Accountability as a Way Forward for Privacy Protection in the Cloud

Siani Pearson[1] and Andrew Charlesworth[2]

[1] HP Labs, Long Down Avenue, Stoke Gifford, Bristol, UK BS34 8QZ
[2] Centre for IT and Law, University of Bristol, Queens Road, Bristol, UK BS8 1RJ
Siani.Pearson@hp.com, a.j.charlesworth@bris.ac.uk

Abstract. The issue of how to provide appropriate privacy protection for cloud computing is important, and as yet unresolved. In this paper we propose an approach in which procedural and technical solutions are co-designed to demonstrate accountability as a path forward to resolving jurisdictional privacy and security risks within the cloud.

Keywords: Accountability, cloud computing, privacy.

1 Introduction

Cloud computing is a means by which highly scalable, technology-enabled services can be easily consumed over the Internet on an as-needed basis [1]. The convenience and efficiency of this approach, however, comes with privacy and security risks [2]. A significant barrier to the adoption of cloud services is thus user fear of confidential data leakage and loss of privacy in the cloud. Furthermore, the cross-jurisdictional nature of clouds presents a new challenge in maintaining the data protection required by current legislation including restrictions on cross-border data transfer.

At the broadest level, privacy is a fundamental human right that encompasses the right to be left alone, although an analysis of the term is complex [3]. 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. For organisations, privacy entails the application of laws, policies, standards and processes by which Personally Identifiable Information (PII) of individuals is managed.

We focus in this paper on: privacy in the sense of data protection, as defined by Directive 95/46/EC [4] (rather than the narrower US sense of data security); data that is PII (information that can be traced to a particular individual, such as a phone number or social security number); the corporate entity seeking to contract for services in the cloud, either for its own use, or to offer to its customers, as this entity is most likely to have resources to use our proposed path of technical and procedural solutions. However, our solution is not EU-specific, and is compatible with privacy principles underlying American and Asia-Pacific regulation and legislation, as well as a self-regulatory approach.

M.G. Jaatun, G. Zhao, and C. Rong (Eds.): CloudCom 2009, LNCS 5931, pp. 131–144, 2009.

This paper proposes the incorporation of complementary regulatory, procedural and technical provisions that demonstrate accountability into a flexible operational framework to address privacy issues in this cloud computing scenario. The structure of the paper is as follows: consideration of open issues that relate to cloud computing and privacy; an explanation of accountability and how this might apply in cloud computing; proposal of legal mechanisms, procedures and technical measures that tie in with this approach; an assessment of this approach and conclusions.

2 Privacy Issues for Cloud Computing

Privacy is a key business risk and compliance issue, as it sits at the intersection of social norms, human rights and legal mandates [5]. Conforming to legal privacy requirements, and meeting client privacy expectations with regard to PII, require corporations to demonstrate a context-appropriate level of control over such data at all stages of its processing, from collection to destruction. The advantages of cloud computing – its ability to scale rapidly (through subcontractors), store data remotely (in unknown places), and share services in a dynamic environment – can thus become disadvantages in maintaining a level of privacy assurance sufficient to sustain confidence in potential customers. For example:

- **Outsourcing.** Outsourcing of data processing invariably raises governance and accountability questions. Which party is responsible (statutorily or contractually) for ensuring legal requirements for PII are observed, or appropriate data handling standards are set and followed [6]? Can they effectively audit third-party compliance with such laws and standards? To what extent can processing be further sub-contracted, and how are the identities, and *bona fides*, of sub-contractors to be confirmed? What rights in the data will be acquired by data processors and their sub-contractors, and are these transferable to other third parties upon bankruptcy, takeover, or merger [7]? 'On-demand' and 'pay-as-you-go' models may be based on weak trust relationships, involve third parties with lax data security practices, expose data widely, and make deletion hard to verify.
- **Offshoring.** Offshoring of data processing increases risk factors and legal complexity [8]. Issues of jurisdiction (whose courts can/will hear a case?), choice of law (whose law applies?) and enforcement (can a legal remedy be effectively applied?) need to be considered [9]. A cloud computing service which combines outsourcing and offshoring may raise very complex issues [10].
- **Virtualization.** There are security risks in sharing machines, e.g. loss of control over data location, and who has access to it. Transactional data is a byproduct with unclear ownership, and it can be hard to anticipate which data to protect. Even innocuous-seeming data can turn out to be commercially sensitive [11].
- **Autonomic technology.** If technological processes are granted a degree of autonomy in decision making, e.g. automatically adapting services to meet changing needs of customers and service providers, this challenges enterprises' abilities to maintain consistent security standards, and to provide appropriate business continuity and back-up, not least as it may not be possible to determine with any specificity where data processing will take place within the cloud [12].

