

Data Protection-Aware Design for Cloud Services

Sadie Creese¹, Paul Hopkins¹, Siani Pearson², and Yun Shen²

¹ International Digital Laboratory, University of Warwick, Coventry, UK.

² HP Labs, Long Down Avenue, Bristol, UK. BS34 8QZ

{Sadie.Creese, P.D.Hopkins}@warwick.ac.uk,

{Siani.Pearson, Yun.Shen}@hp.com

Abstract. The Cloud is a relatively new concept and so it is unsurprising that the information assurance, data protection, network security and privacy concerns have yet to be fully addressed. This paper seeks to begin the process of designing data protection controls into clouds from the outset so as to avoid the costs associated with bolting on security as an afterthought. Our approach is firstly to consider cloud maturity from an enterprise level perspective, describing a novel capability maturity model. We use this model to explore privacy controls within an enterprise cloud deployment, and explore where there may be opportunities to design in data protection controls as exploitation of the Cloud matures. We demonstrate how we might enable such controls via the use of design patterns. Finally, we consider how Service Level Agreements (SLAs) might be used to ensure that third party suppliers act in support of such controls.

Keywords: Data protection, information security, privacy, cloud computing, design pattern, capability maturity model.

1 Introduction

Cloud computing offers a utility model for IT, enabling users to access applications, middleware and hardware via the Internet as opposed to owning it themselves. The vision for the Cloud is one where applications, platforms and infrastructure can all be consumed as and when required. The ability to rapidly scale-up and scale-down is perceived by many to directly lead to cost savings. Other benefits include fast access to new applications, easier ability to try things out before large-scale investment and staying on the leading edge. ‘Cloud nirvana’ is a future where cloud service providers (SPs) utilise the cloud to deliver dynamic capability enhancements, resources are switched on and off like taps, and users can switch suppliers quickly in order to access the best solution on the market. Current expectations of the market potential remain high, with Gartner predicting a services market value of \$150bn by 2013 [1].

The adoption of cloud services will vary across enterprises and users. Early take-up appears to be within the technology sector with other potential users voicing concerns surrounding security and privacy of data. Undoubtedly, any model which involves data assets residing on equipment not within users’ immediate control needs to address security and privacy. In ‘cloud nirvana’ environments this will only become more

acute, and potentially more challenging. Current recommendations and approaches to information security in the cloud are essentially based on today's best practice surrounding traditional outsourcing. Certainly, this is an obvious and valid starting point, and one which is recognised by those operating in the data-centre and secure-hosted service space, since they already possess the relationships, infrastructure, and business models which could easily be extended into a cloud service domain.

However, the cloud vision does offer some particularly challenging privacy problems that are unlikely to be sufficiently addressed by today's best practice [2]. Privacy can be thought of as a human right that is rather complex to analyse [3]. We focus in this paper on the issue of privacy in the sense of data protection (processing of data on identifiable living people), as defined by Directive 95/46/EC [4]. There are differing interpretations of what this may mean in a practical sense, since users of cloud services are likely to have varying expectations of confidentiality, control, and service responsiveness in response to their changing privacy requirements. The protection of these expectations will be met to equally varying degrees by the legal and regulatory structures in operation, which in themselves could vary as a cloud service could transcend national boundaries.

We seek here to begin addressing whether there are opportunities to *design-in* data protection during this cloud start-up phase, so avoiding costly future bolt-ons and sub-optimal protection resulting from design decisions in conflict with data protection needs. We cannot provide a complete treatment in a paper of this size; instead we focus on three aspects of the cloud deployment lifecycle: Firstly, we consider cloud adoption at the enterprise level and the likely maturity characteristics, developing a novel capability maturity model for enterprises exploiting cloud services. We use this maturity model as a basis for identifying opportunities for designing in privacy, and capture this analysis in a privacy maturity model. Secondly, we consider the design stage for a cloud service and how we might use design patterns to enable enterprises to adopt data protection design principles. Finally, we consider how the use of third party suppliers of cloud services might impact upon privacy, and how the associated risks might be mitigated via the use of Service Level Agreements (SLAs).

