HAL open science

Privacy Management in Global Organisations

Siani Pearson

To cite this version:

Siani Pearson. Privacy Management in Global Organisations. 13th International Conference on Communications and Multimedia Security (CMS), Sep 2012, Canterbury, United Kingdom. pp.217-237, 10.1007/978-3-642-32805-3_23 . hal-01540892

HAL Id: hal-01540892

https://inria.hal.science/hal-01540892v1

Submitted on 16 Jun 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed under a Creative Commons Attribution 4.0 International License

Privacy Management in Global Organisations

Siani Pearson

Cloud and Security Lab, HP Labs, Bristol. BS34 8QZ. UK.
Siani.Pearson@hp.com

Abstract. In this paper it is considered how meeting privacy requirements can be challenging for global organisations, particularly where future Internet service provision models are involved. Approaches will be explained that can be used to help address these issues, with a focus on some of the innovative solutions that the author has been involved in developing in HP Labs that are currently being used, rolled out or are the subjects of further research.

Keywords: accountability, governance, privacy impact assessment, regulation

1 Introduction

Privacy protection is currently in a state of change, as a direct result of new technologies, business models and techniques (such as cloud computing, big data processing and extended data mining, location-based services, social computing, radio-frequency identification, etc.) straining the traditional legal frameworks for privacy. In particular, more information is known, recorded and accessible, making it difficult for people not to be judged on the basis of past actions. The bulk of privacy laws across the world were created before the Internet, and this has created gaps between the guidance that laws and regulations can provide and decisions that organisations need to make about the collection and use of information.

Organisations processing personal data need to ensure that their operations are in compliance with applicable privacy regulations as well as with consumer expectations, but this can be very challenging. Contributing factors to this challenge include the factors above, as well as the growing number of privacy regulations around the world, outsourcing and transborder data flow concerns, which together challenge existing governance and security frameworks for handling personal information.

New privacy risks are emerging, and the capacity to create risk and consumer harm has increased dramatically. So, companies must find ways to integrate ethics, values and new forms of risk assessment within the organisation, as well as demonstrating responsible practices. Conforming to legal privacy requirements and meeting client privacy and security expectations with regard to personal information require organisations to demonstrate a context-appropriate level of control over such data at all stages of its processing, from collection to destruction. Privacy protection builds trust between service providers and users, and accountability and privacy by design provide mechanisms to achieve the desired end effects and create this trust. This man-

agement can span a number of layers: policy, process, legal and technological. It is universally accepted as best practice that such mechanisms should be built in as early as possible into a system's lifecycle.

Organisations need to be able to guide appropriate decisions at each stage of the product and service lifecycle. Both large and small organisations can benefit from automated solutions (such as decision support tools) that help them take privacy concerns properly into account for all relevant projects. Both large and small organisations will benefit from broad privacy knowledge encoded in the knowledge base (KB) of such decision support tools as this knowledge is becoming increasingly complex. In addition, for large organisations, tools – unlike manual processes – can scale up to handle hundreds or thousands of projects. Tools can thereby achieve a better level of assurance that most or all their projects are in compliance with regulatory standards and an organisation's policies.

In this paper a decision support tool is described that has been developed for privacy, as well as its generalisation to other compliance domains and other approaches that companies can use in order to employ best practice and be accountable.

First some background is provided about what privacy is.

1.1 What is privacy?

At the broadest level (and particularly from a European standpoint), privacy is a fundamental human right, enshrined in the United Nations Universal Declaration of Human Rights (1948) and subsequently in the European Convention on Human Rights and national constitutions and charters of rights. There are various forms of privacy, ranging from 'the right to be left alone' [1], 'control of information about ourselves' [2], 'the rights and obligations of individuals and organisations with respect to the collection, use, disclosure, and retention of personally identifiable information.' [3], focus on the harms that arise from privacy violations [4] and contextual integrity [5].

In the commercial, consumer context, privacy entails the protection and appropriate use of the personal information of customers, and the meeting of expectations of customers about its use. What is appropriate will depend on the applicable laws, individuals' expectations about the collection, use and disclosure of their personal information and other contextual information, hence one way of thinking about privacy is just as 'the appropriate use of personal information under the circumstances'[6].

Data protection is the management of personal information, and is often used within the European Union in relation to privacy-related laws and regulations (although in US the usage of this term is focussed more on security).

The terms '*personal information*' and '*personal data*' are commonly used within Europe and Asia, whereas in US the term '*Personally Identifiable Information'* (PII) is normally used, but they are generally used to refer to the same concept. This can be defined as information that can be traced to a particular individual, and include such things as: name, address, phone number, social security or national identity number, credit card number, email address, passwords, date of birth. The current European Union (EU) Definition of *personal data* is that:

> ‘personal data shall mean any information relating to an identified or identifiable natural person (‘data subject’); an identifable person is one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity.’ [7]

Some personal data elements are considered more sensitive than others, although the definition of what is considered *sensitive personal information* varies depending upon jurisdiction and even on particular regulations.

Privacy differs from security, in that it relates to handling mechanisms for personal information, although security is one element of that. Security mechanisms, on the other hand, focus on provision of protection mechanisms that include authentication, access controls, availability, confidentiality, integrity, retention, storage, backup, incident response and recovery. Privacy relates to personal information only, whereas security and confidentiality can relate to all information.

Privacy is regarded as a human right in Europe, whereas in America it has been traditionally viewed more in terms of avoiding harm to people in specific contexts. It is a complex but important notion and correspondingly the collection and processing of personal information is subject to regulation in many countries across the world.

The focus of this paper is on corporate governance related to privacy, and its structure is as follows. In the following section privacy issues for global organisations are considered. In section 3 measures are considered that corporate governance puts in place to address these issues. In section 4 a recently evolving approach is discussed that should help address privacy issues in global and complex environments, namely accountability. In section 5 it is considered how technology can help address privacy issues, and in section 6 a number of example solutions are presented. Finally, conclusions are given.

2 Privacy Issues for Global Organisations

For organisations, privacy entails the application of laws, policies, standards and processes by which personal information is managed. The fair information practices developed in US in 1970s [8] and later adopted and declared as principles by the Organisation for Economic Co-operation and Development (OECD) and the Council of Europe [9] form the basis for most data protection and privacy laws around the world. These principles can be broadly described as follows:

1. *Data collection limitation:* data should be collected legally with the consent of the data subject where appropriate and should be limited to the data that is needed.
2. *Data quality:* data should be relevant and kept accurate.
3. *Purpose specification:* the purpose should be stated at the time of data collection.
4. *Use limitation:* personal data should not be used for other purposes unless with the consent of the individual.
5. *Security:* personal data should be protected by a reasonable degree of security.
6. *Openness:* individuals should be able to find out what personal data is held and how it is used by an organisation.

