FROM CREATIVE COMMONS TO SMART NOTICES
Designing User Centric Consent Management Systems for the Cloud

Siani Pearson[1] and Prodromos Tsiavos[2]
[1]*Cloud and Security Research Lab, HP Labs, Long Down Avenue, Bristol, U.K.*
[2]*Management Department, Information Systems and Innovation Group, LSE, London, U.K.*

Keywords: Cloud, Consent, Commons, Creative, Governance, Notice, Policy, Privacy, Notice, Smart.

Abstract: As cloud computing is evolving towards an ecosystem of service provision, in order for end users and customers to retain choice and control, they need to be able to select services, specify their preferences and have these reflected within the contractual framework, ideally enforced via a combination of legal and technical means. This paper presents an approach that builds upon successful methods from initiatives such as Creative Commons in order to improve the process of providing consent for usage of a data subject's personal data, and for achieving an appropriate balance between complexity and simplicity. This approach enhances the notices provided by service providers to advocate *Smart Notices* that provide a simple and transparent way of expressing the terms of service and the options available to the data subject before they share personal information with cloud service providers.

1 INTRODUCTION

Privacy issues are particularly pressing for cloud environments and are difficult to address (Gellman, 2009). There is a growing concern from data subjects, consumer advocates and regulators about the potentially significant impact on personal data protection and the required compliance to local regulations. Cloud can exacerbate the strain on traditional frameworks for privacy that globalization has already started. For example, location matters from a legal point of view, but in the cloud, information might be in multiple places, might be managed by different entities and it may be difficult to know the geographic location and which specific servers or storage devices will be used. It is currently difficult to ascertain and meet compliance requirements, as existing global legislation is complex and includes export restrictions, data retention restrictions, sector-specific restrictions and legislation at state and/or national levels. Legal advice is needed, transborder data flow restrictions need to be taken into account, and care must be taken to delete data and virtual storage devices when appropriate. Moreover, the Patriot Act in particular causes fears about transferring information to US. Most privacy issues are shared with other paradigms, such as service-oriented architectures (SOA), grid, web-based services or outsourcing, but often they are exacerbated by cloud, as traditional solutions like model contracts (to allow certain transborder data flows) take too long to set up and are not suited to these types of dynamic environment.

Context is important, in the sense that different information can have different privacy, security and confidentiality requirements. Privacy need be taken into account only if the cloud service handles personal information. There is a low privacy threat if the cloud services is to process information that is (or is very shortly to be) public. That is why the New York Times conversion of scanned images to pdf from a few years ago, that was at the time often highlighted as a classic demonstration of the benefits of a cloud approach, was a good scenario for cloud computing. However, there is a high privacy threat for cloud services that are dynamically personalized, based on people's location, preferences, calendar and social networks, etc. The same information collected in different contexts by different entities might have completely different data protection requirements. Nevertheless it should be borne in mind that there may be confidentiality issues in the cloud even if there is no "personal data" involved. In this sense, the practices and technologies described in this paper for ensuring appropriate

Pearson S. and Tsiavos P..
FROM CREATIVE COMMONS TO SMART NOTICES - Designing User Centric Consent Management Systems for the Cloud.
DOI: 10.5220/0003979206470660
In *Proceedings of the 2nd International Conference on Cloud Computing and Services Science* (CloudSecGov-2012), pages 647-660
ISBN: 978-989-8565-05-1

personal data handling among the cloud ecosystem would also be beneficial for use in protecting intellectual property and trade secrets.

The central motivation for our approach is to aid user control, choice and transparency within cloud service provision – and more widely, for other scenarios where there is a complex service provision infrastructure. The need for such an approach is underlined not only by core principles of privacy that are included in data protection legislation across the world, but indeed within the development of prospective regulation including the forthcoming EU General Data Protection Regulation, the current draft of which mandates that consent should be opt-in. However, current cloud terms of service and SLAs are not easy to understand for end users and business decision makers and offer little choice (Mowbray, 2009; Alhamad et al, 2011).

Our work has been carried out within the context of the EnCoRe project (EnCoRe, 2012), which is developing mechanisms for user-centric consent. From the point of view of a data subject, or end-user, if an enterprise provides consent and revocation controls, it increases that user's *choice* regarding how his or her personal data is handled. Correspondingly, an enterprise would define its own consent and revocation policy, as a way of informing end-users of the choices available to them. This can be encapsulated within a 'Smart Notice' that is provided to cloud service users.

2 BACKGROUND

In this section relevant background is considered that focuses on the scope, transparency and usability of privacy policy management.

2.1 Privacy by Policy

"Privacy by policy" is the standard current means of protecting privacy rights through laws and organizational privacy policies, which must be enforced. Privacy by policy mechanisms focus on provision of notice, choice, security safeguards, access and accountability (via audits and privacy policy management technology). Often, mechanisms are required to obtain and record consent. The 'privacy by policy' approach is central to the current legislative approach, although there is another approach to privacy protection, which is 'privacy by architecture' (Spiekermann & Cranor, 2009), which relies on technology to provide anonymity. The latter is often viewed as too expensive or restrictive. Although in privacy by policy the elements can more easily be broken down, it is possible to enhance that approach to cover a hybrid approach with privacy by architecture.

Notice and choice are the pillars of good privacy practice, according to the privacy by policy approach. A privacy statement (often called a 'privacy policy') communicates to the data subject which personal data an organization collects, how they are used, to whom they are disclosed, how long they are retained etc. It needs to be accurate, understandable and complete. Achieving all three is challenging in practice. Some attempts to improve this have been developed, such as the concept of layered privacy notices, as discussed in the following subsection. Notice can be prominent or discoverable. Software notice can be presented at different times according to the context, including at installation time or when information is about to be collected ('just in time'). Similarly, user choice can involve a number of different consent options, and these can be expressed in explicit or implicit ways, for example via opt-in or opt-out means.

2.1.1 Layered Notices

Privacy notices can often be difficult to understand, complex and long. Layered notices are designed to be more readable and understandable. The information is structured into multiple parts (typically, two). The first layer is a condensed privacy policy that provides the reader with a clear summary of the policy, and that provides links to more detailed information. The other layers provide the full privacy policy and can contain more detailed or specific information, possibly being broken down into separate web pages for readability or searchability. This helps readers locate the important points without having to struggle with all the detail unless they actually wish to do that.

There is some international support for layered notices, including endorsements from 25th International Data Protection Conference in Sydney, Australia (2003), Article 29 Working Party of the European Union (2004), Asia Pacific Economic Community (APEC) (2005) and OECD Working Party of Information Security and Privacy (2006).

