Decision Support for Selection of Cloud Service Providers

Tomas Sander

HP Labs
5 Vaughn Drive
Princeton, US
+1-609-514-0680
Tomas.Sander@hp.com

Siani Pearson

HP Labs
Long Down Avenue
Bristol, UK
+44-117-3128438
Siani.Pearson@hp.com

Abstract—Clear and consistent assessment of the various capabilities of cloud service providers (CSPs) will become an essential factor in deciding on which CSPs to use in the future, particularly as cloud service provision expands further into more sensitive and regulated areas. This paper describes an approach that is useful in this regard. Specifically, we describe a mechanism in which context is gathered relating to CSPs; this is inputted to a rule-based system and decisions are output about the suitability of each CSP, including an analysis of privacy and security risk and recommended stipulations to be taken into account when negotiating contracts and SLAs.

Keywords—Cloud service provider, provisioning, decision support system

1. INTRODUCTION

There are a range of privacy and security concerns related to cloud computing. In particular, putting data in the cloud may impact privacy rights, obligations and status, and different laws may apply depending upon where information resides in the cloud. However, there is strong evidence that current cloud models do not take into account confidence, trust and assurance requirements [1]. Although currently this is often overlooked, when considering using a CSP, organizations should ask the same questions that they would during a third party vendor or business partner security program review, as well as some additional cloud-specific questions.

For example, privacy-related questions that should be asked in advance by an organization to CSPs if personal information is going to be handled by them might include:

- Who will have access to the organisation's data?
- Where will this data be processed and stored?
- Will it be encrypted at rest and in transit?
- Will it be intermingled with data from other companies?
- Are the backup and recovery processes in place adequate for the organisation's needs?
- What are the availability promises for the cloud service? Are they documented within a Service Level Agreement (SLA)?

- Will the organisation be able to obtain information about data access and associated logs, and how quickly will this be?
- What audit trails are generated and maintained for the data? Does this include logs about data usage and sharing?

In addition, legal issues such as transborder data flow may need to be taken into account. The situation can easily become too complex for a human alone to handle, or it might be that the human does not have the requisite experience to be able to make an informed decision. The latter is particularly relevant for companies who do not yet have extensive experience with outsourcing.

In this paper we describe a solution to this issue, in the form of a decision support system for cloud provisioning. In the following section we provide more details of this approach. Section 3 describes how the knowledge base (KB) can be built up and used within the system. Section 4 provides an example usage case, Section 5 compares this approach with related work and Section 6 assesses its current status. Finally, conclusions are given.

2. DECISION SUPPORT FOR CLOUD COMPUTING

In this section we provide an overview of our approach, details of the system architecture, examine the system from the customer and the CSP viewpoints, and discuss the key issue of how to deal with the CSP responses in a uniform way.

2.1 Overview of our Approach

Our approach is a decision support tool that gathers context relating to cloud service providers (CSPs) and inputs to a rule-based system, to trigger decisions about whether or not to use that CSP and/or additional stipulations that would need to be made. The tool helps to determine appropriate actions that should be allowed and assesses risk before personal information is passed on through the cloud. It is semi-automated to significantly lower the transaction costs for the selection of CSPs for the customer.

Our solution provides a means in which an assessment of the CSP's practices can be made in a semi-automated fashion prior to the decision to use that CSP. The solution may also be useful for the CSPs themselves when composing other CSP's services to deliver their own services.

Our innovation is the usage of a specialised tool in order to aid a customer to make decisions about whether or not to use a particular CSP and optionally also particular clauses that should be added into contracts or service level agreements (SLAs) between the customer and that CSP. When the customer wishes to assess different CSPs offering a service, these providers will use the tool via a web interface in order to provide answers to the questionnaire which they are asked, and the results will be sent back to the enterprise that wishes to choose between the service providers. These results will include reports and automatically generated ratings. Potentially, a compliance team of the customer firm, regulators or other third parties (e.g. auditors) can access a view on the DSS system (focusing on its logs and generated reports). This is a major step towards providing transparency and accountability for assessing compliance.