2 Related Work

Whilst there is no existing published work directly considering how to develop mechanisms for designing in data protection controls in the cloud, there are a range of work areas upon which our research is based. We discuss these here.

The point of a *capability maturity model* (CMM) is generally to understand the maturity of organisations through various characteristics: see [5] for detailed definition and history. Such maturity models can help facilitate process development and enterprise evolution by identifying maturity milestones and benchmarks for comparison. Thus, it is possible to plan, prioritise and invest in order to progress along the maturity model until the most effective and beneficial state is achieved for the enterprise. It should be noted that it is unlikely always to be the case that a higher maturity leads to greater profit in a commercial organisation, or that cloud deployment makes sense for every application (see [6]). By considering a maturity model for cloud exploitation we hope to identify the key developmental stages for a number of enterprise characteristics,

which in turn will have implications for information security and data protection strategies. Hence it may be possible to anticipate future needs and begin delivering techniques for architecting data-protection aware clouds. A number of Cloud maturity models have been proposed, for example see [7,8,9]. In [7] a cloud maturity model is presented that is specifically aimed at data centres. Whilst it offers inspiration when considering a model for exploitation of cloud by an enterprise, it cannot be directly applied. Wardley [9] implies that to achieve cloud maturity the following are likely to exist (accumulatively as maturity grows): resilient architecture, SLAs, an option to run the service in-house, evidential portability between SPs, third party assurance and monitoring of services, a marketplace of providers with easy switching, third party management of cloud market exploitation. However, as for [7], the detail is missing.

Dr. Dobb's Jake Sorofman [8] proposes a slightly different model where: the lowest level of cloud maturity involves adoption of virtualisation for seamless portability of applications and a shared server infrastructure; level two is cloud experimentation where a cloud is deployed (internally or externally) based on controlled and bounded deployments; level three is cloud foundations where governance, controls, procedures, policies and best practice begin to form initially focused on internal and non-mission critical applications; level four cloud advancement sees the scaling up of the volume of cloud applications and broad-based deployments; and level five is cloud actualisation where dynamic workload balancing occurs across multiple utility clouds and applications are distributed based on cloud capacity, cost and proximity to users. This does not break down maturity across enterprise characteristics, but is the closest to what we require and, in combination with a detailed capability model for service-oriented architectures (SOAs) designed by IBM (see [10]) forms the foundation upon which we design our new capability maturity model for cloud.

Privacy design techniques are not a new concept: various companies, notably Microsoft [11], have produced detailed privacy design guidelines. Cannon has described processes and methodologies about how to integrate privacy considerations and engineering into the development process [12]. Privacy design guidelines in specific areas are given in [13,14], and [2] considers the case of cloud computing. In November 2007 the UK Information Commissioners Office (ICO) [15] (an organisation responsible for regulating and enforcing access to and use of personal information), launched a Privacy Impact Assessment (PIA) [15] process (incorporating privacy by design) to help organizations 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 [16]. This methodology aims to combat the slow take-up to design in privacy protections from first principles at the enterprise level, see [17] for further discussion, [18] for further background, and [19] for a useful classification system for online privacy.

However, whilst there is a body of privacy design guidelines, there exist no practical techniques for designing specifically for cloud environments. To do this we choose to focus on the utility of *design patterns* [20]. We believe that in some circumstances they could be useful since the use-cases that drive cloud computing are familiar ones and so design patterns to fit these can be produced [21]. Some previous work has been carried out in the privacy design pattern area, but not for cloud computing:

[22] describes four design patterns that can aid the decision making process for the designers of privacy protecting systems. These design patterns are applicable to the design of anonymity systems for various types of online communication, online data sharing, location monitoring, voting and electronic cash management and do not address use within an enterprise.