7. *Individual participation:* an individual should be able to obtain details of all information about them held by a data controller and challenge it if incorrect.
8. *Accountability:* the data controller should be accountable for complying with these principles.

This framework can enable sharing of personal information across participating jurisdictions without the need for individual contracts. It imposes requirements on organisations including data collection, subject access rights and data flow restrictions.

In Europe, the European Data Protection Directive 95/46/EC (and its supporting country legislation) implements these Fair Information Principles, along with some additional requirements including transborder data flow restrictions. Other privacy-related restrictions may also be imposed (e.g. on cookie usage by the recent EU ePrivacy Directive). Legislation similar to the European Data Protection Directive has been, and continues to be, enacted in many other countries, including Australia, New Zealand, Hong Kong, Japan and APEC. Notably, legislation in Canada, Argentina, Israel, Switzerland, Guernsey, Iceland, Lichtenstein, Norway, Jersey and the Isle of Man is considered strong enough to be 'adequate' by EC. (*Adequacy* defines how a specific country is considered to have an adequate or inadequate level of protection for processing personal data of subjects from within the European Union countries.) In contrast, the US does not have a comprehensive regime of data protection but instead has a variety of laws —such as the Health Insurance Portability and Accountability Act (HIPAA) — which are targeted at the protection of particularly sensitive types of information. This US approach to privacy legislation is historically sector-based or enacted at the state level (for example, the State of Massachusetts has set out appropriate security standards for protecting the personal information of residents of that state) and places few if any restrictions on transborder data flow. The US is considered adequate for data transfer only under the limitation of the Safe Harbor agreement [10].

With regard to security (number 5. in the list above), it is a common requirement under data protection law that if a company outsources the handling of personal information or confidential data to another company, it has some responsibility to make sure the outsourcer uses "reasonable security" to protect those data. This means that any organisation creating, maintaining, using or disseminating records of PII must ensure that the records have not been tampered with, and must take precautions to prevent misuse of the information. Specifically, to ensure the security of the processing of such information, data controllers must implement appropriate technical and organisational measures to protect it against:

- *Unauthorised access or disclosure:* especially for data transmission over a network
- *Destruction:* accidental or unlawful destruction or loss
- *Modification:* inappropriate alteration
- *Unauthorised use:* all other unlawful forms of processing

Mechanisms to do this include risk assessment, implementing an information security program and putting in place effective, reasonable and adequate safeguards that cover physical, administrative and technical aspects of security.

Privacy challenges for businesses include data broaches (which can be costly (on average $204 per record, according to a 2010 Ponemon Institute study), risk of litigation due to country-specific laws, the complexity of managing privacy and negative public attention and loss of brand value if exposures occur. When customers are concerned for the welfare of their privacy (whether that be due to worries about unsolicited marketing, identity theft, surveillance, unwanted inferences about their behaviour or other reasons), it can affect a company's ability to do business.

Privacy issues depend upon the role of the company. For example, an organisation could be a custodian of employee personal data, could collect end-user personal information, or could just be providing outsourcing services for another organisation. Legally, the requirements are quite different depending upon whether the organisation is a data controller or a data processor in that situation (although it might be both).

A *data controller* is an entity (which could be a person, public authority, agency or other body) which alone, jointly or in common with others determines the purposes for which and the manner in which any item of personal information is processed, and this is legally responsible for ensuring compliance requirements are met. Obligations and risks of the data controller include: regulatory fines, criminal liability, civil liability if data subjects enforce their rights, investment risk, business continuity impact and reputational damage. In environments such as cloud computing, a data controller has a responsibility to ensure that the service providers are meeting regulatory obligations and this can be challenging [11].

A *data processor* is an entity which processes personal information on behalf and upon instructions of the data controller. Contractual agreements may add additional responsibilities or constraints with respect to privacy, although data protection laws stipulate that the organisation that is transferring personal information to a third party for processing remains responsible for the personal information. The data processor may also face issues such as lack of training of key personnel and deliberate targeting of sensitive information by criminals.

When considering privacy risks, context is an important aspect, as different information can have different privacy, security and confidentiality requirements and privacy threats differ according to the type of scenario: for example, they would tend to be higher for services that are dynamically personalised, based on people's location, preferences, calendar and social networks, etc. Privacy need be taken into account only if a service handles personal information (in the sense of collecting, transferring, processing, sharing, accessing or storing it). Even if the same information is involved, there may be different data protection requirements in different contexts, due to factors including location and trust in the entities collecting and processing it. There are special laws concerning treatment of sensitive data, and data leakage and loss of privacy are of particular concern to users when sensitive data is processed. In addition, privacy issues vary across different stages of the information lifecycle, e.g. data collection, processing, storage, archival and destruction.

Companies differ in the resources they have available to deal with privacy. Many larger organisations have a Chief Privacy Officer and privacy staff in order to implement compliance in their organisations. Smaller organisations often do not have the resources for hiring qualified privacy experts and instead the person appointed who is

responsible for overseeing the organisations's compliance with applicable privacy legislation could well be the owner or operator. Key elements of privacy management such as defining a corporate privacy policy can often be difficult to achieve in such situations. However, small companies are largely domestically bound, and hence driven by domestic legislation, except in the case for certain small companies in niche areas that might quickly become multinational. For multinational companies, requirements are more diverse and privacy management is more difficult. Nevertheless, data is an asset, so proper privacy management will be valuable for forward-thinking companies, quite apart from being mandatory from a legal point of view.

Some companies might choose to ignore the issue and pay the penalties if they are found to be in breach, but at the time of writing, regulations, enforcement activities and sanctions are currently increasing the world over. The US is introducing a Consumer Privacy Bill of Rights [12] and the EU is revising their Data Protection Directive and regulation [13], with the result that FTC enforcement will be strengthened within US and current plans are that European DPAs will be able to impose fines of up to 2% of worldwide annual turnover to companies that do not have mechanisms in place to underpin regulatory data protection compliance [13].

It has been discussed in the introduction how privacy risks are increasing, and correspondingly there is a need to push compliance and reduce risks throughout organisations, including to untrained people that might expose hundreds of files by the click of a button, lose a laptop containing unencrypted confidential information or switch sensitive information to the cloud almost instantly using a credit card. However, requirements can be complex to ascertain and a privacy staff is typically small, making effective oversight over hundreds or possibly thousands of projects per year difficult. Hence the role of both process and technology is important and in the following sections it is considered how solutions may be provided.

3 Corporate Governance for Privacy

Privacy has been a concern for mainstream corporate entities for at least a decade. Since the 1970s the primary focus of privacy has been personal information, and particularly concerned with protecting individuals from government surveillance and potential mandatory disclosure of privacy databases. In the 1980s concerns were raised related to direct marketing and telemarketing. In the late 90s there was a response in corporate governance to the activities of data protection regulators within EU, Canada, New Zealand and Australia. About ten years ago security measures were introduced to help counter the increasing threat of online identity theft, spam and phishing. More recently, governments and markets are starting to expect privacy and it is becoming a mainstream business activity.