As cloud computing exhibits all the aspects above, privacy solutions need to address a combination of issues, and this may require new and even unique mechanisms rather than just a combination of known techniques for addressing selected aspects. For example, privacy problems when transferring PII across borders within a group of companies can be addressed via Binding Corporate Rules, and yet this approach would not be available to a corporation seeking to adopt a cloud computing solution where PII will be handled by third party cloud service providers.

Overall, the speed and flexibility of adjustment to vendor offerings, which benefits business and motivates cloud computing uptake, brings a higher risk to data privacy and security. This is a key user concern, particularly for financial and health data.

2.1 Mapping Legal and Regulatory Approaches

Effective corporate governance is vital to compliance with the type of regional block regulatory governance models which underpin Binding Corporate Rules in Europe and Cross Border Privacy Rules in Asia-Pacific Economic Cooperation (APEC) countries. Organizations that process PII must safeguard it (including limiting its use and disclosure) or face legal, financial and reputational penalties. Where there are inadequate information governance capabilities in the cloud, this will severely restrict the outsourcing of key business processes using cloud-based service marketplaces.

Companies and governmental organisations are increasingly aware of the need to integrate privacy into the technology design process [13, 14]. However, tools and technical controls alone cannot fully address privacy issues in cloud computing, due to diverse privacy obligations upon, and privacy practices within, organisations [15]. Cloud service providers (SPs) and marketplace providers need to design their processes to ensure those obligations and practices can be mapped against a combination of technical and procedural measures, which together provide broad assurance that appropriate contextual safeguards apply to PII processing in the cloud. Context is key to requirements. Identical information collected in different contexts by different entities might involve totally divergent data protection criteria [16, 17].

Such a mapping exercise requires an understanding of the rationales for, and the objectives of, the protection of PII, and how these translate to the cloud computing environment. If cloud computing is to reach its full potential, where customers are willing to entrust PII to such a service marketplace, these criteria need to be met:

1. a determination of risks and requirements involved in a given interaction situation e.g. consideration of the underlying legal, policy and social context.
2. a determination of what protective measures (procedural and/or technological) are appropriate, based on this information.
3. effective ways of providing assurance and auditing that potential partners protect PII, in accordance with contextually appropriate protective measures.
4. a degree of transparency, in the sense of visibility into the data protection obligations and processes of potential suppliers.

Requirements arising from applying privacy legislation to the cloud are considered in [2]. A key issue is the need to respect cross-border transfer obligations. As this is particularly difficult to ensure within cloud computing, it is suggested that legislation

will need to evolve to allow compliance in dynamic, global environments. The notion of accountability is likely to provide a way forward, as discussed in this paper.

3 Accountability: A Way Forward

In this section we examine what accountability is and how we believe accountability and corporate responsibility with regard to the use of PII might be applicable in cloud computing. In doing so, we present how accountability can help fill the gaps identified above. Finally, we explain what procedural measures are needed, and the basis of a technological approach to provide accountability.

3.1 What Is Accountability?

It is important to clearly define what is meant by 'accountability' as the term is susceptible to a variety of different meanings within and across disciplines. For example, the term has been used for a number of years in computer science to refer to an imprecise requirement that is met by reporting and auditing mechanisms (see for example, [18]). In this paper the context of its use is corporate data governance (the management of the availability, usability, integrity and security of the data used, stored, or processed within an organization), and it refers to the process by which a particular goal – the prevention of disproportionate (in the circumstances) harm to the subjects of PII – can be obtained via a combination of public law (legislation, regulation), private law (contract), self-regulation and the use of privacy technologies (system architectures, access controls, machine readable policies).

To date, national and international privacy protection approaches have been heavily influenced by public law, and premised upon 'command and control' regulatory strategies. However, such legislative and regulatory mechanisms have declined in effectiveness as technological developments render the underlying regulatory techniques obsolete. Effective privacy protection for PII in some business environments is thus heavily compromised, and the ability of organizations to meaningfully quantify, control, and offset, their business risk is significantly impeded.

> It enjoins upon 'data controllers' a set of largely procedural requirements for their processing activities, and therefore conveys the impression that formal compliance will be enough to legitimise their activities. It encourages a box-ticking mentality, rather than a more systemic, and systematic, approach to fulfilling its values. [19]

The EU data protection regime, in particular, lacks effective regulatory responses for key developing technologies, such as mobile e-commerce and cloud computing [20]. Equally, self-regulation, in isolation, has failed to gain traction as a plausible alternative for effective privacy protection, with weak risk assessment and limited compliance checking [21].

Accountability in our sense will be achieved via a combination of *private* and *public* accountability. Public accountability is derived from an active interaction between: subjects of PII; regulatory bodies, such as Information Commissioners; data controllers. It is premised upon highly transparent processes. Private accountability, in contrast, is derived from the interaction between data controllers and data processors, and

is premised on contract law, technological processes, and practical internal compliance requirements. The objective of such accountability is not to meet 'a set of largely procedural requirements for … processing activities' but rather to reduce the risk of disproportionate (in context) harm to the subjects of PII, and thus reduce or permit the amelioration of negative consequences for the data controller. It reflects an acceptance that absolute reduction of harm to the subjects of PII is an impossible goal in a disaggregated environment, such as a cloud service, and that the ability to respond flexibly and efficiently (or systemically and systematically) to harms arising will provide a more efficient form of privacy protection than enforcing blunt and/or static 'tick-box' compliance criteria.