3 Cloud Capability Maturity Model

We begin by considering capability maturity for enterprises exploiting clouds, considering a number of key characteristics: business strategy, governance, procurement methods, applications, information and information security.

Table 1. Capability Maturity Model for Cloud Computing

	Level 1 Initial Services	Level 2 Architected Services	Level 3 Aligned and responsive	Level 4 Measurable	Level 5 Optimised and Dynamically Reconfigurable
Business Perspective	Componentised business units	Strategy and vision for broad adoption across business	Cloud services integrated across business	Impact of cloud on business lines measured	Multiple cloud suppliers used, seamless transition
Cloud Governance	Due diligence of external suppliers	Based on best practice	Cloud and enterprise practice / processes aligned	Compliance metrics established and applied	Dynamically monitored and enforced
Procurement Methods	Standard contractual arrangements utilised	Best practice in cloud procurement adopted, SLAs emerging	Support team established, common service environment, automated SLAs	Cloud procurement cycles measured for deployment efficiency and consumption	Dynamic and changing SLAs and optimisation
Cloud Applications	Additional functionality and siloed applications	Applications enable new enterprise activities	Process integration across enterprise, enhanced productivity	Cloud applications measured for impact	Bespoke and dynamically changing service offerings
Information Perspective	Cloud services offer information based applications to sites	Enterprise service meta-data available	Single enterprise ontology, shared with partners to enhance delivery	Cloud based information services measured for quality	Information services evolving in response to changing enterprise strategy
Infosec	Monitoring and control at gateway and enterprise network boundary	SLAs include infosec, idm across the enterprise and in cloud	Monitoring and auditability integrated across enterprise	Cloud infosec measured for impact on risks	Auto enforcement of changing cloud infosec strategies, multi-level security in place
	<i>single site</i>	<i>business wide on benefits</i>	<i>close alignment between enterprise and suppliers</i>	<i>enterprise deployment and aligned policies</i>	<i>leader in cloud capabilities</i>

We present our capability maturity model for enterprises exploiting cloud services in Table 1 above. Level 1 represents today’s environment where users of cloud services are adopting offerings to enable additional functionality, and controlling the risks via standard outsourcing risk management processes. The cloud service is consumed within a business unit and is typically siloed off from the rest of the enterprise. Information security is focused at the perimeter. At Level 2 best practice begins to emerge surrounding the adoption of cloud services within an enterprise, and the enterprise begins to roll out a broader adoption strategy. This in turn generates

enterprise level metadata which underpins new information based services. At Level 3 cloud and other business processes become aligned, enabling a more integrated management activity. This in turn delivers enhanced productivity. It also facilitates a single enterprise cloud ontology, which when shared with partners and suppliers can directly enhance delivery. The importance of cloud to the enterprise results in a dedicated support function being maintained within the enterprise. The information security function delivers monitoring and audibility across the enterprise. At Level 4 the impact of cloud on the enterprise becomes measurable, compliance metrics are established and services and applications are measured for quality. Information security functions for cloud are also measured for impact on the overall risk mitigation strategy. At Level 5 cloud services become dynamically reconfigurable in response to the increased awareness delivered by the various metrics and changing operating requirements. Governance mechanisms can be dynamically monitored and enforced. Procurement methods become dynamic, with SLAs requiring an agile and perhaps automated solution, in order to provide the agility required by the enterprise. Information security mechanisms also require additional automation and multi-level security solutions will need to be present and effective.

From an enterprise perspective it is at the points of crossing maturity levels that change is likely, for all characteristics. With change comes the potential for introduction of information security vulnerabilities, and alongside opportunities for designing in privacy. Consider the governance perspective: in the lower maturity levels best practice will be based upon existing outsourcing practice. However, as cloud exploitation matures this is unlikely to be sufficient since the dynamic business models and agility of service deployment will move at a faster pace. New best practice in risk management will certainly be required, and this will impact governance.