2.2 Policy Management

Organisations need to cope with a variety of policies and constraints that emerge from many different sources, including legislation (national and international), societal expectations, business requirements and (where appropriate) individual preferences expressed by users and customers. In this paper we focus on those policies relating to the handling of personal data and privacy.

Whilst privacy requirements are in general context dependent, we believe that there are a core set of privacy concepts which are common and underpin the various controls designed to deliver privacy against this varying set of requirements.

We consider policies to fit within a layered model which in itself represents a hierarchy of policies. In this model, high-level policies express general requirements and rights, as embodied typically in law, business and regulatory requirements 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 HIPAA, GLBA, SB 1386, COPPA and various State Breach laws in US and there are export and transborder flow restrictions on personal data that need to be enforced. Privacy laws and regulations are often expressed in natural language as is typically the case with related data subjects' preferences. Security requirements may include adherence to the Sarbanes-Oxley Act (SOX) for financial reporting, or the PCI Data Security Standard (DSS).

The preferences of a data subject are high level policies that need to be taken into consideration, along with contractual obligations, internal organizational policies and legal constraints. Hence the origins of privacy requirements which an enterprise has to meet are very diverse, and they arise at many different levels of abstraction. In an ideal world, lower level policies should always be the result of refinements, or special cases, of the higher level ones. In the real world, high-level requirements change over time. Data subjects and data controllers exercise choices relating to their preferences and risk appetites. This makes it impossible for a system to always be a correct refinement of requirements, as it will take time for choices to be implemented. It will be for the data subjects to decide whether they are being offered appropriate service levels regarding the response to their choices, and for service providers to determine what level of guarantee is appropriate for their business model.

2.2.1 Policy Representation

Translation of legislation/regulation to machine readable policies has proven very difficult, although there are some examples of how translations of principles into machine readable policies can be done: in particular, the REALM project (IBM, 2006) has worked on translating high level policy and compliance constraints into machine readable formats, and research into how to extract privacy rules and regulations from natural language text (Breaux & Antón, 2008). For a summary of progress to date in this field, see (Papanikolaou et al, 2011). It is still an open problem how to interpret and model arbitrary laws. As an alternative, company policies can be mapped to lower level implementable policies, or human-readable output: HP Privacy Advisor represents HP privacy policies in a machine readable format and analyses these to provide human-readable customized output relating to specific circumstances (Pearson, 2010).

Besides the high level policies that describe regulatory and legal constraints, there can be a range of lower-level policies including descriptions of how privacy requirements are implemented in a particular piece of hardware, or in software that handles personal data. Some instantiations may be specific to a particular system. Such policies comprise detailed conditions on how particular data is to be handled within a system: often these are just statements prohibiting particular access to the data, in which case they are referred to as access control policies. These policies can be machine-readable and enforceable by policy management frameworks.

There are a number of existing options including EPAL (IBM, 2004), OASIS XACML (OASIS, 2012) and extensions (Ardagna et al, 2009; Bussard & Becker, 2009; Papanikolaou et al, 2010), W3C P3P (Cranor, 2002), Ponder (Damianou et al, 2001), PRIME (Ardagna et al, 2006), E-P3P (Schunter & Waidner, 2003) and SecPAL4P (Becker et al, 2009). Most of these focus on internal back-end policies, but nevertheless there can – and indeed should – be a mapping between user-defined policies across to these machine-readable policies enforced by service providers. However, the resultant low level privacy policy languages (such as those provided by EPAL and XACML) are not well suited for human user understanding.

2.2.2 Policy Matching

Various approaches have been taken whereby the user can define policies that govern handling of their data that are matched, and even negotiated, against service provider policies. Sometimes this is done prior to release of data, and sometimes the checking may be done by third parties after the data is released in an encrypted form, but before the decryption key is made available to the service provider. Examples of this approach include:

- **Privacy Incorporated Software Agent (PISA)** (Kenny & Borking, 2002): project in which privacy principles derived from (OECD, 1980) were modelled and used as a backbone in conversations between agents
- **P3P** (Cranor, 2002): user privacy preferences were matched against web site privacy statements)
- **Xpref** (Agrawal et al, 2005): a preference language that can be automatically matched against P3P policies)
- **PRIME** (Camenisch, Leenes & Sommer, 2011): project involving the definition and usage of various types of user and service side privacy policies – including specification of user requirements about service side policies and associated service provider assurance policies, with related real-time checking of usage of the backend security provisions specified in these policies (Pearson, 2011)
- **PrimeLife** (Camenisch, Fischer-Hübner & Rannenberg, 2011): project extending work from PRIME, including initial steps at structuring legal data protection policy representation in different contexts (Holtz & Schallaböck, 2011)
- **EnCore** (EnCoRe, 2012): project involving privacy-enhanced access control and obligation policies on the back end, together with 'sticky policies' that specify data usage requirements and that are stuck to data as it passes around the cloud service provision eco-system (Pearson et al, 2011)

One of the key issues is in getting infrastructure providers and service providers to take up such an approach; another is in making it easy for the users to define their policies. An analysis of why P3P has failed to achieve take-up in the marketplace is given in (Jaatun et al, 2010). There has been a range of different work to help with the usability issue, as discussed in the next section, but this is still an open issue that has not been adequately solved.

2.3 Usability

The P3P preference language APPEL is a standard for encoding users' privacy preferences in a machine-readable way, but the syntax of both this, XPref and P3P is difficult for users to deal with directly. Several tools have been developed to help facilitate this process, notably a policy editor to assist service providers to define P3P policies (Bergmann, Rost & Pettersson, 2006) (although this process is still somewhat cumbersome) and AT&Ts Privacy Bird (Cranor et al., 2006). The latter is a plug-in for Internet Explorer that monitors P3P policies for the user; it has an easy to use interface, but with very limited options.

An alternative approach is to ask a series of dynamic questions which the user can answer to inform agents about their privacy preferences and by these means to set user policies (Irwin and Yu, 2005).

It is also worth considering the balance between flexibility in policy definition and usability: for example, a pre-defined set of natural language clauses might be used as the policies and evidence could be provided by the system that these are satisfied on the back end (Elahi and Pearson, 2007). Patrick and Kenny (2003) described the HCI requirements of an effective privacy interface design. The PRIME project (Pettersson et al., 2005) used three UI paradigms – role-centred, relationship-centred and town map-based paradigms – for privacy-enhanced identity management in the PRIME project. Andersson and others (2005) discussed the socio-psychological factors and HCI aspects that influence end users' trust in privacy enhancing identity management. Hawkey and Inkpen (2006) examined the privacy comfort levels of participants if others can view traces of their web browsing activity. At the implementation level, Kobsa (2003) adopted a redundant component array architecture to personalised web systems so that they can dynamically adjust to the current prevailing privacy concerns and requirements without burdening the application with privacy management tasks. Iachello and Hong (2007) summarised previous research and proposed new research directions in privacy-aware HCI. Work is currently being carried out in a number of projects related to how to visualise privacy to the user, notably MobiLife (MobiLife, 2012) and VOME (VOME, 2012).