Our tool offers a number of advantages over currently available solutions. In comparison to manual or Excel spreadsheet-based solutions, our tool has intelligence integrated into the solution so that it can efficiently tailor questionnaires and output to particular customer needs. Decision tree-based solutions are also not powerful enough to model common privacy knowledge conveniently as decision trees have difficulty in dealing simultaneously with the multiple dependencies relevant for privacy decisions (e.g. between the countries involved, the type of data that need to be handled, business processes etc.) On the other hand, general expert systems due to their extensive capabilities lack determinism and sometimes guarantees for termination. The tool we are using is expressive enough to express common privacy knowledge for cloud services, while restricted enough to have provable determinism and termination. We achieve this by using a rule-based system in a controlled way. The basic tool has been built and its KB has been successfully populated with the privacy knowledge of the 300 page rule-book of a global corporation.

2.2 System Architecture

There are various ways in which such a decision support tool might be deployed in cloud environments. One example that provides trust for the customer would be if the customer hosts the Cloud Decision Support System (DSS), and optionally even by the DSS being provided as a service within a private, hybrid or public cloud that is controlled by the customer. More

generally, a cloud-based CSP assessment service could be offered to a number of enterprises, as shown in Figure 1. For each customer enterprise, an administrator will set up the original questionnaire according to the policies that the customer (i.e. the enterprise) wishes to check, or else the enterprise may just use a default setting offered by the assessment service. When the customer (typically, an enterprise) wishes to assess different CSPs offering a service, these providers (CSP1 and CSP2 in Figure 1) will use the tool via a web interface in order to provide answers to the questionnaire which they are asked, and the results will be sent back to the enterprise that wishes to choose between the service providers. These results will include reports and automatically generated ratings (to allow the administrator to easily distinguish between the competitors).

They are generated by means of an inference engine that contains rules which are triggered by parameters (set from the answers to the questions or set by other rules) and that output new parameter settings, refine metrics that are used for the ratings and/or information to be contained within the report.

Fig. 1. High-level architecture.

The particular inputs, settings, reports and values associated with each assessment are stored within the database (DB). An authorised entity may review and also update the knowledge base (KB) if desired, via a management interface and management handler.

The mechanism by which inputs from the service provider side are compared with the requirements of the customer may vary. For example, the CSPs may publish policies and the DSS can use this input within the tool, without the need to engage the CSPs in any interactions with the DSS. Alternatively, as shown in Figure 1 for example, the CSPs might engage with the DSS directly in order that the DSS can produce output from this interaction that is usable by customers. Figure

1 illustrates the case where a CSP might host the DSS as part of a cloud service available for multiple customers; alternatively, as mentioned above, an individual customer might host the DSS.

2.3 The Customer Experience

The customer's task is to provide the context that determines the requirements for the cloud service provision. The customer user interface (UI) asks the customer for information that allows generation of a questionnaire that will be shown to the CSP and provides the context in which the CSP's answers can be rated for their adequacy to the customer's needs.

The customer UI assesses

- the type of data that will be processed (e.g. whether there is any sensitive data involved),
- for which business process the outsourcing is to occur
- the country of origin of the data subjects
- any restrictions the customer might have for countries where the data will be processed or stored
- which level of access to the data are required
- certain security standards and security certifications and audit procedures the customer wishes to see
- whether end users need to be in direct contact with the service provider and who contacts whom over which medium.

Some of these questions – such as for data types, the business process used etc – are made mandatory, whereas others such as certifications and security standards are optional (in cases where the customer does not yet have strict assurance requirements in place). We also allow the customer company to mark certain answers as 'always applicable' for that corporation, for instance because they reflect company policies. Thus, when the customer works on different outsourcing projects, the company profile is always pre-populated with these answers. This is beneficial by ensuring that certain minimum standards are followed with all cloud vendors with whom a customer might be dealing.

In addition, the customer UI allows for certain free text entries about the customer expectations to express requirements that can not yet be expressed with the pre-populated choices. This information will be displayed to the CSP who can respond with free text. However this information cannot yet be used for the automated reasoning in the tool and should thus only be used where necessary.

2.4 The Cloud Service Provider Experience

As mentioned in Section 2.2, there are various alternative mechanisms for providing the information about what the CSP will do: using the DSS directly with the customer, or else using the tool to produce output that is then accessible publicly, e.g. via a website which contains the CSP's answers to common question the DSS would ask. These CSP policies can then be taken as input by the DSS in order to make an evaluation for the current organisation's context. The motivation for CSPs to provide input for the tool is market forces – otherwise, only the 'lowest' default policy within the system would be allocated to their offering.