Weitzner *et al* have previously used the term "information accountability" to refer to checking 'whether the policies that govern data manipulations and inferences were in fact adhered to' [22]. Our usage of the term 'accountability' differs from this to the extent that adherence to policy becomes less critical than achieving a proportionate and responsive process for reacting to context-dependent privacy risks.

Crompton *et al* note that in contrast to the EU's 'adequacy' regime, 'accountability' is increasingly popular in jurisdictions such as Australia, Canada and the US [23]. As discussed below, accountability in this context means placing a legal responsibility upon an organization that uses PII to ensure that contracted partners to whom it supplies the PII are compliant, wherever in the world they may be. Our accountability model reflects the basic premise of this approach, but expands upon it in suggesting ways in which organizations might take the 'accountability' approach further in order to develop a reflexive privacy process.

3.2 How Accountability Might Provide a Way Forward for Privacy Protection within Cloud Computing

Solutions to privacy risks in the cloud involve reintroducing an element of control. For the corporate user, privacy risk in cloud computing can be reduced if organisations involved in cloud provision use a combination of privacy policies and contractual terms to create accountability in the form of transparent, enforceable commitments to responsible data handling [2, 19]. 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 organizations involved in the cloud provision would 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 (SP1) accountable through contractual agreements, requiring in turn that SP1 hold its SPs accountable contractually as well. This is analogous to some existing cases in outsourcing environments, where the transferor is held accountable by regulators even when it is the transferee that does not act in accordance with individuals' wishes [23].

The following elements are key to provision of accountability within the cloud:

- **Transparency.** Individuals should be adequately informed about how their data is handled within the cloud and the responsibilities of people and organisations in relation to the processing of PII should be clearly identified. As with other

disaggregated data environments, transparency in cloud computing is important not only for legal and regulatory reasons, but also to avoid violation of social norms [24]. In the context of this paper, transparency means a level of openness about an entity's handling of PII that permits meaningful accountability.

- **Assurance.** The corporate user provides assurance and transparency to the customer/client through its privacy policy, while requiring similar assurances from the SP through contractual measures and audits.
- **User trust.** Accountability helps foster user trust. When it is not clear to individuals why their personal information is requested, or how and by whom it will be processed, this lack of control will lead to suspicion and ultimately distrust [25]. There are also security-related concerns about whether data in the cloud will be adequately protected [6].
- **Responsibility.** Most data protection regimes require a clear allocation of responsibility for the processing of PII, as existing regulatory mechanisms rely heavily upon user and regulator intervention with responsible parties. Disaggregated data environments, e.g. mobile e-commerce and cloud computing, can hinder determination of that responsibility. Predetermining responsibility, via contract, as information is shared and processed within the cloud, pre-empts perceptions of regulatory failure, which may erode user trust. It also permits companies to assess their trading risks in terms of potential financial losses and data privacy breaches. This knowledge can be used to establish organisational and group privacy and security standards, and to implement due diligence/compliance measures which conform to regulatory parameters, but which are otherwise negotiable between contracting organisations, based on relevant operational criteria [20].
- **Policy compliance.** Accountability helps ensure that the cloud service complies with laws, and also the mechanisms proposed in this paper help compliance with cloud provider organisational policies and user preferences, and with auditing.

With a legal and regulatory approach, location is paramount to enforcement. With accountability, location either becomes less relevant to the customer/client because of assurances that data will be treated as described regardless of jurisdiction or becomes transparent through contracts specifying where data processing will take place. In the accountability model, the corporate user works with legal and regulatory bodies to move data between jurisdictions through mechanisms such as Binding Corporate Rules and intra-company agreements. For the corporate user, the flexibility to move customer/client data between jurisdictions has a big impact on cost.

With accountability, regulators enforce the law on the 'first in the chain' in regard to the misdeeds of anybody in the chain, including those further along. However, whether any regulatory framework will be effective depends upon a number of characteristics including the background of the regulator (country, resources available to prosecute, etc.). This approach is more effective if action can be taken against an organization that has a presence in the regulator's home jurisdiction.

Accountability is included in various privacy frameworks, including Canada and USA and the APEC privacy framework. In the EU it applies in the restricted sense that data controllers (DCs) are directly responsible for the actions of their data processors (DPs) (and thus for clouds of DPs and sub-DPs). The difference in approaches becomes more obvious where there are multiple DCs; if these are responsible separately (DCs in common, but not joint DCs) it is hard to police via the EU model, as

the data subject (DS) may be unable to identify and enforce rights against a specific DC in a cloud computing environment with a mix of DCs and DPs.