The Sparcle project (IBM, 2007) built an editor to support transforming natural language based policies into XML code that can be utilised by enforcement engines. This makes it easier for non-experts to input rules into the system, but the output format itself is not user friendly and is targeted towards machine execution. In HP's Privacy Advisor tool that essentially carries out internal privacy impact assessments, UIs were provided for potentially untrained employees to input contextual information and view system reports, privacy officers to offer advice within a defined workflow, administrators to set access rights and system settings and domain experts to amend company policies within a relatively user-friendly setting, although in some situations complexity could not be avoided (Pearson, 2010).

Mary Rundle has carried out some seminal work on the use of privacy icons (Rundle, 2006), which can be used to help facilitate end user policy definition and understanding. This work has been extended by others, including within the PrimeLife project (Holtz et al, 2011), where icon sets were developed and tested on end users for different use cases including e-commerce, social networks and handling of email. A related technique is to use privacy labelling to convey privacy policies and preferences to data subjects (Kelley et al., 2009). Note however that the use of icons cannot (and indeed is not intended to) replace full, written privacy policies as the basis for informed consent, according to European privacy regulations.

2.4 Privacy Commons

Privacy Commons (PC) draws from the success of Creative Commons (CC) (CC, 2012), but tries to implement some of its key features in the context of personal data rather than copyright. PC differs substantially from CC in the sense that its main objective is not to provide licences but a policy framework which may subsequently be converted into a contract provided the parties agree to it. Another important difference relates to the structure of the licences from a vertical and horizontal perspective. CC comprises three layers: (a) the legal expression of the terms under which the transaction takes place that are described as the “legal code” (b) meta-data that describe the main licence features in Rights Expression Language (REL) that makes the licences findable by search engines such as Google or Yahoo! or tagged in platforms like Flickr and (c) the “human readable” code that is the licence expressed in simple language and a set of standardised icons that reflect their basic features. In addition, the structure of the CC licences is modular, i.e. they comprise of three variable and one fixed element that may be freely combined in order to produce six licences. Finally, the CC licences have local implementations in over 60 jurisdictions around the world.

PC, on the other hand, currently has only a legal and a human readable layer and is still very US-centric. The technical layer could be supplemented by other technology providers and solutions like the ones presented above, but for the time being there is only some very basic discussion about what such tools should be. An important link is between the PC contracts and the ToSBack service of the Electronic Frontier Foundation (EFF) that tracks changes in the ToS of large user-base services such as Facebook or Google (EFF, 2012). While discussions within the PC group acknowledge the limitations of ToSBack in its ability to educate users regarding their privacy rights, it is still considered a good instrument to build on a service that is closer to the PC objectives (PC, 2012). Finally, PC has also produced a series of modules with which it is experimenting such as Complete, Opt-in, No Rights, Certified, No Breaches and Auditable.

Because of its US-centric approach, if a PC policy is converted into a contract, then this is not based on Copyright law but rather on a combination of tort, contract and Intellectual Property (IP) law. While IP is not recognised upon specific data-points it is debatable whether it is possible to assert any Intellectual Property Rights (IPRs) on a data set referring to a specific individual and the meta-data associated with his/her personally identifiable information (PII). Again this may be a rather problematic approach where the data are collected and managed by an entity other than the data subject and in that case it is most likely that the PC agreement will have to be construed as a contract and not as an IP licence.

PC identifies a series of problems with respect to privacy policies that it aims at addressing. First, that the quality of most existing policies is of a low level. Second, that the policies tend to waive rather than to assert privacy rights for the end user. Third, that they are not easily understood by the end user. Finally, US courts have not attached any legal consequences from the violation of privacy policies which are not normally deemed as a contract. Hence further research is required to satisfy these issues.

In a PC scenario, there are two parties, the Data Steward (Steward) and the Data Subject (Subject). The objective of a PC framework is to convert privacy policies into proper contracts including offer, consideration and acceptance. The transactional side of the PC contract is yet another difference between PC and CC, since the latter does not describe, in the jurisdictions where this is possible, its licences as contracts but as “bare” licences that do not require consideration and acceptance.

While PC aims to transform into a non-profit organisation, for the time being it remains a loose network of “grass-root efforts” with no specific organisational affiliation.

The overall objective of PC is to provide awareness to the general public regarding rights over their PII and to make available tools that may empower them in their transaction with service providers. Regarding the latter, PC is considering being active in the areas of machine-readable versions of the PC contracts, icons representing the contracts in simple terms, developing a common vocabulary for privacy contracts and allowing users to express their privacy preferences through software agents. Consistent with the CC vision, PC is geared towards adoption, not enforcement.

PC operates on the basis of a set of core data disclosure requirements that are common for different realms of activity, each one of which has its own policy framework. Each industry may have disclosure requirements additional to the common set of core disclosure requirements. PC has identified a broad range of areas of activity that include: Goods and Services, Healthcare, Financial, Education, Network Provider and Government.

All disclosure requirements are split up into Required, Optional and Prohibited Representations. Each policy suggested by a service provider is to be assessed on the basis of a model of Privacy Policy Requirements set by PC per area of activity.

At the current development stage, PC proceeds by producing Use Cases where a set of goals is to be implemented through the deployment of different technical components. The PC privacy policy is assumed to have been marked up so that it may be “read” by the different components. The scenario which is currently available in the PC site refers to generic Internet browsing. The technical components present in a PC scenario are three: First, a web privacy layer that assesses the privacy level of the site and alerts the user on the basis of an alert threshold set by the user. The privacy level of the site is assessed on the basis of compliance with the PC policy framework suggested for that particular context. Second, there is a set of Privacy Commons Registrars. This involves a “marketplace” for generating and registering PC-compliant privacy policies. This model would operate in a way similar to an SSL certificate marketplace. A certain number of companies could be recognized as trusted authorities with respect to PC-compliant privacy policies (Parsons 2010). Finally, there is the moderation/ reporting layer that involves the provision of web services that report violations of the PC policy, report noncompliant policies, and compile user-requested reports of compliance activity for a specific site.