Current best practice for creating a privacy program is to:

- garner senior management support and establish a comprehensive organisational privacy policy
- establish clear processes and assign responsibilities to individuals, including appointment of a Chief Privacy Officer and a Corporate Privacy Team

- utilise proven, existing standard and frameworks for security and IT management, such as ISO 27001/2 and ITIL, and
- establish proper monitoring and audit practices, in order to verify and assess what is happening in the organisation against the privacy policies, and take action where required to achieve alignment

More specifically, a privacy management program would ideally include the following measures [14]:

- establish reporting mechanisms and reflect these within the organisation's privacy management program controls
- put in place privacy management *program controls*, namely:
  - a *Personal Information Inventory* to allow the organisation to identify the personal information in its custody, its sensitivity and the organisation's authority for its collection, usage and disclosure
  - *policies* relating to: collection, use and disclosure of personal information (including requirements for consent and notification); access to and correction of personal information; retention and disposal of personal information; security controls and role-based access; handling complaints by individuals about the organisation's personal information handling practices
  - *risk assessment* mechanisms
  - *training and education*
  - *breach and incident management*
  - setting privacy *requirements for third parties* that handle personal information
  - procedures for *informing individuals* about their privacy rights and the organisation's program controls
- develop an *oversight and review plan* that describes how the organisation's program controls will be monitored and assessed
- *ongoing assessment and revision* of the program controls above

3.1 Privacy by Design

Privacy by Design refers to the philosophy and approach of embedding privacy into design specifications, as first espoused by Ann Cavoukian and others [15,16]. It applies to products, services and business processes. The main elements are:

1. Recognition that privacy concerns must be addressed
2. Application of basic principles expressing universal spheres of privacy protection
3. Early mitigation of privacy concerns when developing information technologies and systems, across the entire information life cycle
4. Need for qualified privacy input; and
5. Adoption and integration of privacy-enhancing technologies (PETs). These are considered further below.

In essence, companies should build in privacy protections at every stage in developing products, and these should include reasonable security for consumer data, limited

collection and retention of that data, as well as reasonable procedures to promote data accuracy.

In addition to the Canadian regulators, there has been strong emphasis and encouragement from Federal Trade Commission (FTC) and EC amongst others on usage of a privacy by design approach [13,17]. The FTC report [17] calls on companies handling consumer data to implement recommendations for protecting privacy, including greater transparency about collection and usage of consumers' information and provision of simplified choices to consumers so that they can decide what information is shared about them, and with whom. This should include a Do-Not-Track mechanism that would provide a simple and easy way for consumers to control tracking of their online activities.

Various companies have produced detailed privacy design guidelines (see for example [18]). Cannon has described processes and methodologies about how to integrate privacy considerations and engineering into the development process [19]. Privacy design guidelines in specific areas are given in [20,21], and [22] considers the case of cloud computing.

Privacy maturity models may be used to help organisations plan to improve their privacy management over time. The point of a *capability maturity model* (CMM) is generally to understand the maturity of organisations through various characteristics [23]. Such maturity models can help facilitate process development and enterprise evolution by identifying maturity milestones and benchmarks for comparison. It is possible to represent a privacy maturity model by capturing key privacy controls. A simple model for privacy risks was described in [24]. [25] builds upon this by describing a cloud capability maturity model and using it to explore privacy controls within an enterprise cloud deployment, including where there may be opportunities to design in data protection controls as exploitation of the cloud matures. Currently, consultancy is on offer to help organisations define a privacy maturity model [26].

'Privacy by policy' is the standard current means of protecting privacy rights through laws and organisational 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' [27], 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 (and preferable) to enhance that approach to cover a hybrid approach with privacy by architecture.

The Privacy by Design approach strives to reach a "positive sum", which allows privacy, accountability and transparency. This can be achieved by pseudonymity schemes that allow revocation of anonymity for misbehaving users while guaranteeing strong anonymity for honest users [28,29]. It may also be achieved by decision support and audit systems that make decision makers aware and responsible for the consequences of their actions. In November 2007 the UK Information Commissioners Office (ICO) (an organisation responsible for regulating and enforcing access to and

use of personal information), launched a Privacy Impact Assessment (PIA) [30] process (incorporating privacy by design) to help organisations assess the impact of their operations on personal privacy. This process assesses the privacy requirements of new and existing systems; it is primarily intended for use in public sector risk management, but is increasingly seen to be of value to private sector businesses that process personal data. Similar methodologies exist and can have legal status in Australia, Canada and the USA [31]. The methodology aims to combat the slow take-up to design in privacy protections from first principles at the enterprise level.

3.2 Addressing Transborder Data Flow Restrictions

One aspect that organisations need to plan for is restrictions on transborder data flow. It is not just transborder data flow requirements that restrict the flow of information across borders: there may also be trade sanctions and other export restrictions, for example restriction of cryptography and confidential data from US.

Personal information can be transferred from any EU/EEA country if model contracts have been signed and in many instances approved by the country regulator, or Binding Corporate Rules (BCRs) have been approved, or the individual has "freely given" consent. Model contracts are contractual agreements that contain data protection commitments, company liability requirements and liabilities to the individuals concerned. Transfers from other countries with national privacy legislation (e.g. Canada, Argentina) also require contractual agreement. BCRs are binding internal agreements/contracts that obligate all legal entities within a corporate group that will have access to EU personal information to adhere to all obligations of the EU Data Protection Directive.

These techniques (and especially model contracts as currently used) are not well suited to dynamic or cloud environments, because administering and obtaining regulatory approval for model contracts can result in lengthy delays: the notification and prior approval requirements for EU Model Contracts vary significantly across the EU but are burdensome and can take from one to six months to set up. BCRs are suitable for dynamic environments but their scope is limited: they only apply to data movement within a company group, it may be difficult for SMEs to invest in setting these up and there are only a few BCRs to date, although it is a relatively new technique.

4 The Role of Accountability

New approaches to privacy oversight have recently started to emerge, in the form of accountability-based programs recognised across jurisdictions and supported by regulators, society and the private sector. This approach requires greater transparency but in return, removes unnecessary burdens and so resources can be allocated instead to implementation and assurance monitoring. Even though organisations should appoint a Privacy Officer to be responsible for the organisations's privacy management programme, the organisation remains accountable for compliance with applicable privacy legislation and its accountability is not passed on to that individual [14].

