AUTOMATED UNDERSTANDING OF CLOUD TERMS OF SERVICE AND SLAs

Nick Papanikolaou, Siani Pearson, Marco Casassa Mont – Cloud and Security Lab, Hewlett Packard Laboratories, Bristol, UK

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.

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 need to be put into place to protect against data leakage and to enforce legislation and other related restrictions on the storage and transfer of data, especially across national borders.

Our objective is to identify 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.

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. What is practically required for a cloud provider to be accountable is, among other things, 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 demonstrate compliance to the law and adherence to all relevant regulations and other restrictions. More importantly, this capability allows any instances of non-compliance to be detected effectively, so that suitable corrective action can be taken.

Of course, the capability to provide and demonstrate compliance needs to be founded on privacy law and secured based on best practices and industry standards. Any platform to provide accountability needs to be secured so that it cannot be exploited by attackers.

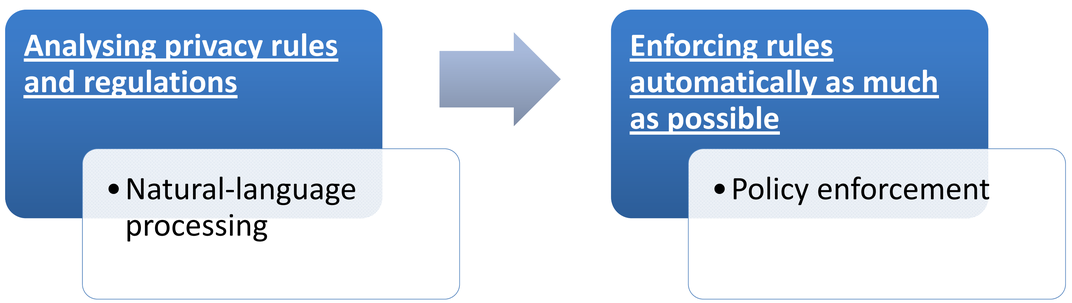

Fig.3. The basis of our approach is to use natural-language processing techniques and tools; once rules have been extracted from texts, we can use existing policy enforcement techniques, as developed in European and UK-wide research projects such as PRIME, PrimeLife and EnCoRe.

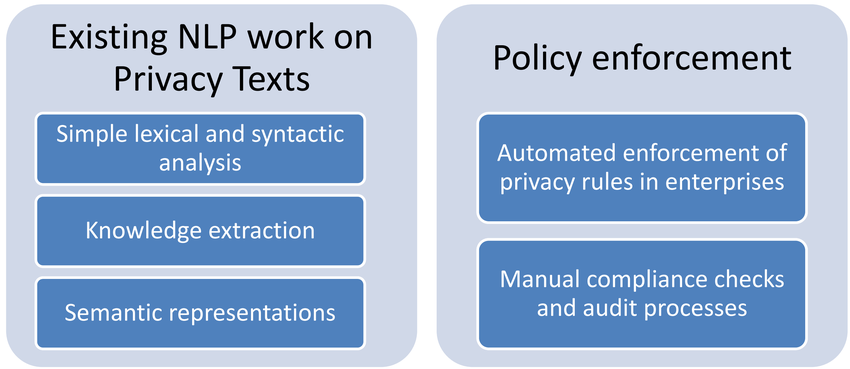

Fig. 1. We've looked at existing work on extracting information from legal and regulatory privacy texts and we have significant experience in policy enforcement, access control and compliance methods. Bringing these techniques together we can develop powerful tools for analysing Cloud Service Providers' terms of service and SLAs.

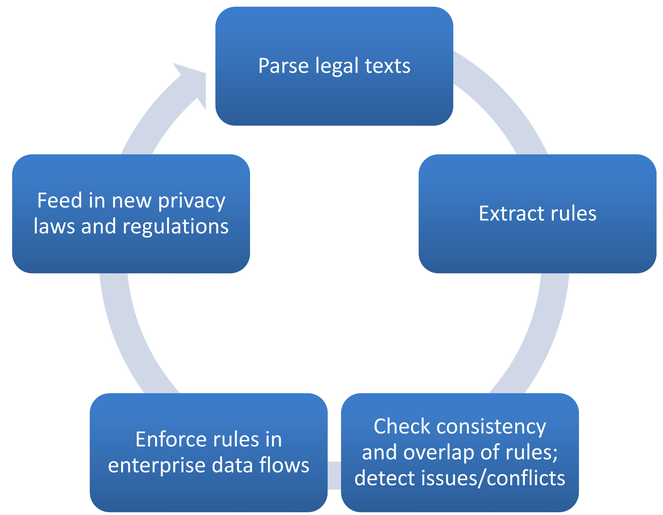

Fig. 2. This diagram depicts a lifecycle for privacy rules that are extracted from texts and stored in a central repository from where they can be drawn for enforcement in a real system.

Contact us at: {nick.papanikolaou, siani.pearson, marco_casassa-mont}@hp.com