The role of rights expression languages in this context is particularly important, although, as the experience from the PRIME project indicates, there are still substantial problems to be tackled. CC is based on REL whereas in the privacy context other tools need to be used such as privacy policies (as considered in subsection 2.2.1). The extent to which the PC policy framework will make use of such languages still remains unclear. The development of a common privacy ontology also seems to be a requirement for moving into a PC-like solution, though PC’s exact contribution toward that direction remains unclear.

Another interesting feature suggested by PC is the employment of user ratings in order to assess the quality of a service provider that acts as data controller. WhatApp (Stanford, 2012) is an open rating system that could act as a blueprint regarding how such services could operate in the future.

3 SMART NOTICES

In this section we propose the new notion of ‘smart notice’ in order to aid end user control, choice and transparency in cloud computing scenarios. A ‘smart notice’ is a customisable and searchable set of related policies that would be shown to end users by service providers, in place of the current standard fixed ‘notice’ approach. It would be generated via the end user being offered a number of different consent options, from which he or she may select. From the user options chosen, machine-readable, human-readable and legal policies would be generated corresponding to these choices and these would comprise the ‘smart notice’. These policies may all be viewed by the end user, although it is most likely that the human-readable policy (that shows a simple version of what the policy does – rather like the ‘top’ level of the layered notice) would be the only one the user actually wished to see. The machine-readable policy may be used within automated policy checking and enforcement mechanisms such as the ones developed within the EnCoRe project (EnCoRe, 2012). The legal policy could be used in case redress were needed, and to help enforce obligations set by the user via legal means. Note that there are only a limited number of possible options that the smart notice can take, and so the different policy layers can either be fixed in advance or generated at the time according to the choices made.

The context of usage is a user-centric approach where the end user controls their personal data stores, even if those are stored remotely in the cloud, and sets preferences that relate to the usage (including sharing) of their data. Rather than the user defining a complicated policy, the policy creation process needs to be easy. So, one instantiation of this approach would be as a wizard that shows a questionnaire to the user, and from the answers that the user selects, the smart notice can be created. An example of this approach is considered in the following section. Another approach would be for a combination of slider bars and drop-down boxes to be offered for user selection, which again then automatically creates the associated policies.

Within this approach we may incorporate the usage of privacy icons (although too much usage and reliance on these can become confusing for the user). The privacy icons can be used to help achieve transparency within the human readable policy – and indeed that policy might even be broken down into two parts comprising such icons as well as a textual description. This is rather similar to the suggestion of the Article 29 Working Party (AWP, 2004) that icons might be used within one layer of layered privacy notices. In addition, some other aspects of layered notices are similar to this concept, in the sense that there are different versions of the notice that are intended for being read or used in different circumstances. Unlike layered notices though, these different versions are targetted for readability to different audiences and there is not such a risk of organisations deliberately hiding 'bad' privacy policies in the lower levels.

A promising option for expressing the policies themselves is in the approach developed by Creative Commons (CC, 2012), although this is not a necessary part of the approach. Even though the main goal and focus of that work relates to copyright and is to create a mechanism for easy specification of data rights by data owners, there are strong parallels with the current problem and potential solution under discussion.

Automation of contractual terms and conditions and simple expression has been very successfully used in the case of the CC licences and the "copymarks" idea (Bing, 2004). The CC licences differ from other End User Licence Agreements (EULAs) in the sense that (a) the end user is the licensor and not necessarily the licensee (b) they are accompanied by what is called the "Commons Deed" or an abstraction of their main terms and conditions in very simple language accompanied by diagrams schematically explaining such features and (c) they are supported by an extensive and growing open source community of developers building software tools allowing the tracing and management of content licensed under such licensing schemes. Such a licensing scheme allows the end-user to reduce the costs of participating in the construction of micro-regulatory regimes controlling his/her own content by providing a range of ready-made legal instruments with technological implementation support that are easy to understand and use.

A similar solution could be used in the case of personal data management though some important differences between the two domains (i.e. Copyright versus Personal Data regulation) entail a slightly different approach. For instance, although the individual has rights as a result of the Data Protection act, there is no standardised licence for him/her to allow the use of such data. The way the whole system works right now is that the end user rather than the data controller agrees to an EULA. It would be interesting to see how a reverse model would work. There are a number of variants on the basic CC approach, including PC and Consent Commons. Building upon PC within the implementation of Smart Notices allows use of existing policy templates, as well as standard terms shared by different organisations. This has a big advantage in that the ontology and usage base is already at least partially established. In particular, there is the option to:

- specify '*PrivacyAlike*', i.e. to share information only with organisations following the same principles
- use the existing representations for specifying that usage should not be for commercial purposes: *NonCommercial*
- request notification in case of the information being shared: *ShareNotification*
- be involved in the benefits accruing from the usage of that information: *BenefitShare*.

Consent Commons is set up to obtain consent at the point of rights collection, and this approach may be exploited as well.

CC already uses the concept of a wizard to set up policies (CC, 2012b), and this concept may be extended directly. For the machine-readable layer, the formatting of REL CC (CC, 2012c) may be used. There is then already defined a mapping across to the human readable form of the policy (CC, 2012d) and to the corresponding lawyer readable form (CC, 2012e), and these could be used to generate the different types of policy within the Smart Notice.

3.1 Implementing Smart Notices

As noted above, CC is designed to define property rights. If we transition that approach to a privacy context, this has implications for the requirements and design of the system.

In particular, our focus is on consent management and not the definition of ownership of data. With regard to his or her personal data, there are three principal things for which an enterprise or other data collector may require the consent of an end-user: collection, processing and sharing of data.

Collection of data refers to the initial process by which data is acquired and stored on the enterprise's information system. Processing includes any access of the data that has been collected and is characterised by a *stated purpose* (e.g. research, marketing, aggregation to derive average customer habits). Data may be shared – internally and externally (e.g. to third parties) so that it can be processed, often elsewhere than the site of data collection. The definition of consent as a wish for data to be collected, processed or shared is too coarse, for it does not account for subtleties such as desires to:

- restrict data collection so that it occurs only in selected jurisdictions.
- expire consent after a fixed period of time.
- restrict processing of data so that it is used for only certain stated purposes.
- share the data only with particular parties.

Thus we claim that consent is parameterised by certain quantities referred to as *consent variables* (Encore, 2012). Examples would be the time for which consent is granted, the data for which consent is granted, the set of stated purposes for which consent is granted and the set of parties who may access the data. Our approach is that consent is fully determined when the following are specified:

- the task for which consent is given (collection, processing, sharing, or any combination thereof)
- for this task, the values of the consent variables of interest.