4.1 The Meaning of Accountability

The term 'accountability' is susceptible to a variety of different meanings within and across disciplines. In particular, Daniel Weitzner has defined 'information accountability' as 'the claim of individuals, groups, and institutions to determine for themselves when, how, and to what extent information about them is used lawfully and appropriately by others' [32]. In general, accountability takes a principles-based approach that focuses on outcomes, and the use of information. Key elements of this notion include responsibility, transparency, remediation and validation. Accountability in relation to privacy focuses on the acceptance of responsibility for protection of personal information.

Accountability is enshrined in regulatory frameworks for data protection across the globe, notably the OECD privacy guidelines (1980) [9], Canada's Personal Information Protection and Electronic Documents Act (2000) [33] and Asia Pacific Economic Cooperation (APEC)'s Privacy Framework (2005) [34]. Within these, accountability is used in the sense that the data controller should be accountable for complying with that particular data protection legislation. Accountability concepts are evolving as the current legal framework responds to globalisation and new technologies [35], and indeed the current drafts of the new EU Data Protection Directive [13] include this concept. Region block governance frameworks such as the EU's Binding Corporate Rules (BCRs) [36] and APEC's Cross Border Privacy Rules (CBPRs) [37] are being developed to provide a cohesive and more practical approach to data protection across disparate regulatory systems. The Galway/Paris project started by privacy regulators and privacy professionals has been for the last two years defining and refining the concept of accountability in the context of the latest regulations [38]. In particular, there has been a shift towards organisations owning data processing risks on behalf of individuals, and being penalised if they are not acting responsibly.

4.2 The Shift Towards Accountable Organisations

The main motivations for an accountability-based approach is that it increases trust (for consumers, clients and regulators) and eases organisational operations. Privacy and trust come from sound information stewardship by service providers for which they are held accountable. It also decreases regulatory complexity in global business environments, especially for cloud. It can provide a clear and consistent framework of data protection rules, and allows avoidance of a complex matrix of national laws and reduces unnecessary layers of complexity for cloud providers.

One aspect for achieving accountability is to have a strong emphasis on auditing. Audit should be able to keep track of where the data has been outsourced, who processed it and for what purpose. These steps are essential in ensuring accountability and gaining user trust. Indeed, Weitzner and others have argued that to provide accountability, there must be a shift from hiding information to ensuring that only appropriate uses occur [32]. So, service providers (SPs) should maintain a history of data manipulation and inferences (providing transparency) that can then be checked against the policies that govern them, thus providing retrospective accountability.

Solutions to privacy risks involve inclusion of an element of control. For the corporate user, privacy risk can be reduced if organisations use a combination of privacy policies and contractual terms to create accountability in the form of transparent, enforceable commitments to responsible data handling. Specifically, accountable organisations will ensure that obligations to protect data (corresponding to user, legal and company policy requirements) are observed by all processors of the data, irrespective of where that processing occurs.

Through contractual agreements, all organisations involved in service provision could be accountable. While the corporate user, as the first corporate entity in the cloud provision, would be held legally accountable, the corporate user would then hold the initial service provider accountable through contractual agreements, requiring in turn that it hold its SPs accountable contractually as well. Thus, the transferor is held accountable by regulators even when it is the transferee that does not act in accordance with individuals' wishes [39,40].

Responsible company governance entails that organisations act as a responsible steward of the data which is entrusted to them within the cloud, ensuring responsible behaviour via accountability mechanisms and balancing innovation with individuals' expectations. Hence Privacy by Design may complement and incorporate corporate accountability mechanisms [41]. The Galway and Paris projects outlined core elements of implementing an accountability project within an organisation [38], which is very similar to the guidance provided by the Privacy Commissioners of Canada, Alberta and British Columbia [14] considered above, but with more emphasis on risk identification, mitigation, and redress. This is not surprising because in order to be an accountable organisation, a privacy management program needs to be rolled out within that institution. Furthermore, it is the organisation's responsibility to understand the risks and build mitigation and abatement programs into their processes as it is no longer the consumer's responsibility to isolate risks. Correspondingly, privacy maturity models shift towards assessment of systems designed to meet clear objectives.

Accountability begins to shift our thinking from only having an obligation to comply with a principle, to an obligation to prove that you can put those principles into effect. Mechanisms can be provided both for internal accountability (within an organisation, for example ensuring privacy compliance is monitored via a Privacy Office) and external accountability (providing assurance to regulators and auditors about the organisation's compliance with policies and regulations). Correspondingly, new laws and regulations [12,13] are tending to include explicit requirements that an organisation not only comply, but that they have programs that put the principles into effect. Therefore, in future companies will need to do more to ensure privacy is considered in their products and services. Technology can provide assistance in ensuring proper implementation of accountability.

5 The Role of Technology

Privacy Enhancing Technologies (PETs) can be defined (here with a UK focus) as "… any technology that exists to protect or enhance an individual's privacy, including

facilitating individuals' access to their rights under the Data Protection Act 1998" [42]. These include privacy management tools that enable inspection of service-side policies about handling of personal data, provision of user-centric choice, control, transparency, etc., audit and accountability as well as pseudonymisation tools that provide confidentiality at the network layer, and anonymisation or pseudonymisation for web browsing, email, payment, voting, etc. For example, some known technologies for web and email privacy include: spam filters, cookie blockers, pop up blockers, anti-spyware, web proxies that strip off identifying information from web traffic, anonymous remailers and Mix Nets (work started by David Chaum [43]). Other technologies are centred around privacy-enhanced identity management [44]. Different approaches depend on weak versus strong trust models and also the extent to which personal and sensitive data is actually needed to be revealed as part of the service provision. There are 'degrees of anonymity' [45] and what is most appropriate depends upon the context. A review of different types of PETs is given in [46].

Technical support for accountability can be provided in a number of areas, including: audit; risk analysis; obligation; service level agreement (SLA), trust and incident management; monitoring; policy enforcement and selective information exchange. One area where technology is very beneficial for privacy management in particular is in helping to provide risk assessment tools. An important part of any organisational privacy management programme is to conduct regular risk assessments to ensure compliance with applicable legislation and company policies. This is because privacy risks change over time, and new services might be provided that collect, user or disclose personal information and that have not been thoroughly vetted from a privacy perspective — it is much better to minimise privacy impacts in this way before deploying or changing services rather than having to fix privacy problems after they have occurred. In the following section a number of different privacy impact assessment tools developed by HP Labs are considered.

6 Example Solutions

In this section some examples of privacy accountability tools developed within HP are presented. HP has a comprehensive privacy management programme in place, including deployment of different tools and procedures for accountability, but there is not the space in this paper to describe all of these so the focus here is on examples of solutions which the author was involved in developing.

