

SECURITY AND PRIVACY GOVERNANCE IN CLOUD COMPUTING VIA SLAs AND A POLICY ORCHESTRATION SERVICE

Marco Casassa Mont¹, Kieran McCorry², Nick Papanikolaou¹, Siani Pearson¹

¹Cloud and Security Lab, HP Labs, Bristol, United Kingdom

²Enterprise Services CTO Office, Hewlett Packard Ltd, United Kingdom
{marco.casassa-mont,nick.papanikolaou,siani.pearson}@hp.com

Keywords: cloud security, privacy, SLAs, decision support systems, enterprise computing, information governance

Abstract: We present in this paper the novel concept of a *policy orchestration service*, which is designed to facilitate security and privacy governance in the enterprise, particularly for the case where various services are provided to the enterprise through external suppliers in the cloud. The orchestration service mediates between the enterprises' internal decision support systems (which incorporate core security and privacy recommendations) and the cloud-based service providers, who are assumed to be bound by contractual service level agreements with the enterprise. The function of the orchestration service, which is intended to be accessed as a trusted service in the cloud, is to ensure that applicable security and privacy recommendations are actioned by service providers through adequate monitoring and enforcement mechanisms.

1 INTRODUCTION

The potential offered by cloud computing for the provision of core business services anytime, anywhere, to anyone is huge. Customers of cloud service providers have access to potentially vast computational resources (not just raw computing power, but also data storage and networking infrastructure), and enterprises are enticed to outsource core business activities and functions to the cloud due to the benefits of reduced cost and increased efficiency that can be thus made. It is actually possible to outsource the work of entire business units to providers of cloud-based business software, such as *Salesforce.com*.

However, security concerns and worries over customer privacy have proven to be fundamental barriers to the adoption of cloud-based services. Enterprises that handle large amounts of customer data need to take numerous measures to comply with security standards, data protection laws, and core business principles with regards to preventing harm to customers' data and their reputations.

The usual model of information security management within the enterprise is given by the lifecycle depicted in Figure 1. This *information*

security lifecycle is a very high-level view of the processes that should be repeatedly carried out in order to assess and implement security requirements;

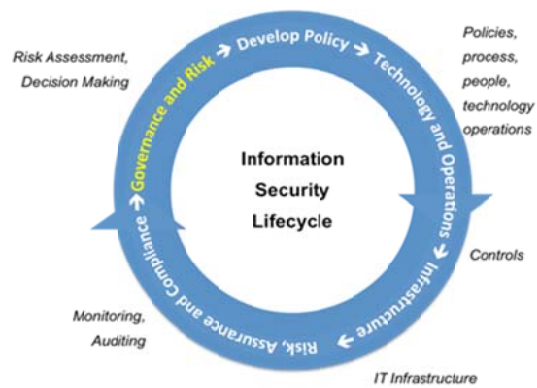


Figure 1. The Information Security Lifecycle.

we note here that specifically for privacy issues, there exist similar models of information management, including for example privacy impact assessments (see ICO 2009), but we refer to the information security lifecycle here as the starting

point for our considerations. The lifecycle does not specify how the processes of, say, risk assessment, monitoring and auditing are implemented in practice, whether through a cloud service or otherwise. Our focus is on the governance of information within the enterprise, and the processes that enable both privacy and security requirements to be satisfied.

What the lifecycle shows is that the handling of such requirements within the enterprise is a cyclic process, which requires constant re-assessment and re-implementation of appropriate measures. For example, risk assessments need to be conducted regularly in order to determine necessary data governance practices, which in turn need to be encoded into policies and enforced using suitable controls.

The objective in this paper is to describe a means for automating information governance in enterprises that outsource important information processing tasks to external cloud service providers. Our contribution is the design of a system component known as a policy orchestrator or *policy orchestration service*, and an accompanying architecture, whose purpose is to automate a significant part of the governance processes (mainly the monitoring and enforcement of privacy and security requirements) in this setting.

1.1 Related Work

Decision support systems for business applications are covered extensively in Turban et al. (2010). We note that there is a body of research on algorithms for decision support systems that support optimal selection of suppliers (see for example Ghodsypour et al.); while in this paper we are not interested in optimality criteria for supplier selection directly, the idea of using decision support systems to manage external suppliers supporting business functions has some relevance to our ideas.