Revocation corresponds to the withholding or withdrawal of consent (manifested in its simplest form as deletion of data).

Any convenient representation can be used for the policies, including the user policy representations. However, as discussed above, we advocate using PC as the basis for the representation of the policies. In order to do this, we need to create abstractions of policy terms, and produce sharing options and link these to consent. We also need to define the type of access that the sharing would entail. The approach used needs to be different from the standard CC licences where there is no filling in, as we need to allow customisation and selection of options, and hence a 'filling in' approach.

There are a number of possible mechanisms for checking and enforcement of the policies created, ranging from technical enforcement to social means of enforcement. For example, the human readable form of the policies can be enforced via peer-to-peer pressure and reputation systems, the legal form of the policies can be enforced via legal and regulatory mechanisms, and the machine readable version of the policies can be enforced via technical enforcement mechanisms such as obligation management and access control.

We now consider in detail an example illustrating this approach.

4 USE CASE

In this section we consider a use case that we have considered within the EnCoRe project (EnCoRe, 2012). This project is a collaborative research project into informational privacy by UK industry and academia. EnCoRe's approach is an interdisciplinary one based on the notion of trust that the limits of an individual's consent will be respected by all those that process his/her personal data, and that tools are needed to manage the consent lifecycle effectively. Although EnCoRe is not primarily a cloud-based approach, similar mechanisms could be applied in cloud scenarios. We have considered situations which could be common to a number of situations, such as health service provision, access to applications and services in the cloud (storage, computing, etc.), and so on: in all these situations, a customer needs to reveal personal and even sensitive information in order to receive a service, but wishes to control the way in which that information is used. In the following sections we concentrate upon on particular example, for reasons of space.

4.1 Biobank Scenario

We shall consider the EnCoRe policies defined for our second case study, which focuses on obtaining user consent for research related to biobanks. In this case, both human tissue and the associated data are shared, for specific research studies and control experiments, in relation to a specific real-life biobank (ORB). Within the policies, the commercial use and users need to be specified, and the research use and users defined. Part of the workflow is that constraints associated with research studies are defined upfront within a 'REC' document approved by an ethics advisory board before the research can start. These obligations need to be reflected in the sharing conditions, along with specific requirements from the UK Data Protection Act (DPA) (1998) and Human Tissue Act (HTA) (2004).

The sharing conditions should ideally take into account access to tissue, personal data, meta-data and research results, what is considered to be private, whether information may be provided to a specific person or people satisfying certain roles, what the primary usage would be and whether the information may be re-used.

4.2 Policy Definition

Policies need to be defined for patients to express their preferences about the handling of their data and samples, and also for biobanks to clarify what options are available and will be respected within the network within which information is shared.

4.2.1 Patient Policies

Patient choices are provided within the Smart Notice and the patient makes a selection (via answering questions generated by a wizard) to customise their policy, per record.

The term policy here describes the high level policy modules that are addressed to the end – user, and potentially also to other data controllers, and these need to be translated into internal organisational policies.

The basic elements of the EnCoRe policy are as follows:

1. *Regulatory Obligations*. This incorporates all elements found in policies that refer to obligations of the sample/ data manager that are the result of legislation/ regulations, including in particular UK DPA and HTA, as well as Research Ethics Comittees.
2. *Management*. These are the actions that the biobank has to perform. Relevant consent variables within the policies are time, aggregation, destruction and retention time. The policies specify the core actions that relate to management in the strict sense of the material or the data and include all Not Access Related Management (NASAM) actions such as:
    - storage
    - documentation
    - destruction of the sample/data
    - retention time of data/sample
    - possible linking of data/samples of the same data-subject/donor (this is a form of aggregation)
3. *Access*. This is a generic term to describe both simple access and re-distribution of the sample/ data. Relevant consent variables within the policies are Commercial, Research, and Specific type of research; ShareBack (monetary or samples); Re-deposit (additional samples, data, research); PrivacyAlike: all the original conditions are to be enforced to anyone using the data/ sample further – with options for additional or less restrictions as the sample is passed to third parties. The basic elements include:
    - Entity type accessing the data (with definition of groups or circles of access)
    - Type of processing/ use (or Purpose)
    - Time of access
    - Redistribution
    - Aggregation/ Linking: conditions as to how much data may be collected about a single individual
4. *Notification/ Contact.* Key element are:
    - Re-consent in the case of re-distribution when the purposes of the third party are different from the original accessor
    - Frequency
    - Mechanism
5. *Revocation*
6. *Anonymisation*
7. *Delegation*: in the event of termination of the donor/ data subject.

The generic form of a minimal list of user consent and revocation options offered to the user for the generic case could be as follows:

I {consent/do not consent/revoke consent} *for this EnCoRe compliant system* to {collect/store/use} my {data/sample/data and sample} for specified purpose (subject to time constraints/notification constraints/usage count restraints)

[*contact*] I {consent/do not consent/revoke consent} *for this EnCoRe compliant system* to contact me via {email, phone, post, and/or GP} about my data {sample} (for Specified purpose)

[*sharing]* I {consent/do not consent/revoke consent} *for this EnCoRe compliant system* to share with/copy my data {sample} for Specified purpose with Specified Data Controller who is an {EnCoRe Compliant Data Controller/non-EnCoRe Compliant Data Controller}

For the use case in question, this list becomes somewhat more complex, as follows:

I {consent /revoke consent} for ORB to {collect/store/use} my personal data for { any research (provided it has been approved by ORB and met all ethical standards of research); DNA specific research; selected clinical trials [list]; not at all} with access by {the research team that contacts me; pharmaceutical companies; others} (subject to time constraints/notification constraints)

I {consent/do not consent/revoke consent} for ORB to {collect/store/use} my {sample and associated digital representations} for {Specified purpose} (subject to time constraints/notification constraints)

[contact] I {consent/do not consent/revoke consent} for ORB to contact me about my data or sample via {e-mail, phone, post, GP} when {my sample is shared, results of the research have gone public}

[sharing] I {consent/do not consent/revoke consent} for ORB to share my sample (or its digital representations) for {Specified purpose} to {direct contacts of the researcher, anyone}

[sharing] I {consent/do not consent/revoke consent} for ORB to share data for {any research (provided it has been approved by ORB and met all ethical standards of research); selected clinical trials [list]; only the research team that contacts me}

This approach allows finer-grained user control and flexibility than existing consent model forms for consent to research, such as at http:// www.p3gobservatory.org/repository/ethics.htm. However, this information needs to be in a more user-friendly form, and ideally to utilise the existing icons provided by CC. Accordingly, we have investigated how EnCoRe might offer a wizard as part of the user inerface shown to end users via the Consent and Revocation Assistant; an IT admin with the relevant permissions in an EnCoRe-compliant organization can reduce the policy options offered by wizard if necessary; the wizard shows a smart consent and revocation form to end users in order to automatically generate from the user options chosen machine-readable, human-readable and legal policies corresponding to these. The human-readable form is a 'smart notice' that is a searchable policy that is customized, but in the sense of there only being a limited number of possible options it could take.