6.1 HP Privacy Advisor (HP PA)

HP PA is an intelligent online rule-driven system that assesses activities that handle personal data within HP and provides privacy by design guidance. It is a web-based decision support system used internally within HP to assess risk and degree of compliance for projects that handle personal data and to guide individual employees in their decisions on how to handle different types of data. HP PA elicits privacy-relevant information about a project via a customised sequence of questions. It uses a

dynamic interface to minimise unnecessary questions and maintains a record of activities. Based on the answers given, HP PA:

- Assesses a project's degree of compliance with corporate privacy policy, ethics and global legislation, and the privacy promises the company makes
- Integrates privacy risk assessment, education, and guidance into the process
- Scores projects for a list of ten privacy compliance indicators including transborder data flows, compliance, business controls, security, transparency, and so forth
- Generates a compliance report for each project and, if appropriate, notifies an appropriate member of the corporate privacy team for further guidance/intervention
- Provides checklists, reminders, customised help and warnings to users.

The scores for different rules in the output report and the compliance indicators can be green (signifying no privacy issues), yellow (indicating possible privacy risks) or red (indicating the project could violate a regulatory requirement or company policy).

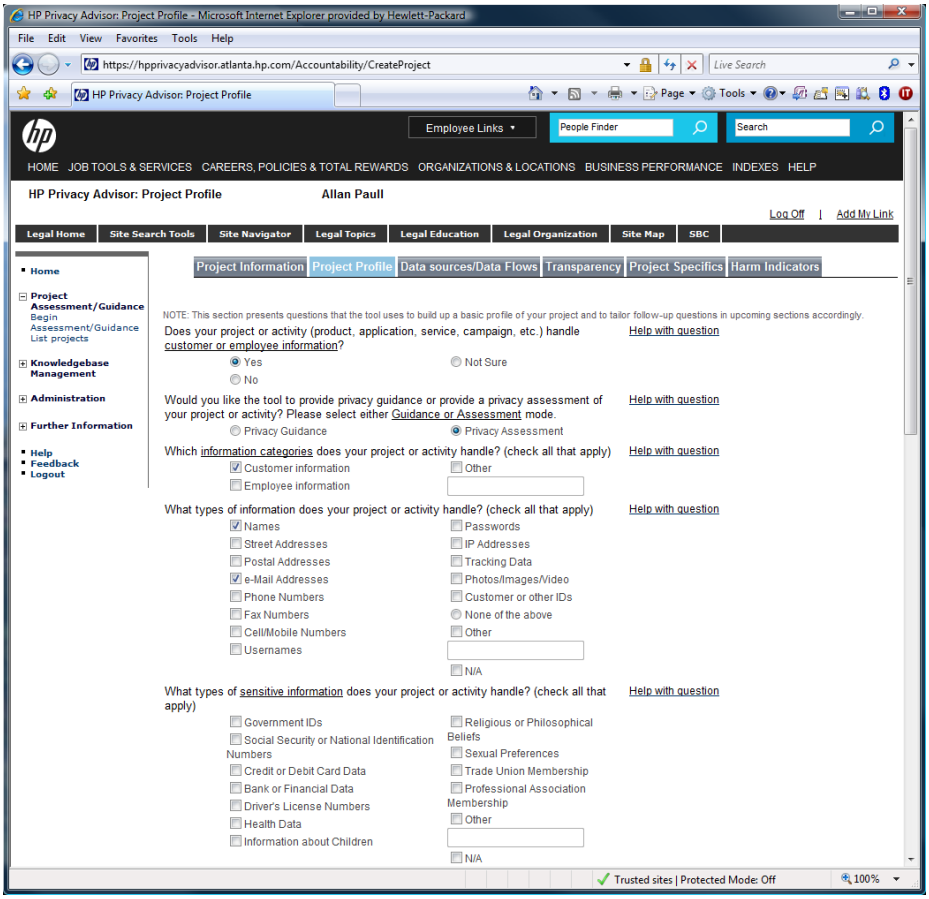

Fig. 1. Dynamic questionnaire

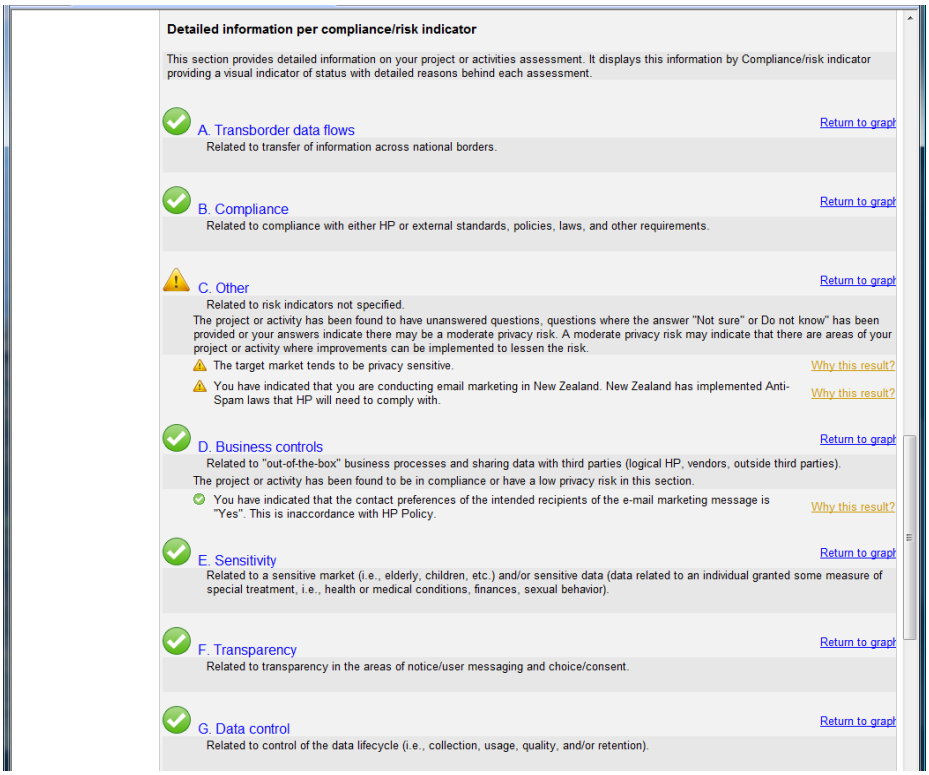

Fig. 2. Part of report

As the assessment is designed to be detailed, a broad range of privacy risks within a project will be flagged up. The distribution of these risks and their severity as indicated by the yellow or red flags gives a good understanding about which privacy risk a project carries.

To use HP PA, employees access a web-based tool and answer a questionnaire. Employees can use the tool to carry out an assessment or obtain privacy guidance for their project. They use *Privacy Assessment* mode if they are about to deploy (or have already deployed) the project and need to determine if it complies with privacy requirements. Alternatively, *Privacy Guidance* mode is used when they are developing a project and need information on how to ensure it will meet privacy requirements. Using the link in the left navigation area takes them to a questionnaire, as shown in Figure 1. After having filled in the project information, the Project Profile section is used to gather a profile of the project. It is used by HPPA to build the remainder of the questionnaire. The first question is a gating question used to determine if the questionnaire needs to be answered. The questionnaire will continually be built as the user answers each question, and will be recalculated dynamically using a rules engine if answers are changed. If the user moves the mouse over underlined text, tool tips are used to display a definition of that term: this is especially helpful for explaining privacy-related terms. Blocks of related questions are grouped together, for readability.

