Chapter 8.6
Securing Real-Time Interactive Applications in Federated Clouds

Michael Boniface
University of Southampton IT Innovation Centre, UK

Bassem Nasser
University of Southampton IT Innovation Centre, UK

Mike Surridge
University of Southampton IT Innovation Centre, UK

Eduardo Oliveros
Telefónica Investigación y Desarollo, Spain

ABSTRACT

Enterprise adoption of cloud computing for real-time interactive applications processes is limited by their ability to meet inter-enterprise security requirements. Although some clouds' offerings comply with security standards, no solution today allows businesses to assess security compliance of applications at the business level and dynamically link to security countermeasures on-demand. In this chapter we examine cloud security, privacy, and trust issues from three levels: business, jurisdiction, and technical. Firstly, we look at the business level to identify issues arising from the motivations and concerns of business stakeholders. Secondly, we explore jurisdictional level to identify risks that arise from legislation, gaps in legislation, or conflicts between legislation in different jurisdictions related to a cloud deployment, given the concerns of stakeholders. Finally, we examine the technical level to identify issues that arise from technical causes such as ICT vulnerabilities, and/or require technical solutions, such as data confidentiality and integrity protection.

DOI: 10.4018/978-1-4666-0879-5.ch8.6

INTRODUCTION

The cloud is such a general-purpose paradigm that it is impossible to consider 'the cloud' as a single set of business models with a single set of security, privacy and trust issues. To some extent, issues with cloud computing are necessarily related to the application purpose. However, using the cloud modalities (IaaS, PaaS and SaaS) it is possible to identify common stakeholders and concerns in each classification (Rimal, Choi, & Lumb, 2009).

Infrastructure as a service (IaaS): the provision of 'raw' machines (servers, storage and other devices) on which the service consumer installs their own software (usually as virtual machine images). The service is billed on a utility computing basis according to the amount of resources consumed. IaaS stakeholders include the IaaS hoster, provider and customer. The IaaS hoster must provide adequate resources in order to meet demands of its customers needs, together with appropriate availability contingencies. It will also need to avoid becoming liable for illegal uses of the software including licence violations. The IaaS provider must provide an interface in order for customers to configure and manage the resources they will use, to upload workloads and data. Utility billing or the resources used together with usage data must be provided. Today it is normal for the IaaS provider to also host the cloud resources, but in some research projects (Edutain@Grid and RESERVOIR) multi-hosting models are being developed in which at least some hosters delegate the provision of customer-facing cloud management services to a separate entity (Ferris, Surridge & Glinka, 2009; Sotomayor, Montero, Llorente & Foster, 2009). The IaaS customer will be able to choose from a range of resources, services and SLAs that best meet their need. This may include tiered pricing for different kinds of configuration (performance, capacity), which may be provisioned as shared or dedicated resources, or levels of security.

The PaaS stakeholders include the PaaS hoster, provider and user (developer). The PaaS hoster must provide adequate resources (typically via an IaaS model) in order to meet demands of its customers needs, together with appropriate availability contingencies. The PaaS provider provides an environment suitable for general developers to build web applications without deep domain expertise of back-end server and front-end client development or website administration. The PaaS user (developer) must have a browser-based development environment, the ability to deployment seamlessly to a hosted runtime environment, management and monitoring tools and pay as you go billing.

The SaaS stakeholders include hoster, provider, customer, consumers and application software vendors. The SaaS hoster needs to protect their infrastructure from misbehaving applications, and avoid becoming liable for illegal uses of the software including licence violations, etc. This stakeholder may be covered under IaaS, of course. The SaaS provider needs to restrict access to paying customers, and ensure the hoster provides the necessary performance and respects other consumer requirements such as confidentiality, personal data protection, etc. The SaaS customer needs to have adequate performance and pay only for what they used, protect sensitive application data (inputs, outputs and stored data) including protection of any personal data used, and may need the fact they used the service to be confidential. The SaaS consumers are people who use the service purchased by the SaaS customer – e.g. friends, colleagues or downstream customers – who may be owners or data subjects for some of the data being used in the service. The application software vendor needs to define and enforce a licensing model that compensates them for use of their application via SaaS, ensure effective user support, etc.