The wizard will show a questionnaire to the user and the user will answer the questions given. An example questionnaire is:

Which types of entity can use your data?

- commercial
- non-commercial

For which purposes can it be used?

- commercial
- research
- non-commercial

How long can it be used for?

- for ever
- duration of study
- reconsent after 2 years

Do you require shareback?

- yes
- no

If 'yes', what type of shareback do you require?

1. sample
2. data
3. research
4. money

Will you allow sharing?

- not at all
- to organizations or individuals with the same policies
- to organizations or individuals with stronger policies
- to organizations or individuals with weaker policies

Would you allow the biobank or the person accessing the information to link/aggregate data/build a profile about you?

- yes
- no

I would like to be contacted:

- never
- when other people access my data
- when other people ask to access my data
- when new consent would need to be given for other people to use my data

I prefer to be contacted:

- by email
- by post
- by telephone

Ancillary help information is integrated: for example, there is an explanation of what shareback is when the question 'do you require shareback?' is asked (and then subsequently the option to donate that to a charity). There is also a link to input about differences in permissions, in a box, if either of the two lower options are selected as answers to the question 'will you allow sharing?' (these correspond to PrivacyAlike+ and PrivacyAlike- respectively, and use could be made of whitelist/blacklists when determining which organizations should be shared with).

Once initial consent is given, if the end user returns to the notice, they would be allowed to view it and also asked:

Would you like to update your choices?

- yes
- no

If the user clicks on 'yes', the smart notice would present the user's 'old' choices and allow changes, as well as adding in deletion options. In addition, customised advice can be provided. For example, a button 'see implications' included within the smart notice takes into consideration the choices selected, the data type and the general scenario and provides targetted information to the user, facilitating informed consent.

4.2.2 Enterprise Policies

The Smart Notice is achieved by means of EnCoRe providing templates, with subsets being defined by an IT administrator within the enterprise, to customise the choices offered to end users. EnCoRe also provides a wizard that interacts with the end users to allow them to easily make these choices, and hence set the Smart Notice parameters.

Enterprise policies include:

1. *Access control policies*: these include an enhanced representation for privacy, and rely on the predefined set of access control policies for the security aspects already defined within the Sapphire system.
2. *Obligation templates and policies*: these are event-driven, within the enterprise, rather than being triggered by access control.
3. *Sticky policies*: these are ongoing obligations associated with data if that is shared beyond the enterprise and that define how that data are to be used/treated
4. *Privacy compliance policies*: these include policies specifying when notifications are needed to data protection authorities and transborder data flow restriction rules.

The format of the access control policies is in general:

Target: Sample and associated data S

if (Data Requestor wants to access usage sample and associated data S for Purpose P)
and (data subject has given consent for this data)
then Allow Access
else Deny Access

Rules like this can be for samples only, data only, or for a combination of samples associated with data. Role-based access control is used to enable specification for example that 'Only people with role Y can access sensitive data in repository Z'.

Sticky policies govern the use of the associated data, and may specify the following:

- Purposes of using data (e.g. for research, transaction processing, etc.).
- Data may only be used within a given set of platforms (with certain security characteristics), a given network or a subset of the enterprise
- Other obligations and prohibitions (allowed third parties, people or processes; blacklists; notification of disclosure; deletion of data after a certain time)
- List of trusted third parties (potentially the result of a negotiation process)

The machine-readable policy may be represented in any convenient format. A simple example in an XML format is:

```
<Sticky Policy>
    <Purpose>
               Research
    </Purpose>
    <Obligation>
          <Notification>
               Yes
          </Notification>
          <Deletion>
               After 3 years
          </Deletion>
      </Obligation>
</Sticky Policy>
```

The semantics of the policies may be defined using a number of approaches, including reference to formal ontologies. A CC approach has the advantage of providing the related semantics as part of the framework, although in places this might need to be extended.

4.3 Enforcement Mechanisms

EnCoRe uses sticky policies to represent and enforce the consent and revocation preferences of end users. Negotiating, setting, changing, and enforcing sticky policies are integrated with the management of security and privacy policies. Compliance checking and auditing are integrated capabilities.

Various EnCoRe components are involved in the processing of personal data and preferences, along with their enforcement: Users disclose their personal data with privacy preferences via the personal consent and revocation assistant; the EnCoRe Privacy-aware Access Control and Obligation components enforce them, when data is accessed or disclosed to third parties; the Data Registry tracks the whereabouts of this data; the External Workflow Manager creates and attaches Sticky Policies to data, before its disclosure to third parties. This approach is applied recursively across chains of organizations. The corresponding set of functionalities and capabilities that EnCoRe provides can be provided as a set of services (for example, in the cloud) or they can be deployed as an overall stand-alone infrastructural solution.

The Smart Notice can be used as a mechanism to generate the sticky policies associated with the user's data that the EnCoRe system sends to other organizations specifying the purposes of using the data and any obligations and prohibitions, including notification and deletion after a certain time. The EnCoRe external workflow manager component could be used to control sharing of the information associated with these policies, and the data registry to record how the data has been distributed.

If the receiving party is EnCoRe-enabled, the system translates the high-level requirements expressed in the sticky policies into fine-grained access and obligation policies to be enforced along with the original privacy choices. To achieve this, the constraints specified in the Smart Notice can be enforced. If the receiving parties do not have EnCoRe-compliant systems, the external workflow manager assesses the extent to which the data can be released for a given purpose, sanitizing it before release if needed. EnCoRe administrators predefine the criteria for sanitizing data—for example, omitting some details or providing statistical information. The criteria for releasing data include evaluating the purpose for which the data was required and the outcome of risk assessment carried out on the receiving parties—for example, their ability to deliver the required privacy controls on specific data items.

To revoke consent, users edit their consent preferences through Web-based UIs. EnCoRe batches and automatically propagates these preferences throughout the system as well as beyond it to the other organizations involved, leveraging the information stored in the data registry. Organizations can apply this approach recursively to disclose information to one another.