The key issue in responsibility (and accountability) terms under EU law is who is making the decision about the particular processing purpose, and not who is carrying out the processing. A central problem in the mobile e-commerce and cloud computing environments is that it is unclear to the DS if, and if so, where, a breach is taking place, so that they can enforce rights against the relevant DC. The contractual approach provides a mechanism for avoiding that accountability-negating uncertainty, in a manner which permits the DC to demonstrate compliance with the substantive law (and boost user trust), without undue reliance upon the flawed mechanism in the legislation. The accountability process is expanded outwards by the initial DC to DPs and other DCs by contract, then information that the initial DC derives from the accountability processes can be passed upwards to the regulator and downwards to the DS, so that both can perform the functions envisaged by the legislation.

In conclusion, accountability can play a role in ensuring that laws that apply to cloud computing are enforced. There is a role for regulators in the form of criminal penalties for misuse. Also, there is a role for technology, as considered below.

3.3 Procedural Approach

Procedural is used here in the sense of governance, business practices (e.g. strong privacy policies) and contractual agreements. Privacy policies can be defined at a number of levels and be reflected within internal and external corporate policy statements, contracts, Service Level Agreements (SLAs), security policies, etc. Policies are passed on when sharing information with third parties and organisational policies are used to help ensure legal compliance. In general, they should be based upon established privacy principles, such as the OECD privacy principles [26] and regulatory requirements specific to the region(s) in which the company is operating.

For our approach, cloud computing providers should move away from terms and conditions of service towards contracts between the client and the initial service provider (SP), and between that SP and other cloud providers. This approach is consistent with industry self-regulation (for example, Truste certification [27]). At issue in cloud computing is that most policies have a clause that frees the company to change its policy at any time, often, but not always, with some form of notice. These clauses may need to be re-examined in the cloud environment, where data is perhaps not as easily destroyed or returned to its original owner.

The corporate user has options that the consumer does not in using contracts as a governance measure for control within the cloud environment. Contractual obligations are those imposed on an entity by incorporation in a contract of similar legally binding agreement between that entity and other party. The corporate user has experience in using contracts to control offshoring and outsourcing relationships. These experiences can be leveraged in the cloud.

SLAs for the cloud are still being developed and there are still a number of open issues [28]. SLAs can be informal or formal with the former being more in the nature of a promise than a contract and the latter being ancillary to a contract between parties, with breach of an SLA term not being in general as severe as a breach of contract. Moreover, third parties (i.e users) would not easily be able to rely on the terms of an SLA between a cloud computing company and a corporation selling such services

onwards (i.e. the customer), as there are processes for varying the terms, without the need to renegotiate the whole SLA with customers.

Nevertheless, specific contractual agreements can be used between the cloud provider and the corporate user, just as contracts are used today with traditional SPs. SPs can pass on obligations to subcontractors via contracts – they would require written permission to subcontract with agreements that must be no less restrictive than the agreement the corporate user has with the SP, and reserve the right to enter at will into additional confidentiality agreements directly with the subcontractors. Such contracts have to be plausibly capable of supporting meaningful enforcement processes, and capable of at least some degree of meaningful oversight/audit. The contracts can be used to:

1. address the issue of location – by requiring prior written consent for transfers to any third country
2. restrict use of data
3. prevent copying or reproducing of data without express written permission, except as technically necessary to fulfil the agreement (e.g. backup protection)
4. restrict employee access to the associated data (e.g. on a need to know basis), require that the SP provide employee privacy training, and require employees to sign confidentiality agreements
5. specify security levels – at least the same level of care applied to the SP's own similar data, but not less than a reasonable level of care, implementation of any security measures required by applicable laws
6. require immediate notification by specified means (e.g., via telephone with written follow-up), for any suspected data breach, and cooperation in resolving
7. reserve the right to audit
8. require upon request or at termination, that PII be delivered back to the data controller or data subject, and all copies be destroyed.

3.4 Co-design Involving Technological Approach

We now explain our technological approach and how it ties in with the procedural approach.

The direction in which we are carrying out research is to underpin the procedural approach above with a technological approach that helps provide accountability. In this, natural language policies in the contract are associated with lower-level policies that are machine-readable and that can be acted upon automatically within the cloud without the need for human intervention. These policies define the usage contraints of the associated PII. In this approach, as with Weitzner's approach [22], the data is accessible, but its usage is constrained. The main problem in the cloud is how this can be enforced: one option is a Creative Commons-type approach [29], where holders are made aware of their obligations and their behaviour can be audited with regard to this. If more enforcement is required, obligation management and identity management [30] could be used to manipulate data and aid data minimisation, deletion and management of notifications to individuals, but it is difficult to envisage how such a technical solution could work within non-constrained cloud environments.