We focus on the solution where a DSS is used. In this case a questionnaire will be presented to an administrator acting on behalf of the CSP asking them in clear natural language (with drop-down choices where possible rather than free text options) information that is relevant to be able to assess the security and trustworthiness of their service. In addition relevant customer requirements are explicitly displayed to the CSP admin so he can clearly understand what the minimal requirements of the customer are. Based on this information the CSP will be asked questions such as the location of the servers to be used, access control mechanisms used, encryption mechanisms used, purposes for which data will be used, whether there is further outsourcing/sharing of information with other organizations, and the mechanisms in place that the CSP has for enforcing the organisation's requirements along the chain, backup provisions, etc.

In addition the CSP can via the UIs specify more than one service, i.e. different degrees of assurance and indicate that there will be additional charges related to them. The tool does not attempt to capture exact prices. These pricing details can be revealed later in the negotiations with the customer. However the tool allows the customer to clearly see what the different security and privacy choices are that a particular vendor makes available for their situation.

2.5 Enforcing Uniformity

In order for the decision support system (DSS) to be truly useful for comparing different CSPs it needs to enforce uniformity and consistency in the information provided by the various CSPs. Given that the customer defines the question set at least partially, there is an opportunity for the CSP completing the questionnaire to enter in data which may be subject to misinterpretation. The DSS deals with this by asking very detailed questions to lead to a detailed and less ambiguous assessment. Instead of asking for example if the CSP follows good security or privacy practices for a given domain the tool will ask a number of specific questions. For example if the CSP is assisting the customer in conducting an email

marketing campaign, a number of pointed questions will be asked, such as

- “Will the subject line of the email be non-misleading?”
- “Will each email have an unsubscribe link?”
- “Will unsubscribe requests be honored within 10 working days (5 days within Australia or New Zealand)?”

etc rather than relying exclusively on general (and thereby potentially ambiguous) questions such as

- “Do you follow all the applicable privacy legislations w.r.t. email marketing?”

Another mechanism we have implemented in the tool is rule-driven help that can be associated with each question. The help information can thus be made specific to the context the customer has specified and contain specific examples and guidance how questions should be answered.

Being able to drill down to such a level of specificity is much easier doable with an automated tool than with paper documents. Where it seems useful standard cloud evaluation criteria as we will describe in Section 3.1 can be refined through further drill-down questions inquiring into details.

3. KNOWLEDGE REPRESENTATION AND INFERENCE

In this section we provide more detail about the knowledge base, and the representation and inference mechanisms used.

3.1 Cloud Evaluation Criteria

The system described in Section 2 generates information to support choices of cloud suppliers. To make reliable data collection possible within this process, it is highly preferable to use a clear set of evaluation criteria that are agreed across industry. There are currently several candidates for this: notably the ENISA standards [2] that are currently being refined and standardized, and alternatively the Shared Assessments tools [3]. It may also be possible to integrate the approach of Jericho Forum [4] when gathering the context: three scales are defined by which to measure types of clouds: whether it is internal or external; whether it is proprietary or open; and whether it has a security perimeter or not; these three scales can be arranged as three axes, creating a ‘cloud cube’ in which clouds can be placed and classified.

3.2 Rule Representation

The tool uses a set of intermediate variables (IMs) to encode meaningful information and to

drive the questionnaire generation. IMs can be thought of as flags, for example ‘transborder data flow’ (which indicates that the current context involves transborder data flow). For provability, we mandate monotonicity for the IMs, i.e. they cannot be retracted once they have been asserted.

We define two kinds of rules: question (i.e. questionnaire generation) rules, and output rules. All the rules have the general form

when condition then action

Question rules have as their conditions a monotonic expression (i.e. Boolean expression built up using & and v as logical operators) in intermediate variables (IMs) and/or (question, answer) pairs and as actions, directives to ask the user some questions or else to set some IMs. The output rules’ condition is a Boolean expression in a set of IMs and answers to questions and they generate as their actions the content of the output report.

