Towards Greater Accountability in Cloud Computing through Natural-Language Analysis and Automated Policy Enforcement

Nick PAPANIKOLAOU[1], Siani PEARSON[1], Marco CASASSA MONT[1], Ryan K. L. KO[2]
[1]*Cloud and Security Lab, HP Labs Bristol, UK*
Email: {nick.papanikolaou,siani.pearson,marco.casassa_mont}@hp.com
[2]*Cloud and Security Lab, HP Labs Singapore*
Email: ryan.ko@hp.com

Abstract: We argue in favour of a set of particular tools and approaches to achieve accountability in cloud computing. Our concern is helping cloud providers achieve their security goals and meeting their customers' security and privacy requirements. The techniques we propose in particular include: natural-language analysis (of legislative and regulatory texts, and corporate security rulebooks) and extraction of enforceable rules, use of sticky policies, automated policy enforcement and active monitoring of data, particularly in cloud environments.

1. Introduction

For cloud services to be adopted on a wide scale by businesses and individuals, it is necessary for vendors to provide adequate security and privacy controls for the data stored in their systems. In order to ensure compliance with applicable law and standards, and adherence to particular customer requirements (e.g. "Certain types of data should not be stored beyond the national boundaries of Canada or in a public cloud"), vendors need to constantly monitor access and use of their infrastructure and protect against an increasing number of threats. The challenge of accountability is a central concern for vendors, and meeting this challenge means being able to trace the location, flows, instances and accesses of the data stored in their infrastructure [1,2].

There is currently no widely accepted methodology or toolset for technically achieving accountability in cloud computing, with potential solutions being heavily dependent on the particular platform and virtualization technology used by a vendor. What is clear is that a variety of mechanisms are needed to protect against data leakage and to enforce legislation and other related restrictions on the storage and transfer of data, especially across national borders.

2. Objectives

Our objective is to identify top-down, automated means for cloud service providers to provide accountability with regards to their data governance practices. In the context of this paper accountability is understood as the goal of preventing harm to a cloud provider's customers by enforcing adequate protections on these customers' data, and having available effective reporting and auditing mechanisms. More definitions of accountability can be found in [2,3].

While accountability in the broadest sense can be guaranteed only through a combination of law, regulation and technical enforcement mechanisms (e.g. in the context of privacy, such mechanisms are Privacy Enhancing Technologies), our focus is on the technical aspects. As stated in the introduction, what is practically required for a cloud provider to be accountable is a set of tools to track the location, flows, and accesses of its customers' data. As we shall see, this capability allows a provider to readily demonstrate compliance to the law and adherence to all relevant regulations and other restrictions. More importantly, this capability should allow any instances of non-compliance to be detected easily, so that suitable corrective action can be taken.

3. Methodology

We argue that it is possible to automate many of the processes required to ensure that a provider is accountable, although we recognise the difficulty of mapping and linking legal and regulatory requirements - which are high-level and expressed in natural language - to technically enforceable policies on particular data items.

Key techniques that can be used to achieve a significant degree of automation include:

- **Natural-language analysis**, in particular, extraction of policy rules from legislative and regulatory texts and corporate rulebooks; these rules should be represented in a form that can be interpreted by a technical enforcement mechanism (esp. a Policy Enforcement Point or PEP), but possibly also so that they can be incorporated into a compliance checker of information governance software (cf. Governance/Risk Management/Compliance (GRC) Platforms, widely used in industry). It should be noted here that no natural-language processing system can operate with 100% accuracy, but use of such systems can help to reduce significantly the overall amount of human intervention in the process of policy creation and management.
- **Usage of sticky policies**: by strongly binding policies to the data they are associated with, it is easier for providers to control accesses to data within their cloud infrastructure and there is no need for a central policy repository. From the point of view of automating accountability, the use of sticky policies is a very useful technique. Sticky policies provide a means of data encryption, since the data which a policy is bound to cannot be accessed unless that policy is complied with. [4]
- **Automated policy enforcement:** the deployment of control points throughout a cloud provider's infrastructure where policy rules can automatically be enforced, and human users only notified in case of failure or error is essential. We refer to the following current and future HP Labs European and TSB research projects for more related work on policy enforcement: EnCoRe (http://www.encore-project.info), Information Stewardship in the Cloud [14], and the TrustDomains project.
- **Active monitoring for compliance:** we believe that it is fundamental for cloud providers to have in their infrastructure mechanisms for automatically detecting compliance problems and potential sources of such problems. It is possible to formulate and regularly check system invariants corresponding to conditions that should never occur at certain end points, such as links between a provider's data centres, and particularly cross-border links. We refer to the TrustCloud research project by HP Labs' Cloud and Security Lab [1,2].