Although we do not in general hide the data within the cloud, there is still the possibility to obscure it in some contexts: for example, sensitive data can in some cases be obfuscated in the cloud [31] and multi-party security (zero knowledge) techniques can be used [32].

In our approach, the machine-readable policies would include preferences or conditions about how PII should be treated (for example, that it is only to be used for particular purposes, by certain people or that the user must be contacted before it is used). When PII is processed, this is done in such a way as to adhere to these constraints. Existing policy specification, modelling and verification tools that can be used as a basis for this representation include EPAL [33], OASIS XACML [34], W3C P3P [35] and Ponder [36]. Policies can be associated with data with various degrees of binding and enforcement. Trusted computing and cryptography can be used to stick policies to data and ensure that that receivers act according to associated policies and constraints, by interacting with trusted third parties [37, 38]. Strong enforcement mechanisms include Digital Rights Management (DRM) techniques [39] and enforceable 'sticky' electronic privacy policies [37, 40].

Accountability and good privacy design go together, in that privacy protecting controls should be build into different aspects of the business process. This should be a reflexive process in that it is underpinned by a non-static compliance mechanism that is an ongoing process of privacy review throughout the contractual chain. There will be developmental, contractual and technical processes in play that encourage an organisation's cloud contractors to review and improve their privacy standards on an ongoing basis – this discourages 'cheating' by contractors, rewards effective privacy protections, and prioritises the prevention of disproportionate (in context) privacy harms over inconsequential, (in context) privacy harms. This contrasts with the application of privacy protection in a 'box-ticking fashion', where checking 'our contractor is 'adequate' according to this set of static criteria' is likely to either waste resources on low risk privacy harms or fail to identify developing high risk privacy harms. Audit information can be produced, e.g. by logging, usage of third parties and tracking [41]. In particular, sticky policy techniques can be used to ensure an audit trail of notification and disclosure of data to third parties [37]. Third party certifiers or auditors can periodically verify data protection controls, and also underpin a range of new accountability-related services that offer a cloud computing infrastructure assurances as to the degree of privacy offered (e.g. analogous to privacy seal provision for web services [42] and mechanisms for privacy assurance on the service provider side [43]).

It is necessary to utilize security techniques within the cloud to protect PII from unauthorised access or modification, and to protect backup, protect and manage multiple data copies and delete PII. To limit who has access to personal information within an organisation, privacy-aware access control [44] deployed within that organization can make decisions and enforce access control policies, intercepting queries to data repositories and returning sanitized views (if any) on requested data.

Policy enforcement within the cloud is a difficult issue – and a combined technical and procedural approach to this is preferable. The strongest binding between these would be if the wording in the contracts can be translated into machine-readable policies that are bound to data, and then enforced within the cloud. However, this binding cannot be an exact one from laws to human-readable policies to machine-readable policies, due to interpretation of the law, and furthermore only a restricted part of this translation process can be easily automated. Translation of legislation/regulation to machine readable policies has proven very difficult, although there are several examples of how translations of principles into machine readable/actionable policies can be done, e.g. Privacy Incorporated Software Agent (PISA) project [45] (deriving and modelling privacy principles from [27]); Sparcle project [46], (transforming natural

based policies into XML code that can be utilized by enforcement engines); REALM project [47] (translating high level policy and compliance constraints into machine readable formats); Breaux and Antón [48] (extracting privacy rules and regulations from natural language text); OASIS LegalXML [49] (creating and managing contract documents and terms).

Our approach is to add a technical descriptor at the bottom of a contract that describes what a cloud SP should do. For example, there could be a policy text in words that forms part of the contract, then a legal XML expression corresponding to this also within the contract [46]. Also, there could be a mapping from legal XML expression to a policy associated with data covered by the contract, and this policy might be expressed in a language like XACML [34]. However, there are currently gaps between these layers, so further work is needed to allow and provide an automatic translation. In addition, the mapping needs to be agreed, perhaps involving a third party to pre-define clauses and their meanings. A similar approach could be taken to that proposed for assurance control policies [43], to avoid having to use a fine-grained ontological approach. In general, there is a tension between flexibility of expression and ease of understanding of such policies. There is a role for standardization as these technical policies need to be understood by multiple parties so that they can be dealt with and enforced by policy decision points and policy enforcement points within the cloud infrastructure. Current technical policies of this type are access control policies, obligations and security policies. More work needs to be done in defining these, and we are working on this within the Encore project [50].

As an extension of this approach, there can be a role for infomediaries, e.g. as a Trust Authority [37], to check policies apply before allowing the decryption of data, and to play a role in auditing at this point. They could help check the situation before authorising access to personal information, e.g. via IBE [37], or else using secret sharing techniques where the decryption key is broken down and shared between multiple parties, a certain number of whom need to agree in order to be able to build up the decryption key, in a process that exploits Shamir's key sharing algorithm (analogous to the approach used in [51]). Potentially, privacy infomediaries [52] could be used in other ways that help provide accountability, e.g. by acting as insurance brokers and paying claims in case of privacy breaches. Who plays the role of privacy infomediary could vary according to the context; it could be a trusted identity provider for a federated set of services, a web proxy at an enterprise boundary, or a consumer organisation.

