Towards User Centric Data Governance and Control in the Cloud

– Extended Abstract –

Stephan Groß and Alexander Schill
Technische Universität Dresden
Institut für Systemarchitektur
01062 Dresden
E-Mail: {stephan.gross, alexander.schill}@tu-dresden.de

Abstract: Cloud Computing, i. e. providing on-demand access to virtualised computing resources over the Internet, is one of the current mega-trends in IT. Today, there are already several providers offering cloud computing infrastructure (IaaS), platform (PaaS) and software (SaaS) services. Although the cloud computing paradigm promises both economical as well as technological advantages, many potential users still have reservations about using cloud services as this would mean to trust a cloud provider to correctly handle their data according to previously negotiated rules. Furthermore, the virtualisation causes a location independence of offered services which could interfere with domain specific legislative regulations. In this paper, we present an approach of putting the cloud user back into power when migrating data and services into and within the cloud. We outline our work in progress, that aims at providing a platform for developing flexible service architectures for cloud computing with special consideration of security and non-functional properties.

1 Motivation

The recent progress in virtualising storage and computing resources combined with service oriented architectures (SOA) and broadband Internet access has led to a renaissance of already known concepts developed in research fields like grid, utility and autonomic computing. Today, the term *cloud computing* describes different ways of providing on-demand and pay-per-use access to elastic virtualised computing resource pools [1]. These resources are abstracted to services so that cloud computing resources can be retrieved as infrastructure (IaaS), platform (PaaS) and software (SaaS) services respectively. The pay-per-use model of such service oriented architectures includes *Service Level Agreements (SLA)* negotiated between service provider and user to establish guarantees for required non-functional properties including mandatory security requirements. The (economical) advantages of this approach are fairly obvious: One saves costly investments for procuring and maintaining probably underused hardware and at the same time gains new flexibility to react on temporal higher demands.

Nevertheless, there are reasonable reservations about the deployment of cloud computing services, e.g. concerning data security and compliance. Most of these concerns result from the fact, that cloud computing describes complex socio-technical systems with a high number of different kinds of stakeholders following different and possibly contradicting objectives. From a user's perspective, one has to hand over the control over his data and services when entering the cloud, i.e. the user has to trust that the cloud provider behaves in compliance with the established SLA. However, to actu-

ally agree on a specific SLA a user first has to assess his organizational risks related to security and resilience [2].

Current solutions that restrict the provision of sensible services to dedicated private, hybrid or so-called national clouds[1] do not go far enough as they reduce the user's flexibility when scaling in or out and still force him to trust the cloud provider. Furthermore, private clouds intensify the vendor lock-in problem. Last but not least, there is no support for deciding which services and data could be safely migrated to which cloud. Instead we demand new methods and technical support to put the user in a position to benefit from the advantages of cloud computing without giving up the sovereignty over his data and applications. In our current work, we follow a system oriented approach focussing on technical means to achieve this goal.

The remainder of this extended abstract is structured as follows. We first refine our problem statement in section 2. Then, in section 3, we sketch our approach of developing a service platform providing the necessary means for building secure, easy to use and flexible service architectures for cloud computing environments. We also outline our idea of a personal cloud, i.e. the conglomerate of a user's resources and devices, that can be controlled by a specialised gateway, the so-called Π-Box. We exemplify, how the Π-Box supports the controlled migration of resources into the cloud.

2 Problem statement

We identified security as a major obstacle that prevents someone to transfer his resources into the cloud. In order to make sound business decisions and to maintain or obtain security certifications, cloud customers need assurance that providers are following sound security practices and behave according to agreed SLAs [3]. Thus, our overall goal is the development of a flexible open source cloud platform that integrates all necessary components for the development of user-controlled and -monitored secure cloud environments. This platform should contain the following components:

1. Mechanisms to enable an user controlled migration of resources and data into the cloud. These mechanisms should support (semi-)automatic configuration of cryptographic algorithms to simplify the enforcement of a user's security requirements as well as the dynamic selection of cloud providers that best fit the user's requirements and trust assumptions. Thus, we need a formalised way to acquire a user's requirements. Furthermore, we need to integrate the user's private resources and different cloud providers in our cloud platform, e.g. by using wrapper mechanisms or standardised interfaces.
2. A sound and trustworthy monitoring system for cloud services that is able to gather all relevant information to detect or even predict SLA violations without manipulations by the cloud provider under control. To support the configuration of the monitoring system, there should be some mechanism that derives relevant monitoring objectives from negotiated

1 The mentioned cloud types define different deployment models of cloud computing systems. In contrast to *public clouds* that make services available to the general public, *private clouds* are operated solely for an organization although the resources used might be outsourced to some service company. *Hybrid clouds* describe a mixture of public and private cloud, i.e. when users complement their internal IT resources with public ones. The term *national clouds* describes a scenario, where the location of the cloud resource pool is restricted to one country or legislative eco-system like the EU.

SLAs. Thus, we need a formalised language for machine-readable SLA focussing on the technical details of a cloud computing environment.

3. The proposed cloud platform should be adaptive, i.e. it should provide mechanisms to react on SLA violations detected by the monitoring system in order to mitigate the resulting negative effects. These mechanisms should include migration tools to transparently transfer resources to another cloud provider as well as adaptation tools that leaves the resources at the chosen provider but transforms them to further meet the user's non-functional and security requirements.

3 FlexCloud – Flexible Service Architectures for Cloud Computing

Within the FlexCloud project we aim at developing such methods and mechanisms to support the development of flexible and secure cloud applications. The foundation of our work is represented by our service platform SPACE/TECЛA [4]. SPACE is an open source platform for the Internet of services which provides basic tools for contract-bound adaptive service execution and acts as a hosting and brokering environment for services. Its latest extension integrates an Amazon EC2-compatible cloud environment as target environment for complex services delivered as virtual machines. Figure 1 sketches our general system design. On the left hand side there is a user's personal

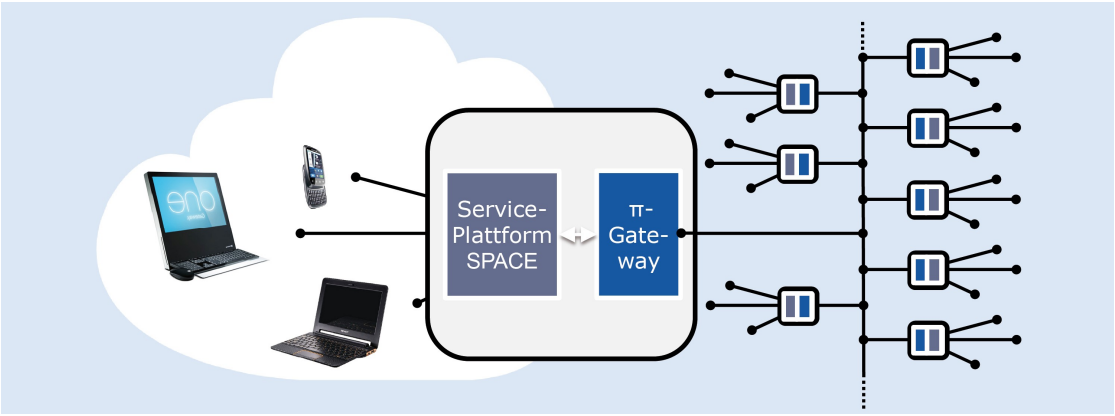

Figure 1: General system design with Π-Box

cloud (Π-Cloud) which is controlled by his Π-Box (rectangle in the middle)[2]. The Π-Box first provides the SPACE-based service execution platform. The relevant security mechanisms described in the previous section are encapsulated in the Π-Gateway component which also interconnects different personal clouds to form a huge virtual cloud environment. Both, the SPACE as well as the Π-Gateway can be provided as virtual machines. Thus, they can either be run on a physical machine, such as an extended DSL router or a dedicated server, or as a virtual appliance within the cloud which means that the Π-Box can also benefit from the cloud's advantages.