Some of the output rules encode the customer’s policies; others encode regulatory privacy requirements. They represent trigger conditions based on Boolean combinations of parameters corresponding to properties relating to cloud service provision and output in the form of risk levels and other information,

All the rules in the system are based upon the production rule system Drools [5]. Let us consider a simple example of the underlying representation, in DRL format (although the rules can automatically be converted to XML format). Assume that a CSP is answering a questionnaire, and that the question “*Will data be stored in encrypted form?*” is answered “Yes”. Assume this question has identifier number 48 in the system. This (question, answer) pair is added to working memory and as a consequence the following question rule is triggered asserting a new IM “Encrypted storage”:

```
rule “IMR21” when QA (id == 48, value
== “Yes”) then insert(new
IM(“Encrypted storage”, “Yes”)); end
```

When the previous IM is asserted to the working memory it triggers the following question rule which adds new questions to the questionnaire:

```
rule “QR17”
when IM (name == “Enc”, value == “Yes”)
then AddToDisplayList_DF(current,
currentQuestion, new long[] {49});
end
```

Question 49 is “What encryption mechanism is being used?”. In other cases, blocks of questions are added to the questionnaire just by adding them to the set in the rule above. The questionnaire generation procedure then iterates through the

questionnaire. The initial (question, answer) pair will also generate a new parameter instance: “Encryption used in storage” with value “Yes”. If we require this to trigger an output, then this can be captured within a output rule. For the example above, when this parameter instance is added to the working memory of the privacy engine it triggers the following output rule:

```
rule “Encrypted storage”
when ParameterInstance ( name ==
“Encryption Used in storage” , value ==
“Yes” )
then
report.addRule(new
RuleFacade().findById(50)); end
```

This rule adds a *Rule* object to the list of rules of the report. The rule can show a flag (to indicate the seriousness of the issue), a reason, a link to more information or other items to be included within the report to the customer e.g.

1. obligations that need to be passed on to other service providers;
2. other information that needs to be provided to the company (potentially including evidence);
3. restrictions about what that service provider should do;
4. clauses that should be added into contracts or service level agreements (SLAs).

3.2.1 Example

An example of a question that can be asked from the Shared Assessments guidelines [3] is: “Are the organisation’s employees and its Third Parties instructed to immediately notify the appropriate individual in the organization if or when target privacy data is, has been or is reasonably likely to have been lost, accessed by, used by or disclosed to unauthorized Third Parties?” This may be represented within our system in the same manner as described for the encryption question (question 48) above; similarly for the other questions in [3]. We can define question rules that link these questions, so that only those questions are asked that are relevant to the given context, and we can define output rules that link the answers to these questions to an expert’s assessment of the risk related to that answer, and more generally to a metric that contributes to the overall assessment of trustworthiness of the CSP that gives that answer (i.e. has that property).

3.3 Inference

Our approach uses an inference engine to run rules in order to produce output. This approach is not reliant on any particular inference engine

or specific format beyond processing of ‘if.. then’ rules, and so a variety of mechanisms could be used here, from production rule systems to Prolog. For our prototyping we have used the JBoss Drools rules engine [5]; this is run after each question is answered by the user and computes additional questions a user will be asked as well as the output reports provided by the system.

4. USE CASE

In this section we describe the use case of a fictitious UK- based company “PeaceOfMind.com” (PoM) that sells meditation supplies (cushions, gongs, incense etc.) via an online store. PoM has recently been undergoing rapid growth, making it necessary to become more professional in the management of its vendors and suppliers, customer database and marketing efforts. PoM would ideally like not to host any applications for these services itself and wishes to evaluate cloud service providers as a cost-effective alternative. PoM does not have particular technical or outsourcing expertise in house and wishes to use the CSP Assessment tool described in this paper to decide whether it might use public cloud services and still be compliant. PoM, due to the nature of its business, also wishes to hold itself accountable to the highest privacy and ethical standards for the management of its customers’ data.

PoM uses the tool to evaluate cloud services for managing its customer data base and for email and telemarketing.

The tool first provides an interface to PoM where an authorized employee of PoM acting on behalf of PoM (henceforth this individual – or multiple individuals sharing this task – is just referred to as ‘PoM’) enters the types of information that would be included in such activities, which in this case would be name, address, email address, home and cell phone, as well as contact preferences for the individuals. PoM answers the next question about the countries of origin of the data subjects by clicking that it has data customers from several European countries, including UK, Spain, Germany and Denmark as it has subsidiaries in all these countries that have collected customer data.

When asked about the business function being outsourced, POM selects customer relationship management and email marketing.

When asked about security and privacy requirements, PoM selects “High”. PoM also specifies that it does not wish the CSP to make any secondary use of the customer data (e.g. data mining), that it requires backup services and disaster recovery mechanisms for its customer data, and that PoM wishes to be notified in case there are any privacy or security breaches.

The tool takes this input and translates it into an online questionnaire to be filled out by the prospective CSPs.