Mechanisms for checking compliance will be a mixture of procedural and technical, involving both auditing and regulatory aspects. There is also a role for risk and trust assessment (including reputation management) [53] before issuing contracts, to help satisfy regulators that best practice is being carried out, and in measuring metrics specified within SLAs. Decision support tools might be useful for lawyers representing the cloud providers, and to determine appropriate actions that should be allowed and to assess risk before PII is passed on (this could be part of a Privacy Impact Assessment [54]). In addition automated access control decision-making could incorporate privacy policy checking.

If trusted infrastructure [38,,55] were available within the cloud, it could help: ensure that the infrastructural building blocks of the cloud are secure, trustworthy and compliant with security best practice; determine and provide assurance regarding

location [56]; provide a basis for enhanced auditing of platforms [38, 55]. Furthermore, trusted virtual machines [57] can support strong enforcement of integrity and security policy controls over a virtual entity; for different groups of cloud services, there could be different personae and virtualized environments on each end user device.

4 Analysis of Our Approach

We believe accountability is a useful basis for enhancing privacy in many cloud computing scenarios. Corporate management can quickly comprehend its links with the recognized concept of, and mechanisms for achieving, corporate responsibility. An effective approach will require a combination of procedural and technical measures to be used and co-designed. In essence, this would use measures to link organisational obligations to machine readable policies, and mechanisms to ensure that these policies are adhered to by the parties that use, store or share that data, irrespective of the jurisdiction in which the information is processed (ideally, with a technical basis for enforcement backing up contractual assurances that incorporate privacy). Companies providing cloud computing services would give a suitable level of contractual assurances, to the organisation that wishes to be accountable, that they are can meet the policies (i.e. obligations) that it has set, particularly PII protection requirements. Technology can provide a stronger level of evidence of compliance, and audit capabilities.

While our approach can provide a practical way forward, it has limitations. First, while contracts provide a solution for an initial SP to enforce its policies along the chain, risks that cannot be addressed contractually will remain. For example, data generally has to be unencrypted at the point of processing, creating a security risk and vulnerability due to the cloud's attractiveness to cybercriminals. Secondly, only large corporate users are likely to have the legal resources to replace generic SLAs with customized contracts. Finally, adding requirements to the vendor chain will increase the cost of the service. Use of contracts will be most effective for more sensitive or more highly regulated data that merits additional and more costly protection. We believe that this approach should be scalable.

Accountability is not a substitute for data protection laws, nor would our approach render other approaches for privacy enhancement unnecessary; rather, it is a practical mechanism for helping reduce end user privacy risk and enhance end user control.

5 Conclusions

The current regulatory structure places too much emphasis on recovering if things go wrong, and not enough on trying to get organizations to 'do the right thing' for privacy in the first place. Provision of a hybrid accountability mechanism via a combination of legal, regulatory and technical means leveraging both public and private forms of accountability could be a practical way of addressing this problem; it is a particularly appropriate mechanism for dealing with some of the privacy issues that arise and are combined within cloud computing. Specifically, we advocate a co-regulation strategy based on a corporate responsibility model that is underpinned primarily by contract, and which thus places the onus upon the data controller to take a more proactive approach to ensuring compliance, but at the same time works to encourage cloud service vendors and their subcontractors to compete in the service

provision arena, at least in part, on the basis of at least maintaining good, and ideally evolving better, privacy enhancing mechanisms and processes. Further work needs to be done to effectively realize this approach, and we are continuing research in this area within the Encore project [50].

Acknowledgments. This work has greatly benefitted from input by Stacy Martin and MariJo Rogers, and broadly related discussions with colleagues from Encore project, Marty Abrams, Malcolm Crompton, Paul Henrion, Wayne Pauley and Scott Taylor.