Governance, risk management and compliance (GRC) platforms (see Gartner 2011) are widely used in enterprises to monitor flows of data and ensure compliance requirements are satisfied throughout businesses processes.

Agent technology (see e.g. Padgham and Winikoff 2004) is currently used at the end-points (systems, apps, solutions and/or processes) to enforce and monitor activities. However, most current approaches involve manual steps to inform these agents about what to do. Information has to be gathered about the affected systems, including required configurations and access rights; relevant people need to be identified and involved to provide this data; somebody has to instantiate

templates/scripts based on the above details; various enforcement and monitoring activities need to be carried out. Potentially new decisions or decision changes require repeating this process all over again.

Collaborative and knowledge management tools are of relevance, particularly due to the complexity of maintaining a large rule base. Knowledge management techniques and tools are widely used (see Alavi and Leidner 1999).

Workflow and scripting solutions exist widely (see Jackson and Twaddle 1997 for an account of the basic principles involved); the very idea of an orchestration service is heavily inspired by workflow management systems, as it can effectively be regarded as a means to implement workflows for decision support system recommendations in the cloud.


2 HANDLING PRIVACY AND SECURITY GOVERNANCE IN THE CLOUD

In order to cope with the deluge of rules and restrictions that need to be met, enterprises use special information governance platforms, namely, tools that enable managers, chief information security officers (CISOs) and privacy officers to monitor how data is stored, handled and processed at different control points within the enterprise. In addition, the use of *decision support systems* is commonplace; these are automated tools that intelligently generate lists of recommendations that, when followed in practice, will ensure that particular requirements (for our purposes, security and privacy requirements) are satisfied.

Decision support systems do not usually include means to enforce or action the recommendations they produce; the output of a decision support system is a visual display of information, often just an itemized list of actions that a human user should carry out. In addition, the usefulness of any decision support system's output is dependent on the accuracy and completeness of its rule base; if the rules that the system incorporates are not up-to-date, or do not take into account all external factors, such as, for example, all the different laws and regulations that impact an enterprise's security practices, then that decision support system will not serve its purpose well, and possibly lead to incorrect decisions on the part of the human user.

We are interested in designing a system architecture that helps manage an enterprise's information governance practices when it utilizes cloud services from external suppliers. Our concrete

contribution in this paper is the design of a core component for such an architecture, namely, the *orchestrator*, whose function is to monitor and enforce security and privacy recommendations. These recommendations would typically be customised for a particular project/application, and it is likely that they would be produced automatically by means of a decision support system – an example of such a system is the HP Privacy Advisor (see Pearson et al. 2009). However, the orchestrator is not intended to be used only with customised, auto-generated inputs; it might use other sources for privacy and security rules, such as databases with pre-defined corporate policies.



The orchestrator takes as its inputs a list of security and privacy recommendations, the set of service level agreements detailing the relationship and expectations of the enterprise with regards to its suppliers, and a set of templates which map recommendations to technical policy rules (such as access control rules for particular types of data). The output of the orchestrator is a set of policies and a set of instructions to be carried out and implemented by the external suppliers. If a particular recommendation cannot be actioned directly, or is not satisfied by any of the SLAs available, then the orchestrator will notify the system user accordingly. Clearly these functions, which the orchestrator provides, are of fundamental practical importance to the enterprise; without them, the entire process of making sure that security and privacy recommendations are in fact implemented in practice becomes fully manual and hence error-prone. Since the orchestrator is a software service, it can perform both monitoring and enforcement functions automatically, with little human intervention as and when needed.

To drive this point home, consider what actions are normally needed to practically implement a decision support system’s recommendations:

- identifying the suitable actions to be carried out

- gathering relevant information at the operational level (for example, about the system affected, the people involved, the access rights)
- delegating steps and activities to IT administrators in charge of specific fields (so that they can then monitor compliance against the recommended changes).

Current approaches to security and privacy governance within the enterprise provide little visibility about the involved ecosystem (people, systems, activities); they are primarily manual; and they do not actively involve all the necessary stakeholders to deal with the overall process.