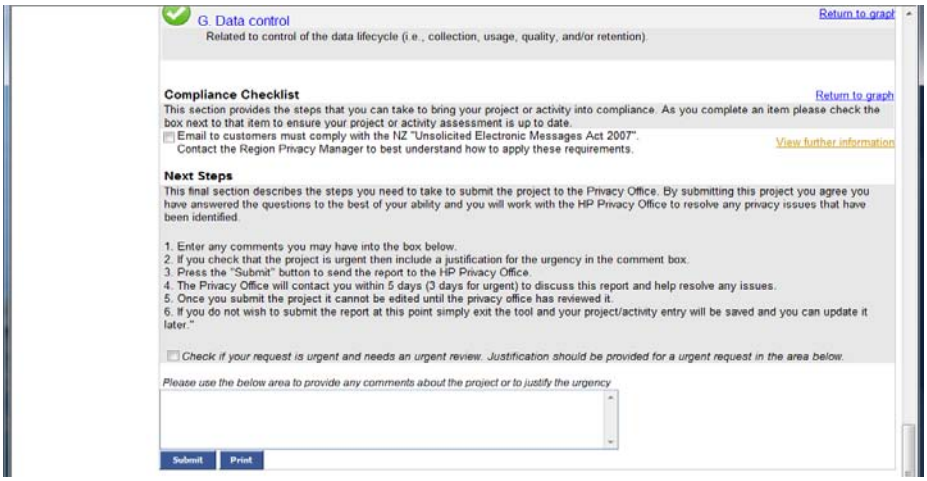

Fig. 3. Action checklist and report submission

Help is available on any question to clarify its meaning, and warnings and informational messages can be associated with any question answers. A 'question is unclear' option in the questionnaire allows administrators to identify questions that users find difficult to understand or answer, and furthermore any unanswered questions are highlighted and the user is made to provide an answer. The user can navigate to any part of the questionnaire using the section tabs.

The assessment report contains several sections that display: the report status; instructions on how to use the report; the project information; an assessment summary. A compliance and risk indicators graph displays a graphical representation of the assessment, showing the number of compliant or low risk (green) responses, "not sure" answers or moderate risk responses and non-compliant or high risk responses. As shown in Figure 2, detailed information is provided on the compliance and risk indicators. Clicking on a 'Why this result?' link displays a window showing details for the reason for the assessment. Part of the report is a Compliance Checklist, which lists actions an employee can take to bring their project into compliance (cf. Figure 3).

An employee can enter a message for the Privacy Office Approver and can indicate if their project is urgent or not. They can submit the project and report for assessment by the Privacy Office. To ensure the integrity of the project once submitted the project is locked and cannot be altered until at a later stage in the workflow, i.e. when the Privacy Office has reviewed the project and unlocks it. Employees can also add additional documents to the project to assist in the assessment, and print the report.

To access projects (at any stage of submission), a list of projects for which the user has permission to view can be displayed (see Figure 4). Projects can be edited, viewed, deleted (in some circumstances) and shared with other members of a team in order to help complete the assessment. HP PA has several layers of access, depending upon whether the user is an employee, a compliance officer (who can access and approve projects, and amend KB content), or an administrator.

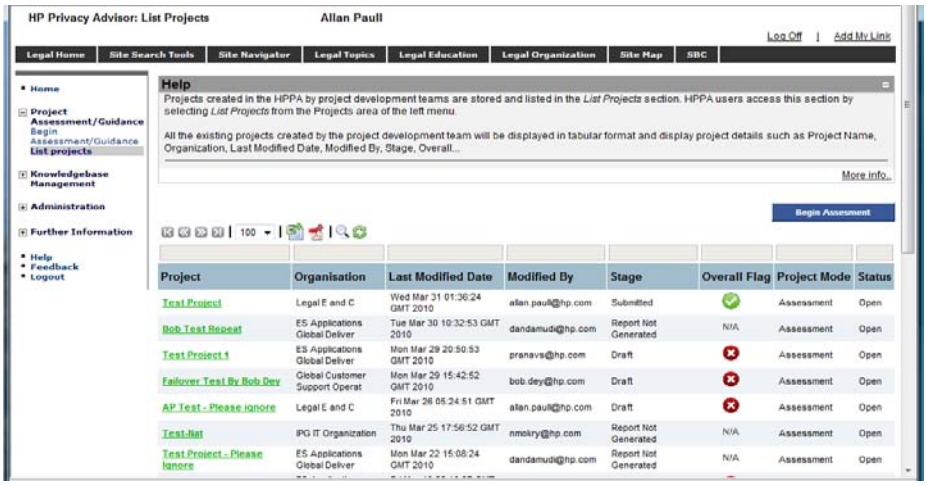

Fig. 4. List of projects

A dashboard view for compliance officers shows graphs and statistics about submitted projects based on business units and/or regions. A range of other controls are provided, including for administration, Privacy Office review and approval and knowledge base (KB) management. Further details about the underlying knowledge representation of HP PA and KB management is given in [47,48]. HP PA provides the KB management via an Expert Mode and a user-friendly Simple Mode that both can be used by domain experts to flexibly customise several aspects of the tool. Using this, the 300-page HP privacy rule book was encoded within the knowledge base (KB) of HP PA – as well as other policy documents – and extensive testing confirmed that it can be used to risk-assess projects within HP that might pose privacy risks. HP PA has been deployed and is available to all HP employees.

HP PA helps HP deal with large volumes of projects and address accountability requirements. HP PA is integrated into standard business processes so that key decisions must pass through the tool before project deployment. At predefined periods, users with non-compliant projects are reminded of their responsibilities to bring their project into compliance. In addition, formal HP Privacy Office analysis and review of project reports is undertaken, via a workflow generated via the tool.

The core technology that underpins this tool is the result of a joint effort by HP Labs and the HP Privacy Office. The major areas of technical innovation are in knowledge representation and inference and in simplifying knowledge management. In particular, an accurate representation of organisational privacy policies is provided, desirable system properties are ensured such as deterministic behavior of questionnaire and report generation, tailoring, and completeness of the questionnaire generation [47,48]. Heuristics and usability techniques have been encapsulated in order to aid non-trained users in creating the knowledgebase (KB) and have addressed complexity, including the ramifications of KB updates, KB versioning and quarantine of parts of the KB [48,49]. Thereby, complex compliance tasks and processes have been

automated within a system that is reliable, comprehensive, and simple, balancing complexity of analysis with ease of use.