This solution is applicable in a variety of business contexts, and it is especially valuable where sensitive information is involved. First we enable the user via a Smart Notice to define policies which are preferences or conditions about how that information should be treated. Personal, private or confidential information that is stored and used in the cloud will be associated with the machine-readable part of the Smart Notice, in such a way that we aim to prevent this being compromised. When information is processed, this is done in such a way as to adhere to these constraints. As the data is replicated or shared within the cloud in order to fulfil the service provision request, mechanisms will be in place to ensure that the customer's preferences are respected right along the chain.

Policies can be associated with data with various degrees of binding and enforcement. A variety of techniques for binding data to disclosure policies specifying or constraining how it is to be used are possible, ranging from relatively weak logical bindings to strong bindings that use cryptography to encrypt the data, and only provide the decryption key if the conditions specified by the preferences are verified. The personal data and policies can be digitally signed to provide evidence about the conditions under which the data may be used. One approach to provide a strong binding that enhances integrity is to bind policies to data by encrypting the data under a symmetric key, conditionally shared by sender and receiver (i.e. based on fulfilment of policies), and sticking the data to the policy using public key enveloping techniques similar to Public Key Cryptography Standard (PKCS) 7 (Pearson et al, 2011); we have also considered alternative mechanisms using other types of encryption (Pearson and Casassa Mont, 2011).

5 CONCLUSIONS AND FURTHER WORK

In this paper we have introduced the notion of a 'Smart Notice' as a mechanism for aiding user control and transparency of data propagation and usage within the cloud. This provides the cornerstone of a novel, simple, transparent way exploiting Privacy Commons of expressing the terms of service and the options available to the data subject. This is conceptual work in progress that is to be further developed and tested within the context of the EnCoRe project. This idea is not focused exclusively on cloud infrastructures as it is applicable within any ecosystem of service provision, but the issues are even more pressing within cloud environments due to the probable lack of trust and prior history of engagements between the service providers and the user. Indeed, lack of trust is a widely perceived barrier to moving to cloud, especially where sensitive information is involved, and so such techniques are especially needed in that domain. Furthermore, terms of service currently tend to be fixed and set in favour of cloud service providers (Mowbray, 2009) and we think it important that mechanisms are developed and rolled out that help to address this unfair balance.

In addition, cloud environments allow greater control over the flow of information compared to the public World Wide Web; in that sense, the enforcement capabilities of the Smart Notice are substantially increased in the context of cloud. The Smart Notice follows the CC paradigm but goes beyond it precisely because it allows greater control by the user over his/her personal information and more enforcement. Both CC and the Smart Notice mechanism share the same premises, i.e. greater autonomy of the end user and facilitation of the sharing of information. However, the Smart Notice approach is addressed to a data subject that is going to provide his/her data to a data controller that will then process and disseminate these data, whereas CC allows the individual to share his/her information in an unmediated fashion with the rest of the world. As the CC approach indicates, the Smart Notice and CC licences could be also used in conjunction in order to cover both privacy and IPR aspects of information dissemination. Finally, further work is required in order to explore the degree to which modules similar to the ones we have seen in the CC case could be established in the Smart Notice case as well as to explore ways in which the organisational roll out of CC (i.e. through academic institutions) could be transferred or adapted to fit the needs of the Smart Notice model.

ACKNOWLEDGEMENTS

We gratefully acknowledge feedback from various members of the EnCoRe project within discussions that contributed to the development of the ideas espoused within this paper, and most notably from Ioannis Agrafiotis, Patrizia Bertini, Mark Josephs, Nadja Kanellopoulou and Edgar Whitley.

REFERENCES

Agrawal, R., Kiernan, J., Srikant, R., Xu, Y., 2005. Xpref: a preference language for P3P. *Computer Networks,* 48(5), pp. 809-827.

Alhamad, M., Dillon, T., Chang, E., 2011. A Survey on SLA and Performance Measurement in Cloud Computing. In: *OTM 2011*, Part II, LNCS 7045, Springer-Verlag, pp. 469-477.

Andersson, C., Camenisch, J., Crane, S., Fischer-Hubner, S., Leenes, R., Pearson, S., Pettersson, J., Sommer, D., 2005. Trust in prime. In *Signal Processing and Information Technology*, pp. 552–559, IEEE.

Ardagna, C., Vimercati, S., Samarati, P., 2006. Enhancing user privacy through data handling policies. In: *DAS*, volume 4127, LNCS, pp. 224–236.

Ardagna, C. *et al.,* 2009. PrimeLife Policy Language, ACAS, W3C, http://www.w3.org/2009/policy-ws/

Article 29 Working Party, 2004. Opinion 10/2004 on more harmonised information provisions. 11987/04/EN, WP 100, http://ec.europa.eu/justice/policies/privacu/docs/wpdocs/2004/wp100_en.pdf

Becker, M.Y., Malkis, A., Bussard, L., 2009. A Framework for Privacy Preferences and Data-Handling Policies, *MSR-TR-2009-128* http://research.microsoft.com/apps/pubs/default.aspx?id=102614

Bergmann, M., Rost, M., Pettersson, J.S., 2006. Exploring the feasibility of a spatial user interface paradigm for privacy-enhancing technology. In: *Bridging the Gap between Academia and Industry*, pp. 437-448.

Bing, J., 2004. Copymarks: A suggestion for simple management for copyrighted material. In: *International Review of Law, Computers & Technology*. 18(3), pp. 347-374.

Breaux, T. & Antón, A., 2008. Analysing Regulatory Rules for Privacy and Security Requirements. IEEE Transactions on Software Engineering, 34(1), pp. 5-20.

Bussard, L. Becker, M.Y., 2009. Can access control be extended to deal with data handling in privacy scenarios? ACAS, W3C http://www.w3.org/2009/policy-ws/

Camenisch, J., Leenes, R., Sommer, D. (eds.), 2011. *Digital Privacy: PRIME – Privacy and Identity Management for Europe,* LNCS 6545, Springer.

Camenisch, J., Fischer-Hübner, S., Rannenberg, K. (eds.), 2011. *Privacy and Identity Management for Life*, Springer.

Cranor, L., 2002. *Web Privacy with P3P.* O'Reilly & Associates.

Cranor, L.F., Guduru, P., Arjula, M., 2006. User interfaces for privacy agents. *ACM Trans. Comput.-Hum. Interact.*, 13(2), pp. 135–178.