We can use this cloud exploitation capability maturity model to motivate a privacy maturity model for clouds which elucidates the enterprise architecture characteristics which will offer opportunities to deliver privacy preserving functionality, and will necessarily vary as cloud adoption matures.

4 Examples of Privacy Controls in Cloud Computing

It is possible to represent a privacy maturity model by capturing key privacy controls that have been identified in Table 1 above. These controls are shown in Table 2, and are loosely based upon the simpler model for privacy risks (in general) described in [24]. The controls are focused at an appropriate level to allay potential concerns relating to why personal information is collected, and how it will be used in the cloud or passed on to affiliated clouds at different maturity levels. However, the relative control level is selected according to the cloud maturity level. As an example, obligation management can evolve with increasing privacy risks, such that at preliminary stage contracts are used for legal compliance for data treatment, this would correspond to the initial services maturity level defined in Table 1. Before any transition to the use of 'architected services', obligations must be defined within the organisation and for third parties with whom information is shared, which may assure the user in a more mature or advanced level. However, obligations should be automated to facilitate

management processes and therefore enable the transition for cloud service providers to a level at which they can be dynamically composed and measured (e.g. level 4 and above). Finally, at the highest level the obligation management procedure is continually refined and improved with respect to the enterprise’s continuous exploitation of cloud services while responsibly protecting the individuals’ private information [25]. Hence, in order to transition across the relative maturity levels; the privacy controls are also required to transition; however, given the relationship is business- and regulatory- context dependant, the mapping cannot be guaranteed to be linear with businesses given freedom on their adoption of controls relative to their maturity in the use of cloud services. We hope to develop this mapping as part of future work.

In the next section, we demonstrate how guidance about designing such privacy controls into cloud services may be achieved by means of design patterns, and discuss a *sticky privacy policies* pattern in detail.

Table 2. Examples of Privacy Controls in Cloud Computing

Controls	Level 1	Level 2	Level 3	Level 4	Level 5
Usage of Privacy Policy	Privacy policy devised case by case	Privacy policy defined	Privacy policy and organization in place	Privacy requirements defined and managed	Continual improvement in policy definition
Privacy Policy Enforcement	Associated data with policy	Legal/regulation support	Systematic management of privacy policy associated with data	Sticky policy used	Policy enforcement along the data chain
Procurement Methods	Procedures applied to specific situation	General awareness, ad hoc procedures	Risk assessments/PIAs performed	Periodic risk-based reviews, monitoring on an organisational and functional level	Ongoing assessment and assurance; changes systematically scrutinised for privacy impact
Trust Management	Registration and authentication for certain purposes	Standard registration and authentication processes	Persistent compliance check and identity management	Built-in systematic accountability and reassurance processes	Adaptable architecture in place to handle regulations/business processes/standards changes
User Control Management	Standard data management procedure defined	Treatment of user data by organisation is made transparent to users	Users can define preferences over the treatment of their data, and organisation publish policies over data treatment	Usage of data reported/audited with regard to user preferences	Automatic policy negotiation between parties
Obligation Management	Standard contracts with ad hoc compliance to certain cases	Contracts used for legal compliance with data treatment	Obligations defined for organisation and for third parties with whom information is shared	Automated obligation management systems deployed	Refinement and improvement in obligation management

5 Designing Privacy into the Cloud via Design Patterns

The examples of privacy controls in cloud computing given above show that there will be multiple opportunities to *design in* data protection. In order to exploit these opportunities we require methods that can support an enterprise through its evolution towards maturity, which can incorporate anecdotal advice as well as more formal prescriptive solutions. Such methods also need to be flexible enough to incorporate solutions of varying types, including: processes; techniques; methods; training; software configurations; applications; communications protocols. We have selected *design patterns* [20] since they meet all of these requirements.