HP PA provides a framework for a comprehensive regulatory compliance environment, in that it can be used as the basis for a flexible compliance tool that can be used across multiple different domains (e.g., security, compliance, finance, healthcare, etc.), as considered further in the following section.

6.2 Other Privacy Risk Assessment Tools

In order to improve governance practices and reduce organisational risk, the author has been involved in research on various other accountability mechanisms for risk assessment, namely:

Regulatory compliance manager for financial services: this assesses global privacy, bank secrecy and cross border data movement. The core decision support technology is integrated with other system components that include: workflow, document management and project management systems and an external reporting engine, to provide a broad compliance and audit environment. The workflow can be driven by the output of the assessment engine (e.g. for dependency of the workflow on the project risk level as determined by the assessment engine). The KB representation is enhanced to allow for more sophisticated authoring and display of questionnaires. It provides an end-to-end system for accountability for all stages of a project lifecycle that is industry agnostic, to be delivered either via a SaaS model or as a standalone instance.

Decision support system for business process outsourcing: elucidating global privacy requirements corresponding to deal pursuit and due diligence phases and suggestion of corresponding privacy and security controls [50]; this system was deployed within HP but is superseded by HP PA, which now includes outsourcing.

UK privacy impact assessment tool for organisations based upon ICO guidelines related to UK Data Protection Act, allowing appropriate stakeholder views and input and using confidences within the knowledge representation to allow assessment of the value of the input as well as customisation of risk indicator values [51].

Tools for cloud assessment: privacy impact assessment of cloud environments [52] and decision support tools for cloud service provisioning [53].

6.3 Additional Accountability Mechanisms

Apart from the examples considered above, the author is engaged in researching and developing a number of accountability mechanisms:

– *monitoring for information use*: this can occur at different levels [54,55]

- *data obfuscation:* a trade-off can be made of efficiency against security, to obfuscate some of the data before transferring it for processing, using a key that is not revealed to the service provider, and with the degree of this obfuscation dependent upon the context [56].
- *consent management*: consumer preferences are gathered about usage of information and these are then mapped to machine readable policies associated with data. Privacy-enhanced access control and obligation management are then used to help enforce these machine-readable policies and link data usage to auditing [57]
- *sticky policies:* a preventive technique that can be used to provide a chain of accountability. Machine-readable policies are bound to data so that obligations are kept travelling with data along the service provision chain. Access to data can be as fine-grained as necessary, based on policy definitions, underlying encryption mechanisms (supporting the stickiness of policies to the data) and a related key management approach that allows (sets of) data attribute(s) to be encrypted specifically based on the policy [58]. Access to data is mediated by a Trust Authority that checks for compliance to policies in order to release decryption keys. Strong enforcement of such policies is still research in progress that typically requires trusted infrastructure [59,60].

These mechanisms can be used independently or in combination; for example, obligations that apply to a given situation may be deduced with the help of a decision support system and then automatically enforced and monitored.

7 Conclusions

An explanation has been given of why privacy management can be challenging in global organisations and the importance of accountability and technology in addressing this problem has been highlighted. An important new approach is for co-design of legal, procedural and technical mechanisms to provide accountability.

Some examples of tools for privacy management have been provided that have recently been developed and deployed within HP, including a tool for use by employees that asks contextual questions and outputs guidance on specific requirements for compliance with laws, regulations, ethics, and company values. The author is engaged in ongoing research, interactions and collaborations with regulators, policy makers, academics and other institutions on a number of accountability mechanisms.

Acknowledgements

HP PA and RCA are the result of collaboration between HP Labs, HP Privacy Office and HP divisions, involving an extended team with input from many individuals.

References

1. Warren, S., Brandeis, L.: The Right to Privacy. 4 Harvard Law Review 193 (1890)

2. Westin, A.: Privacy and Freedom. New York, USA, Atheneum (1967)
3. American Institute of Certified Public Accountants (AICPA) and CICA: Generally Accepted Privacy Principles. August (2009)
4. Solove, D.J.: A Taxonomy of Privacy. University of Pennyslavania Law Review, 154(3):477, January (2006)
5. Nissenbaum, H.: Privacy as Contextual Integrity. Washington Law Review, pp. 101-139 (2004)
6. Swire, P,, Bermann, S.: Information Privacy. Official Reference for the Certified Information Privacy Professional, CIPP (2007)
7. European Commission (EC): Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data. (1995)
8. Privacy Protection Study Commission: Personal Privacy in an Information Society, United Statues Privacy Protection Study Commission Fair Information Practices. (1977)
9. Organisation for Economic Co-operation and Development (OECD): Guidelines for the Protection of Personal Data and Transborder Data Flows. (1980)
10. Safe Harbor website, `http://export.gov/safeharbor/`
11. Pearson, S.: Privacy, Security and Trust in Cloud Computing. In: Privacy and Security for Cloud Computing, Computer Communications and Networks, S. Pearson and G. Yee (eds.), Springer (2012)
12. The White House: Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy, February (2012)
13. European Commission: Proposal for a Directive of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data by competent authorities for the purposes of prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, and the free movement of such data, January (2012)
14. Office of the Information and Privacy Commissioner of Alberta, Office of the Privacy Commissioner of Canada, Office of the Information and Privacy Commissioner for British Colombia: Getting Accountability Right with a Privacy Management Program. April (2012)
15. Cavoukian, A.: Privacy by Design: Origins, Meaning, and Prospects for Assuring Privacy and Trust in the Information Era. In: Privacy Protection Measures and Technologies in Business Organisations: Aspects and Standards, G. Yee (ed), pp. 170-208, IGI Global (2012)
16. Information Commissioners Office (ICO): Privacy by Design. Report, `www.ico.gov.uk` (2008)
17. Federal Trade Commission (FTC): Protecting Consumer Privacy in an Age of Rapid Change: Recommendations for Business and PolicyMakers. FTC Report, March (2012)
18. Microsoft Corporation: Privacy Guidelines for Developing Software Products and Services, Version 2.1a (2007)
19. Cannon, J.C.: Privacy: What Developers and IT Professionals Should Know. Addison Wesley (2004)
20. Patrick, A., Kenny, S.: From Privacy Legislation to Interface Design: Implementing Information Privacy in Human-Computer Interactions. R. Dingledine (ed.), PET 2003, LNCS 2760, pp. 107-124, Springer-Verlag Berlin Heidelberg (2003)
21. Belloti, V., Sellen, A.: Design for Privacy in Ubiquitous Computing Environments. Proc. 3rd conference on European Conference on Computer-Supported Cooperative Work, pp. 77-92 (1993)