Creative Commons, 2012. http://creativecommons.org/

Creative Commons (CC), 2012b. http://creativecommons.org/choose/

Creative Commons (CC), 2012c. http://wiki.creativecommons.org/CC_REL

Creative Commons (CC), 2012d. http://creativecommons.org/licenses/by/3.0/

Creative Commons (CC), 2012e. http://creativecommons.org/licenses/by/3.0/legalcode

Damianou, N., Dulay, N., Lupu, E., Sloman, M., 2001. The Ponder Policy Specification Language http://

wwwdse.doc.ic.ac.uk/research/policies/index.shtml

Elahi, T.E., Pearson, S., 2007. Privacy assurance: Bridging the gap between preference and practice. In: *TrustBus*, pp. 65–74.

Electronic Frontier Foundation (EFF), 2012. TOSBack: The Terms of Service Tracker, http://www.tosback.org/ timeline.php

EnCoRe project, 2012. www.encore-project.info

Gellman, R., 2009. Privacy in the Clouds: Risks to Privacy and Confidentiality from Cloud Computing, *World Privacy Forum*, www.worldprivacyforum.org/pdf/WPF_Cloud_Privacy_Report.pdf

Hawkey, K., Inkpen, K. Examining the content and privacy of web browsing incidental information. In *WWW '06*, pp. 123–132, New York, NY, USA.

Holtz, L.E., Zwingelberg, H., Hansen, M., 2011. Privacy Policy Icons. In: *Privacy and Identity Management for Life*, Camenisch, J., Fischer-Hübner, S., Rannenberg, K. (eds.), Springer, pp. 279-285.

Holtz, L.E., Schallaböck, 2011. Legal Policy Mechanisms. In: *Privacy and Identity Management for Life*, Camenisch, J., Fischer-Hübner, S., Rannenberg, K. (eds.), Springer, pp. 343-354.

IBM, 2004. The Enterprise Privacy Authorization Language (EPAL), EPAL specification, v1.2, http://www.zurich.ibm.com/security/enterprise-privacy/epal/

IBM, 2006. REALM project, http://www.zurich.ibm.com/security/publications/2006/REALM-at-IRIS2006-20060217.pdf

IBM, 2007. Sparcle project, http://domino.research.ibm.com/comm/research_projects.nsf/pages/sparcle.index.html

Iachello, G., Hong, J. *End-User Privacy in Human-Computer Interaction*. Now Publishers Inc., Hanover, MA, USA, 2007.

Irwin, K., Yu, T., 2005. Determining user privacy preferences by asking the right questions: an automated approach. In *WPES '05: Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, pp. 47–50, New York, NY, USA, ACM.

Jaatun, M.G., Tϕndel, I.A., Bernsmed, K., Nyre, A.A., 2012. Privacy Enhancing Technologies for Information Control. In: *Privacy Protection Measures and Technologies in Business Organisations*, G. Yee (ed.), pp. 1-31, IGI Global.

Kelley, P.G., Bresee, J., Cranor, L.F., Reeder, R.W., 2009. A "nutrition label" for privacy. In: *SOUPS '09: Proceedings of the 5th Symposium on Usable Privacy and Security*, pp. 1–12, New York, NY, USA.

Kenny, S., Borking, J., 2002. The Value of Privacy Engineering. *Journal of Information, Law and Technology (JILT)*, 1. http://elj.warwick.ac.uk/jilt/02-/kenny.html.

Kobsa, A., 2003. A component architecture for dynamically managing privacy constraints in personalized web-based systems. In *Privacy Enhancing Technologies, Third International Workshop, PET 2003, Dresden, Germany, March 26-28, 2003, Revised Papers*, *LNCS 2760*, Springer, pp. 177–188.

MobiLife project. http://www.ist-mobilife.org/

Mowbray, M., 2009. The Fog over the Grimpen Mire: Cloud Computing and the Law. *SCRIPT-ed J Law Technol Soc*, 6(1), pp. 132-146.

OASIS, 2012. XACML, http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml

Organisation for Economic Co-operation and Development (OECD), 1980. Guidelines Governing the Protection of Privacy and Transborder Flow of Personal Data, OECD, Geneva.

Papanikolaou, N., Creese, S., Goldsmith, M., Casassa Mont, M., Pearson, S., 2010. ENCORE: Towards a holistic approach to privacy, *Proc. SECRYPT*.

Papanikalaou, N., Pearson, S., Casassa Mont, M., 2011. Towards Natural-Language Understanding and Automated Enforcement of Privacy Rules and Regulations in the Cloud: Survey and Bibliography. *Secure and Trust Computing, Data Management and Applications*, Communications in Computer and Information Science, 187, Springer, pp. 166-173.

Patrick, A.S., Kenny, S., 2003. From privacy legislation to interface design: Implementing information privacy in human-computer interactions. In *Privacy Enhancing Technologies*, *LNCS 2760*, Springer, pp. 107–124.

Parsons, C., 2010. APIs, End-Users, and the Privacy Commons, http://www.christopher-parsons.com/blog/privacy/apis-end-users-and-the-privacy-commons/

Pearson, S., 2011. Privacy Models and Languages: Assurance Checking Policies. *Digital Privacy*, ed. J. Camenisch, D. Sommer and R. Leenes, LNCS 6545, Springer, pp. 363-375.

Pearson, S., Casassa Mont, M., Chen, L., Reed, A., 2011. End-to-End Policy-Based Encryption and Management of Data in the Cloud. In: *Proc. CloudCom 2011*, IEEE.

Pearson, S., Casassa Mont, M., 2011. Sticky Policies: An Approach for Privacy Management across Multiple Parties, *IEEE Computer*, 44(9), IEEE, pp. 60-68.

Pearson, S., 2010. Addressing Complexity in a Privacy Expert System. In: E. Hüllermeier, R. Kruse, and F. Hoffmann (Eds.), *Proc. IPMU 2010*, Part II, CCIS 81, Springer, pp. 612–621.

Pettersson, J.S., et al, 2005. Making PRIME Usable. In *SOUPS '05*, New York, pp. 53–64.

Privacy Commons, 2012. Privacy Commons Incubator, http://groups.google.com/group/privacy-commons-incubator/browse_thread/thread/28c56ad776f37cb2?hl=en

Rundle, M., 2006. International data protection and digital identity management tools. Presentation at IGF 2006, Privacy Workshop 1, Athens.

Schunter, M., Waidner, M., 2003.Platform for Enterprise Privacy Practices, *PET*, LNCS 2482.

Spiekermann, S., Cranor, L., 2009. Engineering Privacy. *IEEE Transactions on Software Engineering*, 35(1), January/February.

Stanford University, 2012. WhatApp project. https://whatapp.org/about

VOME project. http://www.vome.org.uk/