The tool presents each CSP with a description of the service the customer is looking for and that asks the CSP a

number of questions: for example which countries the CSP uses for processing, storage and backup for data and whether the CSP subcontracts some of the relevant operations.

The high security and privacy requirements requested by PoM are translated into a variety of questions including which security certifications the vendor holds, whether it deploys encryption for data in transit and at rest, what type of identity and access management is used etc. In addition, questions about incident management procedures and secondary uses of data will be explicitly asked. For the email marketing aspect of this project, a number of detailed questions are asked to confirm that the CSP is in compliance with applicable regulations and generally applies best practice.

In our example, two CSPs – CSP 1 and CSP 2 – fill out the questionnaire. Both CSPs confirm that the data processing locations they use are within EU countries. The ‘transborder data flow’ risk indicator is therefore green, i.e. there are no significant risks in this domain. However, CSP2 – unlike CSP1 – uses subcontractors for data backup services. This creates a yellow flag within the output report for CSP2, pointing out that the acceptability of the practices of the subcontractor need to be verified by PoM. A respective follow up item is created as part of a checklist. The security indicator for CSP1 is green as it turns out that CSP1 has a number of relevant security certifications and gives satisfactory answers to a number of security-related questions. CSP2’s security risk indicator is yellow due to the fact that the answers suggest possible weaknesses in its security procedures.

After reviewing these and other risk indicators, PoM decides that CSP1 fulfills its requirements much better than CSP2. Although CSP1 turns out to be more expensive than CSP2, PoM decides that CSP1 will be the best provider to which to entrust its customer data.

5. RELATED WORK

A “Standardized Information Gathering” vendor security assessment questionnaire has been developed by BITS [3]. There is another vendor security review tool available at [6]. Relevant standards have already been discussed in Section 3.1.

There has been some relevant related work on scanning cloud solutions networks, operating systems, and web applications and performing automated penetration testing [7]. This work may test for cloud performance or availability. These automated tests could complement and be integrated into the type of intelligent decision support we are suggesting.

There has also been work in meta-scheduling for Grid applications: in particular, heuristics for scheduling parallel applications on Utility Grids that manage and optimize the trade-off between time and cost constraints [8]; a meta-scheduler that maps user applications to suitable distributed resources using a Continuous Double Auction, using a valuation metric for a user’s applications and computational resources based on multi-criteria requirements of users and resource load [9]. However these systems are not suitable to

address the compliance and policy issues we are concerned about in a comprehensive fashion.

In our system (as with expert systems [10]), problem expertise is encoded in the data structures rather than the programs and the inference rules are authored by a domain expert. Most of those systems are far too complex to allow for provable completeness or predictability. By restricting the structure of conditions and actions in the rules we can prove completeness and predictability.

There has also been some work on dynamic question generation in the expert system community [11] but their concerns and methods are very different from ours.

There has been extensive work carried out to define different types of security and privacy policy: such policy specification, modeling and verification tools include EPAL [12], OASIS XACML [13], W3C P3P [14] and Ponder [15]. These policies are formulated at a different level than the ones we are dealing with, as for example they deal with operational policies, access control constraints, etc. and cannot comfortably express context-specific privacy requirements for business processes such as marketing or software development. In circumstances where policies are specified by one party and enforced by another, the meaning must be agreed: this is achievable for example via standardization or via ontologies. There has been prior work to allow automated checking of user side policies with service side policies, e.g. via P3P [14] and Privacy and Identity Management in Europe (PRIME) project [16], but this is not directly translatable to the problem above.

Most importantly, these policies are not human understandable, which is what we need, even though they can be machine executed. IBM and Sun have done some research on privacy policy management such as EPAL [12] and XACML [13] which are low level privacy policy languages and not well suited for human user understanding. In the Sparcle project, see e.g. [17], IBM Research built an editor to support transforming natural based policies into XML. This makes it easier for non-experts to input rules into the system, but the output format itself is not user friendly. The REALM project [18] from IBM Research suffers the same shortcomings. OASIS LegalXML [19] has worked on creation and management of contract documents and terms, but this converts legal documents into an XML format that is too verbose. Breaux and Antón [20] have also carried out some work on how to extract privacy rules and regulations from natural language text. Their work has a different focus, but could be complementary for helping to populate the KB more easily.