22. Pearson, S.: Taking Account of Privacy when Designing Cloud Computing Services. In: ICSE-Cloud'09, Vancouver. IEEE (2009)
23. Wikipedia, http://en.wikipedia.org/wiki/Capability_Maturity_Model
24. The Institute of Internal Auditors: Managing and Auditing Privacy Risks, http://www.theiia.org/download.cfm?file=33917
25. Creese, S., Hopkins, P., Pearson, S., Shen, Y.: Data Protection-Aware Design for Cloud Computing", Proc. 1st CloudCom 2009, M.G. Jaatun, G. Zhao, C. Rong (eds.), Beijing, Springer LNCS 5931, pp. 119-130, December (2009)
26. Minnesota Privacy Consultants, http://www.minnesotaprivacy.com/
27. Spiekermann, S., Cranor, L. F.: Engineering privacy. IEEE Transactions on Software Engineering, pp. 1-42. IEEE (2008)
28. Chaum, D., Fiat, A., Naor, M.: Untraceable Electronic Cash, CRYPTO 1988, pp. 319-327 (1988)
29. Camenisch, J., Groß, T., Heydt-Benjamin, T.: Accountable privacy supporting services, Identity in the Information Society, 2(3), Springer, pp. 244-267 (2009)
30. Information Commissioner's Office: PIA handbook. http://www.ico.gov.uk/ (2007)
31. Office of the Privacy Commissioner of Canada: Privacy impact assessments (2007)
32. Weitzner, D.J., Abelson, H., Berners-Lee, T., Feigenbaum, J., Hendler, J., Sussman, G.J.: Information accountability. Communications of ACM, 51: 6, p. 87, June (2008)
33. PIPEDA (2000) http://laws-lois.justice.gc.ca/eng/acts/P-8.6/
34. Asia-Pacific Economic Co-operation (APEC): APEC Privacy Framework, Dec (2005) http://publications.apec.org/publication-detail.php?pub_id=390
35. European Commission (EC) EU Article 29 Working Party: Opinion 3/2010 on the principle of accountability. WP 173, July (2010)
36. Information Commissioner's Office (ICO): Binding Corporate Rules http://www.ico.gov.uk/for_organisations/data_protection/overseas/binding_corporate_rules.aspx
37. Australian Government Department of the Prime Minister & Cabinet (PM&C): APEC Cross-Border Privacy Rules, http://www.dpmc.gov.au/privacy/apec/cross-border.cfm
38. Center for Information Policy Leadership (CIPL): Demonstrating and Measuring Accountability: A Discussion Document. Accountability Phase II –The Paris Project (2010)
39. Crompton, M., Cowper, C., Jefferis, C.: The Australian Dodo Case: an insight for data protection regulation. World Data Protection Report, 9:1, BNA (2009)
40. Pearson, S., A. Charlesworth: Accountability as a Way Forward for Privacy Protection in the Cloud. Proc. CloudCom 2009, M.G. Jaatun, G. Zhao, C. Rong (eds.), Beijing, Springer LNCS 5931, pp. 131-144, December (2009)
41. Cavoukian, A., Taylor, S., Abrams, M.: Privacy by Design: Essential for Organisational Accountability and Strong Business Practices. Identity in the Information Society. 3:2, pp. 405-413 (2010)
42. Information Commissioner's Office UK (ICO): Data protection guidance note: Privacy enhancing technologies. (2007)
43. Chaum, D.: Intraceable electronic mail, return addresses and digital pseudonyms. Communications of the ACM (1981)
44. Camenisch, J., Fischer-Hubner, S., Rannenberg, K. (eds.): Privacy and Identity Management for Life, Springer (2011)

45. Reiter, M.K., Rubin, A.D.:Anonymous Web transactions with Crowds. Commun. ACM, 42, Feb, pp. 32-48 (1999)
46. Shen, Y., Pearson, S.: Privacy Enhancing Technologies: A Review. HPL-2011-113, http://www.hpl.hp.com/techreports/2011/HPL-2011-113.html
47. Pearson, S., Rao, P., Sander, T., Parry, A., Paull, A., Patruni, S., Dandamudi-Ratnakar, V., Sharma, P.: Scalable, Accountable Privacy Management for Large Organisations. INSPEC 2009, IEEE, pp. 168-175, September (2009)
48. Pearson, S.: Addressing Complexity in a Privacy Expert System. E. Hüllermeier, R. Kruse, and F. Hoffmann (eds.), IPMU 2010, Part II, CCIS 81, pp. 612–621, Springer-Verlag Berlin Heidelberg (2010)
49. Pearson, S., Sander, T.: A Decision Support System for Privacy Compliance. In: Threats, Countermeasures, and Advances in Applied Information Security, Manish Gupta, John Walp, and Raj Sharman (eds.), Information Science Reference, IGI Global, New York, pp. 158-180 (2012)
50. Pearson, S., Sander, T., Sharma, R.: Privacy Management for Global Organisations. Data Privacy Management and Autonomous Spontaneous Security, J. Garcia-Alfaro et al. (eds.), LNCS 5939, pp. 9-17, Springer-Verlag Berlin Heidelberg (2010)
51. Tancock, D., Pearson, S., Charlesworth, A.: Analysis of Privacy Impact Assessments within Major Jurisdictions. *Proc. PST 2010*, Ottawa, Canada, IEEE, August (2010)
52. Tancock, D., Pearson, S., Charlesworth, A.: Privacy Impact Assessments for Cloud Computing. Privacy and Security for Cloud Computing, Computer Communications and Networks, Springer, 2012.
53. Sander, T., Pearson, S.: Decision Support for Selection of Cloud Service Providers, *International Journal on Computing (JoC)*, GTSF, 2010.
54. Pearson, S., Allison, D.: Privacy Compliance Checking using a Model-Based Approach. E-Business Applications for Product Development and Competitive Growth: Emerging Technologies, I. Lee (ed.), Business Science Reference, IGI Global, pp. 199-220 (2011)
55. Ko, R., Lee, B.S., Pearson, S.: Towards achieving accountability, auditability and trust in cloud computing. A. Abraham et al. (eds.), ACC 2011, Part IV, CCIS 193, pp. 432–444, Springer-Verlag Berlin Heidelberg (2011)
56. Mowbray, M., Pearson, S., Shen, Y.: Enhancing Privacy in Cloud Computing via Policy-based Obfuscation. Journal of Supercomputing, March (2010)
57. EnCoRe project, www.encore-project.info
58. Pearson, S., Casassa Mont, M.: Sticky Policies: An Approach for Privacy Management across Multiple Parties. IEEE Computer, 44: 9, pp. 60-68, September (2011)
59. Trusted Computing Group, http://www.trustedcomputinggroup.org
60. Pearson, S., Casassa Mont, M., Novoa, M.: Securing Information Transfer within Distributed Computing Environments. IEEE Security & Privacy Magazine, Jan/Feb issue, 6:1, pp. 34-42, IEEE (2008)