Key aspects of design patterns have already been introduced in Section 2. There are multiple approaches one might take to how solutions for different maturity levels of Table 1 are reflected into the corresponding design patterns. Our approach is to have the patterns correspond to giving further details of techniques in each ‘cell’ of Table 2, so that for each maturity level there would be a set of patterns. Where closely related techniques could be used across more than one level of Table 1, a single pattern may be used and the distinction between maturity levels made within the *context* and *solution* descriptions (see example below); in such a case a subjective judgement is needed, in the sense that if the variation is great then a new pattern would be created.

We describe below a draft design pattern for building a data protection mechanism into a cloud, specifically *Sticky Privacy Policies*. This pattern provides a method for addressing maturity within the enterprise use of privacy policy (identified by our capability maturity analysis outlined above); it corresponds to the control for Privacy Policy Enforcement used at level 4 in Table 2. Due to space limitations we concentrate on this example. We have also defined a number of others in a similar manner, including: obligation management, data fragmentation, user interface design techniques, risk assessment, reputation management, and user anonymisation.

Sticky Privacy Policy Example

Name: Sticky Privacy Policies **Classification:** Privacy Policy Enforcement

Intent: to bind a privacy policy to the data to which it refers

Motivation: The sticky privacy policy would ensure that policies relating to data are propagated and enforced along all supply chains in the cloud ecosystem and all mechanisms through which the data is stored, processed and shared.

Context: You are designing a cloud service solution and want to make sure that multiple parties are aware of and act in accordance with your policies as personal and sensitive data is passed along the chain of parties storing, using and sharing that data.

Problem: Data could be treated by receivers in ways that the data subject or initiator would not like. The policy could be ignored or separated from the data.

Forces: Factors related to Privacy Policy specification and maintenance and user control (in Table 2) are relevant, as well as contextual factors, notably user trust. For example, in situations where the user has low trust in the service providers, or they have unknown length and composition, the level of user control required increases, gradually implementing the solution set from this pattern.

Solution: Enforceable ‘sticky’ electronic privacy policies: personal information is associated with machine-readable policies, which are preferences or conditions about how that information 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) in such a way that this cannot be compromised. When information is processed,

this is done in such a way as to adhere to these constraints. These policies are associated with data using cryptographic mechanisms. At level 5, Identifier-Based Encryption (IBE) [26] is particularly appropriate as it means that a third party needs to check certain properties at the time of decryption, before a decryption key is released.

Design Issues

- To what level of granularity of data should the policy be attached? It could be anything from a personal data element (e.g. name, etc.) to a whole database
- It might be better to have a reference to a policy bound to the data rather than the actual policy bound to the data, for practicality reasons
- Need to be compatible with current/legacy systems
- Need to provide mechanism to enforce and audit between parties
- Need to provide mechanism for the parties to assess their enforcement abilities
- Need to provide economically feasible mechanism to enforce the policy.

Consequences: *Benefits:* Policies can be propagated throughout the cloud, strong enforcement of these policies, strong binding of data to policies, traceability (for the IBE approach [26]). Multiple copies of data each have the policy attached.

Liabilities: Scalability and practicality: if data is bonded with the policy, this makes data heavier and potentially not compatible to current information systems. It may be difficult to update the policy once the data is sent to the cloud, as there can be multiple copies of data and it might not be known where these are. Once the data is decrypted and in clear, the enforcement mechanism becomes weak, i.e. it is hard to enforce that the data cannot be shared further in clear, but must instead be passed on in the sticky policy form; therefore, audit must be used to check that this does not happen.

Known Uses: Policy specification, modelling and verification tools include EPAL [27], OASIS XACML [28], W3C P3P [29] and Ponder [30]. Various different solutions to sticky policies are compared in [31]. Notably, a technical solution for sticky policies and tracing services is introduced in [26] that leverages Identifier-Based Encryption (IBE) and trusted technologies. This solution requires enforcement for third party tracing and auditing parties. An alternative solution that relies on a Merkle hash tree has been proposed by Pöhls in [32]. A unified policy model is discussed in [33], which discusses steps towards privacy management for an organisation or across organisations within a federated identity management environment.