References

1. HP cloud website, http://h71028.www7.hp.com/enterprise/us/en/technologies/cloud-computing.html?jumpid=ex_r2858_us/en/large/tsg/go_cloud
2. Pearson, S.: Taking Account of Privacy when Designing Cloud Computing Services. In: ICSE-Cloud 2009, Vancouver. IEEE, Los Alamitos (2009); HP Labs Technical Report, HPL-2009-54 (2009), http://www.hpl.hp.com/techreports/2009/HPL-2009-54.html
3. Solove, D.J.: A Taxonomy of Privacy. University of Pennsylvania Law Review 154(3), 477–564 (2006)
4. Council Directive 95/46/EC: On the protection of individuals with regard to the processing of personal data and on the free movement of such data. OJ, L281, pp. 31–50 (1995)
5. Ackerman, M., Darrell, T., Weitzner, D.: Privacy in Context. Human Computer Interaction 16(2), 167–176 (2001)
6. Cloud Security Alliance: Security Guidance for Critical Areas of Focus in Cloud Computing (2009), http://www.cloudsecurityalliance.org/guidance/csaguide.pdf
7. Gellman, R.: Privacy in the Clouds: Risks to Privacy and Confidentiality from Cloud Computing. World Privacy Forum (2009), http://www.worldprivacyforum.org/pdf/WPF_Cloud_Privacy_Report.pdf
8. Abrams, M.: A Perspective: Data Flow Governance in Asia Pacific & APEC Framework (2008), http://ec.europa.eu/justice_home/news/information_dossiers/personal_data_workshop/speeches_en.htm
9. Kohl, U.: Jurisdiction and the Internet. Cambridge University Press, Cambridge (2007)
10. Mowbray, M.: The Fog over the Grimpen Mire: Cloud Computing and the Law. Script-ed Journal of Law, Technology and Society 6(1) (April 2009)
11. Hall, J.A., Liedtka, S.L.: The Sarbanes-Oxley Act: implications for large-scale IT outsourcing. Communications of the ACM 50(3), 95–100 (2007)
12. McKinley, P.K., Samimi, F.A., Shapiro, J.K., Chiping, T.: Service Clouds: A Distributed Infrastructure for Constructing Autonomic Communication Services. In: Dependable, Autonomic and Secure Computing, pp. 341–348. IEEE, Los Alamitos (2006)
13. Microsoft Corporation: Privacy Guidelines for Developing Software Products and Services, v2.1a (2007), http://www.microsoft.com/Downloads/details.aspx?FamilyID=c48cf80f-6e87-48f5-83ec-a18d1ad2fc1f&displaylang=en
14. Information Commissioners Office: Privacy by Design, Report (2008), http://www.ico.gov.uk

15. Bamberger, K., Mulligan, D.: Privacy Decision-making in Administrative Agencies. University of Chicago Law Review 75(1) (2008)
16. Nissenbaum, H.: Privacy as Contextual Integrity. Washington Law Review 79(1), 119–158 (2004)
17. 6, P.: Who wants privacy protection, and what do they want? Journal of Consumer Behaviour 2(1), 80–100 (2002)
18. Cederquist, J.G., Conn, R., Dekker, M.A.C., Etalle, S., den Hartog, J.I.: An audit logic for accountability. In: Policies for Distributed Systems and Networks, pp. 34–43. IEEE, Los Alamitos (2005)
19. UK Information Commissioner's Office A Report on the Surveillance Society (2006)
20. Charlesworth, A.: The Future of UK Data Protection Regulation. Information Security Technical Report 11(1), 46–54 (2006)
21. Charlesworth, A.: Information Privacy Law in the European Union: E. Pluribus Unum. or Ex. Uno. Plures. Hastings Law Review 54, 931–969 (2003)
22. Weitzner, D., Abelson, H., Berners-Lee, T., Hanson, C., Hendler, J.A., Kagal, L., McGuinness, D.L., Sussman, G.J., Waterman, K.K.: Transparent Accountable Data Mining: New Strategies for Privacy Protection. In: Proceedings of AAAI Spring Symposium on The Semantic Web meets eGovernment. AAAI Press, Menlo Park (2006)
23. Crompton, M., Cowper, C., Jefferis, C.: The Australian Dodo Case: an insight for data protection regulation. World Data Protection Report 9(1) (2009)
24. Dolnicar, S., Jordaan, Y.: Protecting Consumer Privacy in the Company's Best Interest. Australasian Marketing Journal 14(1), 39–61 (2006)
25. Tweney, A., Crane, S.: Trustguide2: An exploration of privacy preferences in an online world. In: Cunningham, P., Cunningham, M. (eds.) Expanding the Knowledge Economy. IOS Press, Amsterdam (2007)
26. Organization for Economic Co-operation and Development: Guidelines Governing the Protection of Privacy and Transborder Flow of Personal Data. OECD, Geneva (1980)
27. Truste: Website (2009), http://www.truste.org/
28. SLA@SOI: Website (2009), http://sla-at-soi.eu/
29. Creative Commons: Creative Commons Home Page (2009), http://creativecommons.org
30. Casassa Mont, M.: Dealing with privacy obligations: Important aspects and technical approaches. In: Katsikas, S.K., López, J., Pernul, G. (eds.) TrustBus 2004. LNCS, vol. 3184, pp. 120–131. Springer, Heidelberg (2004)
31. Mowbray, M., Pearson, S.: A Client-Based Privacy Manager for Cloud Computing. In: Proc. COMSWARE 2009. ACM, New York (2009)
32. Yao, A.C.: How to Generate and Exchange Secrets. In: Proc. FoCS, pp. 162–167. IEEE, Los Alamitos (1986)
33. IBM: The Enterprise Privacy Authorization Language (EPAL), EPAL specification, v1.2 (2004), http://www.zurich.ibm.com/security/enterprise-privacy/epal/
34. OASIS: XACML, http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml
35. Cranor, L.: Web Privacy with P3P. O'Reilly & Associates, Sebastopol (2002)
36. Damianou, N., Dulay, N., Lupu, E., Sloman, M.: The Ponder Policy Specification Language (2001), http://wwwdse.doc.ic.ac.uk/research/policies/index.shtml
37. Casassa Mont, M., Pearson, S., Bramhall, P.: Towards Accountable Management of Identity and Privacy: Sticky Policies and Enforceable Tracing Services. In: Mařík, V., Štěpánková, O., Retschitzegger, W. (eds.) DEXA 2003. LNCS, vol. 2736, pp. 377–382. Springer, Heidelberg (2003)