Translation of legislation/regulation to machine readable policies has proven very difficult, although there are some examples of how translations of principles into machine readable policies can be done: Privacy Incorporated Software Agent (PISA) project [21] (where privacy principles derived from [22] were modelled and used as a backbone in conversations between agents [23], P3P [14] (where user privacy preferences were matched against web site privacy statements) and PRIME [16] (involving the definition and usage of various types of user and service side privacy

policies). However our research did not have to address the problem of interpreting and modelling arbitrary laws and adding them into our KB.

Perhaps the most relevant prior work is a policy framework for global enforcement of data assurance controls that enables the expression of both service providers' capabilities and customers' requirements, and enforcement of the agreed-upon requirements in service providers' environments [24]. This describes how the data owner can specify policies in terms of state machines that are placed in the metadata. These state machines are interpreted/executed by the service provider and allow specific policies provided by the data owner to be checked when dealing with the data. The data owner provides policies to the service provider along with the data, but the service provider is responsible for enforcing them, deciding how the high level requirements map down to technical controls and the data owner is not given assurance that the enforcement has taken place in a correct manner.

The DSS does not actually carry out checks that the CSP policies are indeed accurate – it treats the assertions made at face value. To enhance our approach, this work could be integrated with other research that can provide a greater degree of assurance of such policies (for example, [16]), and with audit mechanisms and reputation management technologies. A method of automated trust negotiation within virtual computing environments has been proposed: this establishes a trust relationship between two strangers by exchanging their access control policies and credentials [25]. Our approach takes a much broader range of factors into consideration when calculating the trustworthiness of a CSP.

6. CURRENT STATUS

Our solution is based on decision support technology the authors developed that can provide privacy assessments in complex intra-company scenarios [26]. This 'HP Privacy Advisor' is a rule-driven system that is currently being rolled out to employees for privacy assessments and accountability within HP.

We are applying this system to the cloud environment. This includes generating different UIs and a KB for this new system based upon industry-wide recommendations as to properties of CSPs that should be checked (see Subsection 3.1). From a knowledge representation perspective the cloud knowledge we wish to model in the KB does not appear more complex than the worldwide privacy knowledge we already have been able to successfully express. We are thus confident that our knowledge management capabilities can be used to express the cloud knowledge as well.

6.1 Next steps

Open issues that we will explore next include:

- Ongoing assessment of the cloud services that an organization is using. The DSS would not just provide an assessment at the beginning of the service, but would provide its output reports, warnings, etc. on an ongoing basis, quickly adapting when new information

comes in, such as that certain servers have become unreliable or insecure.

- Capturing additional relevant risks in the cloud space.
- Modelling dependencies among different CSPs; for example, if one provider goes down, which impact does this have?
- By modifying the ruleset one can run through 'what if' scenarios about policy changes and evaluate their effects.

7. CONCLUSION

Cloud computing has the potential to reduce protections for personal data, and to compromise data security. This paper describes a system that aims to aid organizations in ensuring that their – and their customers' – information will be treated appropriately before committing it to cloud service provision.

ACKNOWLEDGMENT

Our thanks to the extended team involved in specifying, developing and testing the HP Privacy Advisor system.

REFERENCES

- [1] D. Blum, Cloud Computing Security in the Enterprise, Research Overview, Burton Group, July 2009
- [2] ENISA, Cloud Computing: Benefits, risks and recommendations for information security, Ed. Daniele Catteddu and Giles Hogben, November 2009
- [3] BITS., Shared Assessments, 2010, <http://www.sharedassessments.org/>
- [4] Jericho Forum, "Cloud Cube Model: Selecting Cloud Formations for Secure Collaboration", v1.0, April 2009. www.jerichoforum.org
- [5] Drools, 2010. <http://jboss.org/drools/>.
- [6] R. Herold, Vendor Security Review tool, 2010. <http://www.privacyguidance.com>
- [7] HP Cloud Assure Solution
- [8] S.K. Garg, R. Buyya and H.J. Siegel, Time and cost trade-off management for scheduling parallel applications on Utility Grids, *Future Generation Computer Systems*, Elsevier, 2009.
- [9] S.K. Garg, S. Venugopal and R. Buyya, A Meta-scheduler with Auction Based Resource Allocation for Global Grids, *Proc. 14th IEEE International Conference on Parallel and Distributed Systems*, IEEE, 2008.
- [10] S. Russel and P. Norvig, Artificial Intelligence – A Modern Approach, 2nd edition, Prentice Hall, Englewood Cliffs, 2003.
- [11] J. McGough, J. Mortensen, J. Johnson and S. Fadali, A web-based testing system with dynamic question generation, *Proc. Frontiers in Education Conference*, Reno, NV, IEEE, 2001.
- [12] IBM, The Enterprise Privacy Authorization Language (EPAL), EPAL specification, v1.2, <http://www.zurich.ibm.com/security/enterprise-privacy/epal/>, 2004.
- [13] OASIS, eXtensible Access Control Markup Language (XACML), 2009. http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml.
- [14] L. Cranor, Web Privacy with P3P, O'Reilly & Associates, 2002.
- [15] N. Damianou, N. Dulay, E. Lupu, M. Sloman, The Ponder Policy Specification Language, <http://www-dse.doc.ic.ac.uk/research/policies/index.shtml>, 2001.
- [16] PRIME, Privacy and Identity Management for Europe, <http://www.prime-project.org.eu>, 2008.