Related Patterns: obligations (obligations can be stuck to data), identity management (e.g. polices bound to data managed in identity management system), audit, Digital Rights Management (DRM).

Our conclusion is that a pattern approach is viable and scalable. However, patterns can only be as good as experience and analysis allow and so they will need to be evolved and refined. But they could offer a practical approach to enabling the adoption of best practice in discrete steps as an enterprise builds towards their optimum level of cloud exploitation maturity.

6 Maintaining Data Protection in the Cloud via SLAs

Whilst patterns provide an intuitive way to engage with system architects and policy developers during cloud service design, they may not be ideal for enabling contractual risk management maturity in deployed cloud services. Service level agreements (SLAs) are an industry standard approach for controlling risk, and so are a more natural starting point. For many outsourced services the SLA is a key document as it attempts to define the relationship between the customer and provider of a service, the service itself and the parameters which can be used to define performance of the service supplier [34]. In practice the SLA can have many more functions, dependant upon the service type and level (of Table 1) to which it is targeted. An enterprise may have a number of SLA agreements which can either be standalone or with multiple dependencies. For example, the hosting of a single server on which Human Resources (HR) data is stored may have a separate SLA from that of the database that it hosts and from that of the provision of the supported HR service.

In general, SLAs can be split into three functional areas [35], which may be simplified as: service description (including roles and responsibilities); SLA governance (identifying metrics and the process for dispute resolution); SLA change management (managing uncertainty and renegotiating services in the agreement). In order to be effective SLAs are broken down into specific service objectives with key performance indicators (KPIs) specifying the service delivered by the outsourced service. One of the key difficulties [36] is mapping the service KPI to a meaningful metric and ensuring that there is a shared understanding of that metric with the provider and the customer.

While recent work [37] has attempted to define KPIs for information security, to the best of our knowledge KPIs for information privacy have yet to be adequately tackled. The reason for this is probably twofold: firstly, as considered in Section 3, privacy is a broader topic and there are many different interpretations based upon societal, cultural and contextual factors; secondly, the privacy of an individual is interpreted through a number of data protection laws, which can potentially be contradictory, sector specific and vary between countries even when they interpret the same principles or directive (as is the case in the EEA).

At the lowest level (Level 1) of maturity cloud services (cf Table 1) are not too dissimilar to the current services we use today; standard contractual methods are used and SLAs will typically be natural language documents; thus high-level requirements will be directly translated into the SLA, where possible. For example, the UK Data Protection eighth principle states that personal data “should not be transferred outside the EEA unless an adequate level of data protection is ensured” and this can be directly expressed as a condition against which certain actions could constitute a serious breach of SLA.

The SP has the responsibility of designing and operating a system such that this breach is not the case and yet currently no standards or mechanisms exist for assessing the effectiveness or suitability of the design. By contrast as the cloud matures (cf Table 1) so a broader number of suppliers become interdependent and are used in a more dynamic manner. Increasing dynamics for quality of service and choice has already been recognised within other projects examining Grid [38] and SOA [39]. These projects have highlighted the need to be able to handle: SLAs that are composed of a hierarchy of technologies to comprise an overall service level objective

within an agreement; service provision that may need to change between providers either for functional, price or quality reasons; the SLA subsequently being negotiated, planned and deployed. A requirement of the SLA is that they are expressed in parameters that are tangible and can be processed by machines rather than requiring interpretation at human quarterly review meetings.