The orchestrator, or *policy orchestration service*, that we propose performs the following in an automated manner:

- (1) it enables a collaborative exchange of information between relevant stakeholders;
- (2) it ‘orchestrates’ the translation of a decision support system’s recommendations into a set of commands corresponding to enforcement and monitoring activities to be executed by the external service provider;
- (3) it supports the execution and monitoring of these activities;
- (4) it traces and audits them.

We detail these functionalities in the next section.

3 A CLOUD-BASED ORCHESTRATION SERVICE

The key idea here, as we have discussed in the previous sections, is to achieve automation in the way decision support systems’ recommendations are transformed into enforceable activities by (1) leveraging workflow capabilities and (2) integrating multiple sources of information. Specifically we introduce a *framework* (in particular, an *assurance service*) that provides assurance and “*orchestration capabilities*” to achieve this.

There are two processes carried out by the orchestration service:

- The *collaborative creation of Templates* to map various types of decision support system outputs into enforceable and monitorable actions
- The *collaborative instantiation of these Templates*, for a specific set of decision support system recommendations, in a given project, into a set for enforceable and monitorable actions.

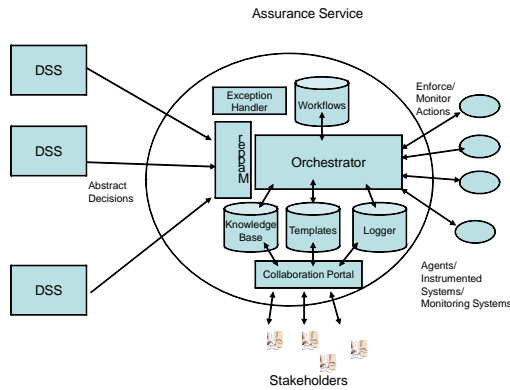


Figure 3. Architecture for the Assurance Service (where DSS = Decision Support System)

Figure 3 illustrates the high level components of this Assurance Service. We assume the scenario of an enterprise, with a set of security and privacy requirements, which are fed into a decision support system in order to produce guidelines and recommendations. In the context of this paper, the requirements relate to a business process that needs to be outsourced to external, cloud-based suppliers or service providers.

The suppliers are bound by service level agreements, whose statements constitute restrictions and obligations as to how the business process will be carried out. The SLAs are fixed and are not expected, under normal circumstances, to change.

We assume that various organisational decision support systems interface with the orchestration service by providing a “formal” representation of the decisions that have been made, such as lists of abstract, required actions and activities to be enforced.

At the very core, within the Assurance Service, there is an Orchestration Engine, driven by various workflows (referred to as “workflow templates”) that can be defined by administrators or as the result of collaborative activities between various stakeholders.

The Assurance Service, by means of the Mapper component, maps the inputs from decision support systems into suitable workflows that need to be carried out by the orchestrator.

Specifically, the orchestrator is able to identify potential issues in the refinement process, for example due to lack of information, conflicting inputs or system unavailability and raise exceptions to be handled by users.

The purpose of the orchestrator component is to determine whether (and how) a particular business process can be outsourced subject to the constraints of the SLAs, and issues appropriate monitoring and enforcement commands to the external service

providers. The orchestrator also takes a set of templates (whose exact format depends on the intended usage scenario – we will not specify this here) that define how each type of recommendation produced by the decision support system is mapped to a technical policy and/or action command to be executed automatically. For example, a security recommendation that requires a given data set to be restricted for exclusive access by a specified group of users will be mapped to a technical access control policy (for example, in the standard eXtended Access Control Markup Language – XACML), and a set of commands that can be executed to enforce this access control policy.

In the ideal situation, when all information is available, the orchestrator (by using these workflows) will automate the following steps:

- identify the suitable templates required to map abstract actions into enforceable and/or monitoring activities. We will refer to these as “refinement templates”;
- collect from the internal knowledge base the actual information necessary to instantiate the templates (e.g. involved systems, required configurations, involved access rights, etc.) and instantiate them;
- interact with the relevant entities (people, systems, applications, processes, etc.) to ensure that specific actions/items are enforced or monitored. This step could be mediated by suitably deployed agents;
- collect and consolidate audit logs from these various entities.

