that country or territory ensures an adequate level of protection for the rights and freedoms of data subjects in relation to the processing of personal data.

Such legislation may be translated into policies, but for some of these, refinement and/or interpretation will be necessary in order to translate these into operational (technical) policies, and this is easier to do with some policies than others. For example, it is straightforward to do this with notification requirements and some forms of transborder data flow, but not for transparency and adequacy requirements.

An example of a simple related privacy-aware access control policies could be expressed as an IF..THEN rule:

Target: Personal Data D

IF (Data Requestor wants to access personal data D for Purpose P) AND (data subject has given consent for this data) THEN Allow Access ELSE Deny Access

Similarly, for transborder data flow, rules may also be represented in the same form, such as:

IF (all source countries are members of EEA and all target countries are members of EEA) THEN: (no problems with transborder data flow)

This type of rule is not an access control policy or an obligation policy, but is a different type of policy – a 'compliance policy'.

For notice and notification, what would be needed is to provide the notification (or in some circumstances, check if it has already been provided), if the trigger conditions of an IF.. THEN rule are triggered. For example:

IF (<country legal entity resides in> ismemberof [Belgium, Portugal])THEN (provide notification)

This is more like an obligation policy, but note that it is not triggered by access control [1]. 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 business policies can be made. These relate to the treatment of information throughout its lifecycle, and that are also relevant to consider as background. These include: availability and recovery time policies, change control policies, binding contractual arrangements with third parties (e.g. standard 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 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)

Examples of requirements from PCI-DSS are:

- *Restrict access to cardholder data by business need to know.*
- *Track and monitor all access to network resources and cardholder data*

Usually these security requirements dictate constraints on who can do what on which protected resource, given a specific context. Conceptually this can be expressed in terms of access control policies

Target: Resource X

IF (Data Requestor is User U/Role R in Context C) THEN: Allow Access to X ELSE Deny Access

At a conceptual level we notice similarities about how to represent these constraints across different domains.

In case of personal data, conceptually privacy and security concepts can be easily bundled in a uniform representation.

For example, both privacy and security constraints could be represented in the same IF..THEN rule model:

Target: Personal Data X

IF (Data Requestor is User U/Role R in Context C) AND

(Data Requestor wants to access personal data D for Purpose P) AND (data subject has given consent for this data) THEN Allow Access to X ELSE Deny Access

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 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, etc.

An example of detail to be taken into account and related technical policies is given in Figure 1.

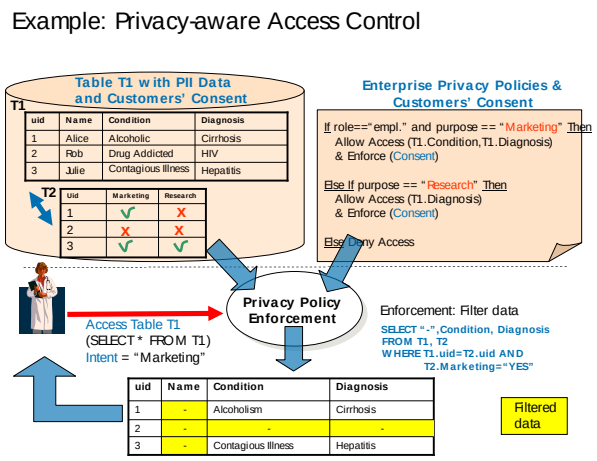

Figure 1. Technical Policies and implementation details

In this example, a basic privacy-aware access control policy (in a pseudo-technical language) could look like the following:

Target: <Database:DB1, Table:T1>

If DataRequestor.role=="employee" and DataRequestor.intent == "Marketing" **Then**

Allow Access (T1.Condition,T1.Diagnosis)

& Enforce (Consent)

Else If DataRequestor.intent == "Research" **Then**

Allow Access (T1.Diagnosis)

& Enforce (Consent)

Else Deny Access

This policy could be mapped in technical policies such as XACML. However, an accurate analysis of the example policy above [4] highlights that "conditional YES" might be required i.e. postponing the check of consent at the policy enforcement point.

This cannot be easily achieved with the current XACML representation. This implied that in the EnCoRe project we had to "twist" the language and framework to achieve the desired outcomes.

All this illustrates the power of a conceptual representation of policies to enable reasoning about and identification of constraints for the underlying levels.