4. Conclusions

We believe that it is beneficial and possible for cloud service providers to automate a number of tasks related to the requirement of accountability. We have identified some specific techniques, namely: natural-language analysis of law, regulation and corporate

guidelines on security and privacy of customer data in order to generate technically enforceable policies; use of sticky policies to achieve a strong binding between data and the stipulations that apply to the use and dissemination of that data; and active monitoring of a cloud provider's infrastructure to detect potential compliance problems.

We are actively working on the development of all these techniques [3-13] which, combined with the deployment of technical policy enforcement mechanisms in a cloud provider's infrastructure, can help achieve accountability, which is a major concern in cloud computing today.

Bibliography

1. Ryan K L Ko, Bu Sung Lee, Siani Pearson, "Towards Achieving Accountability, Auditability and Trust in Cloud Computing", International Workshop on Cloud Computing-Architecture, Algorithms and Applications (Cloudcomp 2011), Springer, Kochi, India, 22-24 July 2011.
2. Ryan K L Ko, Peter Jagadpramana, Miranda Mowbray, Siani Pearson, Markus Kirchberg, Qianhui Liang, Bu Sung Lee, "TrustCloud - A Framework for Accountability and Trust in Cloud Computing", IEEE 2nd Cloud Forum for Practitioners (IEEE ICFP 2011), IEEE Computer Society, Washington DC, USA, 7-8 July 2011.
3. Siani Pearson and Azzedine Benameur, Privacy, Security and Trust Issues Arising from Cloud Computing, Proc. CPSRT 10, CloudCom, IEEE, November 2010.
4. Siani Pearson, Marco Casassa Mont and Gina Kounga, Enhancing Accountability in the Cloud via Sticky Policies, STAVE, Springer, June 2011.
5. Siani Pearson and Andrew Charlesworth, Accountability as a Way Forward for Privacy Protection in the Cloud, Proc. 1st CloudCom 2009, ed. M.G. Jaatun, G. Zhao, C. Rong, Beijing, Springer LNCS 5931, pp. 131-144, December 2009.
6. Siani Pearson, Marco Casassa Mont and Manny Novoa, Securing Information Transfer within Distributed Computing Environments, *IEEE Security & Privacy Magazine*, Jan/Feb issue, volume 6, number 1, pp. 34-42, IEEE, 2008.
7. Tariq Ehsan Elahi and Siani Pearson, Privacy Assurance: Bridging the Gap Between Preference and Practice↓, C. Lambrinoudakis, G. Pernul, A.M. Tjoa (eds.), *Proc. TrustBus 2007*, LNCS 4657, pp. 65-74, Springer-Verlag Berlin Heidelberg, 2007.
8. Marco Casassa Mont, Siani Pearson, Robert Thyne, A Systematic Approach to Privacy Enforcement and Policy Compliance Checking in Enterprises, *Trust and Privacy in Digital Business*, LNCS 4083, Springer Berlin/Heidelberg, pp. 91-102, 2006
9. Marco Casassa-Mont, Siani Pearson and Pete Bramhall, Towards User Control and Accountable Management of Privacy and Identity Information, *Proc. ESORICS*, pp. 146-161, LNCS 2808, Springer, 2003.
10. Marco Casassa Mont, Filipe Beato - On Parametric Obligation Policies: Enabling Privacy-aware Information Lifecycle Management in Enterprises, 8th IEEE Workshop on Policies for Distributed Systems and Networks, Policy 2007, 13-15 June 2007, Bologna, Italy [Conference Presentation], 2007
11. Adrian Baldwin, Marco Casassa Mont, Yolanta Beres, Simon Shiu - On Identity Assurance in the Presence of Federated Identity Management Systems, ACS CCS 2007 Workshop on Digital Identity Management, DIM 2007, 2 November 2007, George Mason University, Fairfax, VA, US, 2007
12. Marco Casassa Mont, Robert Thyne - Privacy Policy Enforcement in Enterprises with Identity Management Solutions - 4th International Conference on Privacy, Security and Trust 2006, PST 2006, 30 October, 01 November 2006, Toronto, Canada
13. Marco Casassa Mont - On the Need to Explicitly Manage Privacy Obligation Policies as Part of Good Data Handling Practices - W3C Workshop on Languages for Privacy Policy Negotiation and Semantic-Driven Enforcement, W3C Privacy Workshop 2006, 17-18 October 2006, Ispra, Italy, 2006.
14. D. Pym and M. Sadler, “Information Stewardship in Cloud Computing,” *Int’l J. Service Science, Management, Engineering and Technology*, vol. 1, no. 1, 2010, pp. 50–67.