38. Pearson, S.: Trusted computing: Strengths, weaknesses and further opportunities for enhancing privacy. In: Herrmann, P., Issarny, V., Shiu, S.C.K. (eds.) iTrust 2005. LNCS, vol. 3477, pp. 305–320. Springer, Heidelberg (2005)
39. Kenny, S., Korba, L.: Applying Digital Rights Management Systems to Privacy Rights Management Computers & Security 21(7) (2002)
40. Tang, Q.: On Using Encryption Techniques to Enhance Sticky Policies Enforcement. TR-CTIT-08-64, Centre for Telematics and Information Technology, Uni. Twente (2008)
41. Golle, P., McSherry, F., Mironov, I.: Data Collection with self-enforcing privacy. In: CCS 2006, Alexandria, Virginia, USA. ACM, New York (2006)
42. Cavoukian, A., Crompton, M.: Web Seals: A review of Online Privacy Programs. In: Privacy and Data Protection (2000), http://www.privacy.gov.au/publications/seals.pdf
43. Elahi, T., Pearson, S.: Privacy Assurance: Bridging the Gap between Preference and Practice. In: Lambrinoudakis, C., Pernul, G., Tjoa, A.M. (eds.) TrustBus. LNCS, vol. 4657, pp. 65–74. Springer, Heidelberg (2007)
44. Casassa Mont, M., Thyne, R.: A Systemic Approach to Automate Privacy Policy Enforcement in Enterprises. In: Danezis, G., Golle, P. (eds.) PET 2006. LNCS, vol. 4258, pp. 118–134. Springer, Heidelberg (2006)
45. Kenny, S., Borking, J.: The Value of Privacy Engineering. JILT, 1 (2002), http://elj.warwick.ac.uk/jilt/02-1/kenny.html
46. IBM: Sparcle project, http://domino.research.ibm.com/comm/research_projects.nsf/pages/sparcle.index.html
47. IBM: REALM project, http://www.zurich.ibm.com/security/publications/2006/REALM-at-IRIS2006-20060217.pdf
48. Travis, D., Breaux, T.D., Antón, A.I.: Analyzing Regulatory Rules for Privacy and Security Requirements. Transactions on Software Engineering 34(1), 5–20 (2008)
49. OASIS: eContracts Specification v1.0 (2007), http://www.oasis-open.org/apps/org/workgroup/legalxml-econtracts
50. EnCoRe: Ensuring Consent and Revocation project (2008), http://www.encore-project.info
51. Flegel, U.: Pseudonymising Unix Log Files. In: Davida, G.I., Frankel, Y., Rees, O. (eds.) InfraSec 2002. LNCS, vol. 2437, pp. 162–179. Springer, Heidelberg (2002)
52. Gritzalis, D., Moulinos, K., Kostis, K.: A Privacy-Enhancing e-Business Model Based on Infomediaries. In: Gorodetski, V.I., Skormin, V.A., Popyack, L.J. (eds.) MMM-ACNS 2001. LNCS, vol. 2052, pp. 72–83. Springer, Heidelberg (2001)
53. Pearson, S., Sander, T., Sharma, R.: A Privacy Management Tool for Global Outsourcing. In: DPM 2009 (2009)
54. Warren, A., Bayley, R., Charlesworth, A., Bennett, C., Clarke, R., Oppenheim, C.: Privacy Impact Assessments: international experience as a basis for UK guidance. Computer Law and Security Report 24(3), 233–242 (2008)
55. Trusted Computing Group (2009), https://www.trustedcomputinggroup.org
56. Pearson, S., Casassa Mont, M.: A System for Privacy-aware Resource Allocation and Data Processing in Dynamic Environments. In: I-NetSec 2006, vol. 201, pp. 471–482. Springer, Heidelberg (2006)
57. Dalton, C., Plaquin, D., Weidner, W., Kuhlmann, D., Balacheff, B., Brown, R.: Trusted virtual platforms: a key enabler for converged client devices. Operating Systems Review 43(1), 36–43 (2009)