We believe that by exposing and interpreting these clauses as service level objectives within an SLA, we will provide engineering requirements against which we can map design solutions. For example, the ‘sticky policy’ design pattern can be used to ensure that the user can be assured that the data processor has correct instructions for each individual data item as to where it may be transferred and processed (e.g. outside of the EEA/Safe Harbour etc) and gain assurance in so doing. However, challenges still remain in strongly defining both the design pattern properties (such as the ability to strongly identify processing parties, via IBE) and the detail required by the service level objective to satisfy the service user (such as identifying all data processors). Clearly, both objective and data protection properties must be expressed in a solution neutral format as well as having a common ontology that can be encoded within a machine readable format around which communication can take place. Due to space limitations we are unable to present any further examples. It may be possible to define a pattern which communicates how an enterprise should seek to utilise SLAs as it matures; however, at this time the technical capability does not exist to actually support the higher levels of maturity envisaged.

7 Conclusion and Acknowledgements

We have demonstrated that through the creation of novel capability maturity models for cloud exploitation and associated privacy requirements in an enterprise we can begin to identify opportunities to produce methods for designing data protection into clouds. Design patterns provide a good mechanism for expressing such techniques in a manner which could be useful at all levels of maturity. We believe that such patterns could also be applicable for architecting SLAs. Our future work will include a more complete analysis of where the maturity models indicate opportunities for designing in privacy, including an assessment of those which should be considered high priority. It is clear that privacy controls (cf. Table 2) could be highly context-dependent. To avoid overlooking the subtleties of individual privacy concerns, which may vary with context, we must be careful not to use too general a template. Further research is required to elucidate the user related contextual factors affecting the degree of privacy protection that is necessary for a given context. Such factors include: sensitivity of data, location of data, sector, contractual restrictions, cultural expectations, user trust (in organisations, etc.), trustworthiness of partners, security deployed in the infrastructure, etc. We will explore the use of recommendations which could be deduced via a decision based support system that assesses context, and then outputs a list of recommendations and controls, as has been done with [23]. Further analysis is also required to study the utility of the design pattern approach by application to a case study, and the issues surrounding legal protections and their inclusion in SLA design patterns.

This research is being conducted as part of the EnCoRe project [40], focused on delivering usable consent and revocation controls for managing personal data and considering amongst others cloud service operating environments.

References

1. Gartner: Forecast: Sizing the Cloud; Understanding the Opportunities in Cloud Services (March 2009)
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 (2006), http://papers.ssrn.com/sol3/papers.cfm?abstract_id=667622
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. Wikipedia (2009), http://en.wikipedia.org/wiki/Capability_Maturity_Model
6. Smith, R.: Cloud Maturity Models Don't Make Sense (2008), http://www.informationweek.com/blog/main/archives/2008/12/cloud_maturity.html;jsessionid=OL1NSZLUOGDMCQSNLDPCKHSCJUNN2JVN
7. Urquhart, J.: A maturity model for cloud computing (2009), http://news.cnet.com/8301-19413_3-10122295-240.html
8. Sorofman, J.: The cloud computing adoption model (2009), <http://www.ddj.com/architect/211201818>
9. Wardley, S.: Maturity models for the cloud (2009), <http://blog.gardeviance.org/2008/12/maturity-models-for-cloud.html>
10. OpenGroup: A Maturity Model for SOA (2009), <http://www.opengroup.org/projects/soa-book/page.tpl?CALLER=faq.tpl&ggid=1319>
11. Microsoft Corporation: Privacy Guidelines for Developing Software Products and Services, Version 2.1a (2007), <http://www.microsoft.com/Downloads/details.aspx?FamilyID=c48cf80f-6e87-48f5-83ec-a18d1ad2fc1f&displaylang=en>
12. Cannon, J.C.: Privacy: What Developers and IT Professionals Should Know. Addison-Wesley, Reading (2004)
13. Patrick, A., Kenny, S.: From Privacy Legislation to Interface Design: Implementing Information Privacy in Human-Computer Interactions. In: Dingledine, R. (ed.) PET 2003. LNCS, vol. 2760, pp. 107-124. Springer, Heidelberg (2003)
14. Bellotti, V., Sellen, A.: Design for Privacy in Ubiquitous Computing Environments. In: Proc. 3rd conference on European Conference on Computer-Supported Cooperative Work, pp. 77-92 (1993)
15. Information Commissioner's Office: PIA handbook (2007), <http://www.ico.gov.uk/>
16. Office of the Privacy Commissioner of Canada: Fact sheet: Privacy impact assessments (2007), <http://www.privcom.gc.ca/>
17. Information Commissioners Office: Privacy by Design. Report (2008), <http://www.ico.gov.uk>
18. Jutla, D.N., Bodorik, P.: Sociotechnical architecture for online privacy. *IEEE Security and Privacy* 3(2), 29-39 (2005)
19. Spiekermann, S., Cranor, L.F.: Engineering privacy. *IEEE Transactions on Software Engineering*, 1-42 (2008)
20. Alexander, C., Ishikawa, S., Silverstein, M., Jacobson, M., Fiksdahl-King, I., Angel, S.: *A Pattern Language: Towns, Buildings, Construction*. Oxford University Press, Oxford (1977)