- [17] K. Vaniea, C. Karat, J.B. Gross, J. Karat and C. Brodie, Evaluating assistance of natural language policy authoring, *Proc. SOUPS '08*, vol. 337, 2008.
- [18] IBM, REALM project, <http://www.zurich.ibm.com/security/publications/2006/REALM-at-IRIS2006-20060217.pdf>
- [19] OASIS, eContracts Specification v1.0, www.oasis-open.org/apps/org/workgroup/legalxml-econtracts, 2007.
- [20] D. Travis, T. Breaux and A. Antón, Analyzing Regulatory Rules for Privacy and Security Requirements. *IEEE Transactions on Software Engineering*, 34(1), pp. 5-20, 2008.
- [21] S. Kenny and J. Borking, The Value of Privacy Engineering, *Journal of Information, Law and Technology (JILT)*, 1. <http://elj.warwick.ac.uk/jilt/02-/kenny.html>, 2002.
- [22] Organization for Economic Co-operation and Development (OECD), Guidelines Governing the Protection of Privacy and Transborder Flow of Personal Data, OECD, Geneva, 1980.
- [23] J. Borking, Privacy Rules: A Steeple Chase for Systems Architects, www.w3.org/2006/07/privacy-ws/papers/04-borking-rules/, 2007.
- [24] J. Li, B. Stephenson, H.R. Motahari-Nezhad and S. Singhal, "A Data Assurance Policy Specification and Enforcement Framework for Outsourced Services", HP Labs Research Report, HPL-2009-357, 2009.
- [25] D. Zou, S. Du, W. Zheng and H. Jin, Building Automated Trust Negotiation architecture in virtual computing environment. *Springer J Supercomput.*, Nov. 2009.
- [26] S. Pearson, P. Rao, T. Sander, A. Parry, A. Paull, S. Patrui, V. Dandamudi-Ratnakar and P. Sharma, Scalable, Accountable Privacy Management for Large Organizations, *INSPEC 2009: 2nd International Workshop on Security and Privacy Distributed Computing, Enterprise Distributed Object Conference Workshops (EDOCW 2009)*, IEEE, pp. 168-175, September 2009.

Tomas Sander received his doctoral degree in mathematics from the University of Dortmund, Germany in 1996.

He is a Senior Researcher at HP Labs in Princeton, New Jersey. He is a member of the Systems Security Lab at HP which conducts research in trust, security and privacy technologies. Before joining HP, he worked for STAR Lab, the research lab of InterTrust Technologies in Santa Clara, California on a broad range of topics relevant to advanced digital rights

management (DRM). From September 1996 to September 1999 he was a postdoctoral researcher at the International Computer Science Institute in Berkeley, California. His research interests include privacy, computer security, cryptography and digital rights management. In the last few years he has been researching and developing technology that assists implementing good privacy practices in large organizations.

Siani L. Pearson (SM'08) has an MA in mathematics and philosophy from Oxford and a PhD in artificial intelligence from Edinburgh.

She is a Senior Researcher in Systems Security Lab (HP Labs Bristol). She was a Fellow at the Computer Lab in Cambridge University, and for the last 16 years has worked at HP Labs in a variety of research

and development programs including collaborations with HP business units and EU PRIME (Privacy and Identity Management for Europe) project. Her research focus is on privacy enhancing technologies. She is currently a technical lead on an accountability modelling project with HP Privacy Office, and on the collaborative TSB-funded EnCoRe (Ensuring Consent and Revocation) project.