The Π-Gateway's functionality is illustrated by an example: confidential cloud storage. Although there already exist first cloud storage solutions providing client-side encryption, these systems introduce new potential for human errors as they ignore available access control systems and must be

2 Please note, that the Π-Box could also control a company's private cloud.

manually configured, e.g. by (re)defining the encryption keys for authorised users. Figure 2 sketches how the Π-Gateway incorporates different tools, such as existing access control and user management systems, face recognition or information retrieval tools to determine the users authorised to access a specific file. For example, it could search a text file for a specific confidentiality note, analyse the people displayed on a foto, or simply check the file system's access rights in combination with the operating systems user database to determine the user identities to be granted access. By using these identities it is able to retrieve the necessary public keys from a public key infrastructure and to encrypt the data accordingly before storing it into the cloud.

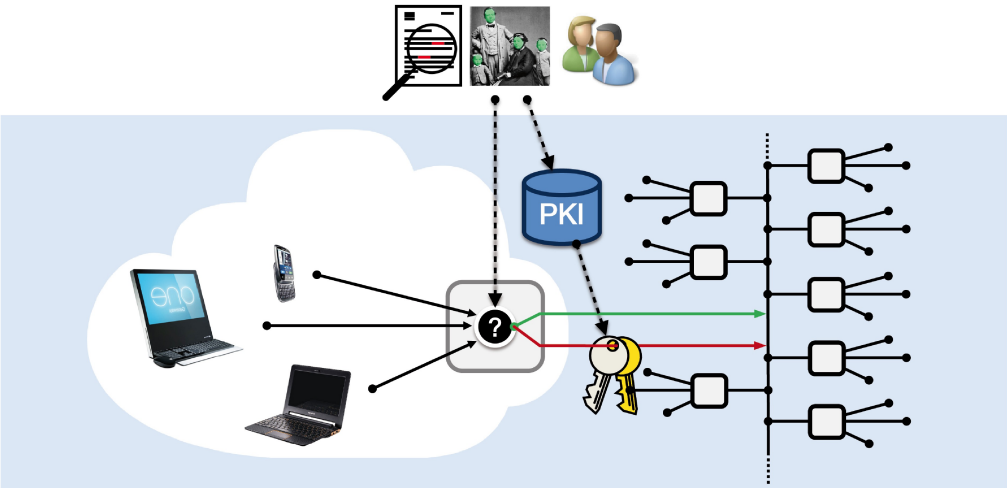

Figure 2: Achieving confidential cloud storage with the Π-Box

4 References

[1] P. Mell and T. Grance: *The NIST Definition of Cloud Computing (Draft).* Recommendations of the National Institute of Standards and Technology (NIST), Special Publication 800—145 (Draft), available at http://csrc.nist.gov/publications/drafts/800-145/Draft-SP-800-145_cloud-definition.pdf, January 2011.

[2] D. Catteddu (Ed.): *Security & Resilience in Governmental Clouds – Making an informed decision.* ENISA Report, January 2011.

[3] D. Catteddu (Ed.): *Cloud Computing – Benefits, risks and recommendations for information security.* ENISA Report, November 2009.

[4] J. Spillner and A. Schill: The TECЛA Service Platform: Web Service Sharing based on Modular Platform Services. In: *FIS - Future Internet Symposium*, Berlin, Germany, September 2009.

5 Acknowledgement

The author would like to express his gratitude to the FlexCloud research group, especially its former member Gerald Hübsch, for many fruitful discussions that contributed to the development of the ideas presented in this paper. Tribute to Marc Mosch for the graphical design of the presented figures.

This work has received funding under project number 080949277 by means of the European Regional Development Fund (ERDF), the European Social Fund (ESF) and the German Free State of Saxony. The information in this document is provided as is, and no guarantee or warranty is given that the information is for any particular purpose.