However, as previously mentioned, this is not so trivial. The details required for mapping actions into actionable items might not be available; specific configurations might still be needed; specific inputs and authorizations might need to be provided by the key stakeholders.

In this context, the assurance service, by means of the orchestrator and an associated Portal, orchestrates the interactions between the various stakeholders to obtain the relevant information. This is again driven by a set of Templates indicating what information is required to enable the mapping of abstract decisions into monitorable and enforceable policies/actions.

The relevant templates might have previously built by using the same collaborative service, by factoring in input from different stakeholders, on how to effectively map specific types of decisions support system recommendations into enforceable and monitorable actions.

There might be situations where the relevant templates are not available. In this case, a process is initiated by the assurance service/orchestrator to

generate these Templates, by involving the relevant stakeholders.

Whilst traditional workflows aim at carrying out various business process steps, in this context the orchestrator's main objective is to:

1. enable the refinement of abstract decisions made by a decision support system, by integrating know-how from different sources and triggering the relevant steps, including interactions with people and systems;
2. support the creation of relevant templates, for various types of decision support system outputs.

The Portal is a collaborative web service where various people in the organisation can register, access various relevant information, based on their roles (templates, knowledge bases, etc.), and contribute to the specification of these information. The Portal is not passive: it is used by the orchestrator to require stakeholder interventions, if necessary. As previously described, the orchestrator (driven by workflow templates) will prompt the stakeholders for the necessary information.

We envisage two main interaction mechanisms:

- Stakeholders are asked to intervene and provide information based on need – i.e. during the execution of workflows. This is pretty much standard practice;
- The assurance service, via the portal, actually provides a collaborative service where the various involved people can interact upfront, share information and collaborate to create the various templates necessary during the mapping process. In other words, this portal provides an additional way to generate the “templates”. These interactions might be triggered by challenges raised by administrators and decision makers rather than just the orchestrations and/or involved workflows.

In the latter case the assurance service and orchestrator provide an active ecosystem where various stakeholders, with different skills and expertise, can collaboratively discuss and create material that is relevant for the enforcement of decisions, within the organisation.

4 REVIEW AND CONCLUSIONS

The main advantage of the orchestration service is that it addresses the problem of enforcing and monitoring decisions by:

- Integrating input and contributions of various stakeholders, by means of an automated process
- Enabling compliance with agreed decisions

Whilst common solutions in this space rely on manual interventions and/or automated but very specific solutions (with static knowledge bases valid only in very restricted domains – subject to expensive maintenance/extensions), the proposed solution solves these problems by collaboratively involving the various stakeholders in the process; getting their input to update knowledge bases; ensuring that templates, scripts and mapping mechanisms evolve over time, based on needs.

ACKNOWLEDGEMENTS

We acknowledge the valuable input of Tomas Sander, Prasad Rao and Richard Brown, whose comments helped shape our ideas around the policy orchestration service.

REFERENCES

- Alavi, M., Leidner, D.E. 1999. *Knowledge management systems: issues, challenges, and benefits*. Communications of the AIS 1 (2).
- Gartner. 2011. *Magic Quadrant for Enterprise Governance, Risk and Compliance Platforms*.
- Ghodspour, S.H., O'Brien, C. 1998. *A decision support system for supplier selection using an integrated analytic hierarchy process and linear programming*, International Journal of Production Economics, Volumes 56–57, 20 September 1998, Pages 199-212.
- ICO. 2009. *Privacy Impact Assessment Handbook*. Available online at http://www.ico.gov.uk/upload/documents/pia_handbook_html_v2/files/PIAhandbookV2.pdf.
- Jackson, M., Twaddle, G. 1997. *Business Process Implementation: Building Workflow Systems*, Addison-Wesley.
- Padgham, L., Winikoff, M. 2004. *Developing Intelligent Agent Systems: A Practical Guide*. Wiley-Blackwell.
- Pearson, S., Rao, P., Sander, T., Parry, A., Paull, A., Patruni, S., Dandamudi-Ratnakar, V. and Sharma, P., 2009. *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.
- Turban, E., Sharda, R., Delen, D. 2010. *Decision Support and Business Intelligence Systems*. Pearson.