21. Arista, Cloud Networking: Design Patterns for 'Cloud Centric' Application Environments (2009),
<http://www.aristanetworks.com/en/CloudCentricDesignPatterns.pdf>
22. Hafiz, M.: A collection of privacy design patterns. *Pattern Languages of Programs*, 1–13 (2006)
23. Pearson, S., Sander, T., Sharma, R.: A Privacy Management Tool for Global Outsourcing. In: *DPM 2009*. LNCS, vol. 5939. Springer, Heidelberg (2009)
24. The Institute of Internal Auditors: Managing and Auditing Privacy Risks,
<http://www.theiia.org/download.cfm?file=33917>
25. 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)
26. 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)
27. IBM, The Enterprise Privacy Authorization Language (EPAL), EPAL specification, v1.2 (2004),
<http://www.zurich.ibm.com/security/enterprise-privacy/epal/>
28. OASIS, eXtensible Access Control Markup Language (XACML),
http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml
29. Cranor, L.: *Web Privacy with P3P*. O'Reilly & Associates, Sebastopol (2002)
30. Damianou, N., Dulay, N., Lupu, E., Sloman, M.: The Ponder Policy Specification Language (2001),
<http://wwwdse.doc.ic.ac.uk/research/policies/index.shtml>
31. Tang, Q.: On Using Encryption Techniques to Enhance Sticky Policies Enforcement, Technical Report TR-CTIT-08-64, Centre for Telematics and Information Technology, University of Twente, Enschede (2008)
32. Pöhls, H.C.: Verifiable and Revocable Expression of Consent to Processing of Aggregated Personal Data. In: Chen, L., Ryan, M.D., Wang, G. (eds.) *ICICS 2008*. LNCS, vol. 5308, pp. 279–293. Springer, Heidelberg (2008)
33. Schunter, M., Waidner, M.: Simplified privacy controls for aggregated services - suspend and resume of personal data. In: Borisov, N., Golle, P. (eds.) *PET 2007*. LNCS, vol. 4776, pp. 218–232. Springer, Heidelberg (2007)
34. Clarke, I., Miller, S.G.: Protecting Free Expression Online with Freenet. *IEEE Computing* (2002)
35. Rhea, S., Eaton, P., Geels, D., Weatherspoon, H., Zhao, B., Kubiawicz, J.: Pond: the OceanStore Prototype. In: *FAST 2003* (2003)
36. Huang, C.D., Goo, J.: Rescuing IT Outsourcing- Strategic Use of Service Level Agreements, *IT Pro.* (2009)
37. Yearworth, M., Monahan, B., Pym, D.: Predictive Modelling for Security Operations Economics, HPL-2006-125 (2006)
38. EU FP7 Network of Excellence(2009) ,
<http://www.coregrid.net/>
39. EU FP7 Project SLA Aware Infrastructure (2009),
<http://sla-at-soi.eu/>
40. EnCoRe: Ensuring Consent and Revocation project (2008),
<http://www.encore-project.info>