What is desirable is to have a uniform conceptual representation of the policies that arise in the different layers.We have seen some of the distinctive features of the different types of policies, and for some we have identified the general format; this is precisely what is needed for a conceptual representation. In future work we hope to find the structures that are common to all the different types of policies and to characterise them in a formal manner.

4. POLICY MANAGEMENT APPROACHES

One might classify current research in privacy policy description, management and enforcement by using the diagram in Figure 2. There are two axes in this figure, corresponding to the two aspects we have discussed in this section. We have a vertical axis along which policies range from high-level (legal, regulatory) to low-level (security/access control policies and user preferences), and a horizontal axis which characterises the approaches to policy management as described previously (ranging from pragmatic to technical).

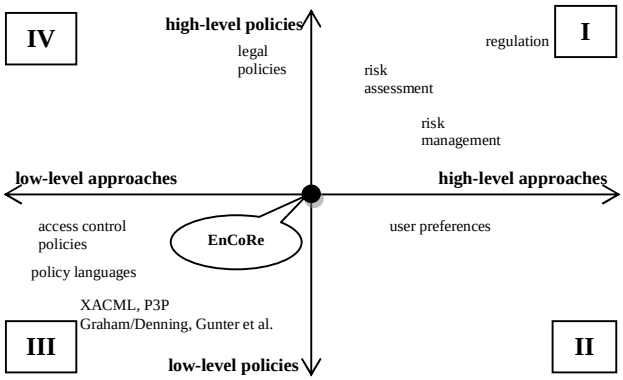

Figure 2. Classification of Privacy Policy Management Approaches versus Policy Levels.

It can be argued that a significant amount of research in this space falls in quadrant II of Figure I; this is no surprise, as the development of policy languages goes hand in hand with the need for machine readable descriptions of low-level technical policies.

It is evident from the figure that there are many other viewpoints and levels of abstraction that are of concern in policy management, and that there is scope for much work in the areas labelled as quadrants I, III and IV.

EnCoRe project is, among other things, aiming to explore how to build a conceptual model of policies bridging the existing gaps, in the specific area of management of consent and revocation [2].

This 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) and taking into account all the levels of policy pertaining to personal data including legal, security and business angles.

As indicated in Figure 2, EnCoRe aims to provide a solution that takes into account all the different angles; 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 and threats 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.

Expertise within EnCoRe varies widely, so that it is uniquely placed to advance aspects of privacy policies at a legal and regulatory level as well as at a technical implementation level.

What is particularly desirable is to devise an intermediate representation of policies that embodies high-level requirements while being directly translatable to a low-level policy or access control language such as XACML. Such a representation should not be tied to a particular implementation language. We will use the expertise available in EnCoRe, along with the approach outlined in Section 3, to achieve this goal.

5. CONCLUSIONS AND FUTURE WORK

We have discussed in this position paper several 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 strongly believe this gap needs to be filled to enable continuity of requirements and constraints across all these levels and enable proper enforcement of policies. To achieve this we proposed the adoption of a conceptual policy model, to enable reasoning and mapping of concepts at lower levels of abstraction.

Our analysis has given rise to: a model showing the range of privacy policy levels that exist (each level corresponds to a different layer of abstraction) versus the approaches used to describe and administer policies; elements of a conceptual model which characterises the format of the rules arising in different types of policy; first thoughts on an access control model which might be used to describe different types of policy rules in a uniform way, with an emphasis on user preferences.

This paper also proposes many avenues of investigation for further work.

6. ACKNOWLEDGMENTS

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.

7. REFERENCES

[1] Marco Casassa Mont (2006). *On the Need to Explicitly Manage Privacy Obligation Policies as Part of Good Data Handling Practices.* Proceedings of W3C Workshop on Languages for Privacy Policy Negotiation and Semantics-Driven Enforcement, 17-18 October 2006, Ispra, Italy.

[2] 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.

[3] 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 NOTE-P3P11-20061113.

[4] Marco Casassa Mont, Robert Thyne, Privacy Policy Enforcement in Enterprises with Identity Management Solutions, PST 2006, 2006.

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

[6] 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, June 20 - 22, 2007). ACM, New York, pp. 41-50.

[7] 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, May 6-10, 2009), pp. 63-68.

[8] 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).