

Cloud Computing Scenario

1. Background

Wikipedia defines the “Cloud” (or “Cloud computing”) as a metaphor for the Internet: an abstraction of the complex infrastructure it conceals. It stands for a style of computing in which IT-related capabilities are provided “as a service”, allowing users to access these services from the Internet without knowledge of, expertise in, or control over the technology that supports them. In this context, the terms “Software as a Service” or the “Internet of Services” are also used.

Cloud computing has generated much interest recently and developments in this space are quickly ramping up. Services have started appearing and infrastructures are being developed. However, the successful evolution of Cloud infrastructures and service offerings face complex multi-party security requirements; there is also a risk of extensive service deployment which lack respect for these requirements.

At another level, Clouds are defined as, “...a large pool of easily usable and accessible virtualized resources (such as hardware, development platforms and/or services). These resources can be dynamically configured to adjust to a variable load (scale), allowing also for optimum resource utilization. This pool of resources is typically exploited by a pay-per-use model in which guarantees are offered by the Infrastructure Provider by means of customized SLAs.”¹ (Replace resource with infrastructure and this definition is generalised to the previous Wikipedia description.) The realization of Cloud computing is, thus, a complex task requiring many technical and business skills, as well as significant financial investments. Components include: the building and operation of vast data centres supporting the Clouds; high-capacity broadband networks and networks of massively distributed computing systems; highly flexible service-oriented-architectures, supporting dynamic service composition and provisioning based on Web Services Security (WS-Security) and other open standards; highly distributed information infrastructures and virtualization technology.

As part of the Cloud development, unprecedented issues concerning jurisdiction, law enforcement, national security and privacy will emerge due to the global nature of the Cloud. From a security and trust perspective, Cloud computing brings about many new challenges: data protection and security; data ownership, access and retention; liability, accountability and transparency of data usage policies and management; secure separation of virtual environments; identity management and identification of objects and entities to enable ownership management, auditing and security policies.

In the ideal “service-oriented” world, users not need be aware of the complex inner structure of the Cloud from which they are consuming services. As such, consumers largely already behave today as if they were dealing with a location-neutral service without a pre-defined legal environment or a risk context. Although they expect - wrongly in many cases - to deal with a single, defined business partner who they can make responsible for the business partnership in which they are engaging, when using the service.

Uptake issues for Cloud computing are more prevalent on the business-user side, mainly due to a lack of accountability and the difficulty of enforcing security, privacy and trust properties for the services. For example: the storage of business-critical information “in the Cloud”. In principle, this

¹ [Vaq09] L.M. Vaquero, L. Rodero-Merino, J. Cáceres, M. Lindner. A Break in the Clouds: Towards a Cloud Definition. ACM Computer Communication Reviews. January 2009

is a good idea with (virtual) storage capacity provided on the basis of daily changing (physical) storage prices. But how is the confidentiality of the stored information ensured, when it is physically stored one day in China and the next in the Philippines? And this is not to even mention business services such as Customer Relationship Management (CRM) or Enterprise Resource Planning (ERP) offered by the Cloud, where the demand for security, integrity and privacy may be even more critical.

It is possible that the current debate about Cloud computing business cases are taking place without considering the inherent security, privacy and trust requirements. This would explain the difference between the end-user and the business-user uptake: businesses are realizing that security properties may be business-critical when their information is maintained and managed in the Cloud, whereas it is taking longer, collectively, for end-users to realise that these properties may be of benefit to individuals.

Barriers to a sizeable uptake of Cloud computing, from a security and trust point of view include (in no specific order): the multitude of legal frameworks required, to observe and eventually to follow when developing, using and/or migrating services world-wide; the missing (technical) ability to address multi-party security requirements for even very simple scenarios (security and privacy not being a zero-sum game); the weak readiness of existing (software-as-a-) service providers to manage their customers' risk as part of the business case.

2. Vision

Let us take two scenario examples to develop a set of characteristic requirements that shall be met by future Cloud computing architectures (technical, business and policy-related). Note that the focus of the description shall not be on the requirements for the individual, but much more on the requirements for the (Cloud) service providers. The first scenario is the following:

1. Smart Shopping

A woman is walking down the street of her local town centre. The RFID tag in her jacket is contacted by a local reader. The reader sends the tag's data to a localization service. The localization service sends this data to a meta-CRM-system that handles consumer related data for that particular area. The CRM system recognizes the consumer, looks up her preferences and offers her - via an SMS to her mobile phone - a 20% SALE reduction a nearby shop, if she orders immediately (via WAP/mobile phone), which she does after having had a look at the item in the shop.

The payment is processed by a payment service provider, who knows the consumer's details who charges her credit card. For security purposes, an alert is sent (via a web service) to a credit card clearance agency, who checks the credit card number against recent fraud. Unfortunately, there has been a fraudulent

action using this credit card, so the agency informs the police (again via web service). The police management system accesses the location service to get the location of the consumer and sends two policemen from the closest office to speak to the consumer.

The second scenario is the following:

2. Nomadic Business Organisation

A network of business people is spread across the world, but they are working as if they were one company. They do not have any specific physical office space and meetings are held by using on-line conferencing tools provided by the Cloud. They use on-line storage for documents, a service-based customer relationship management (CRM) system, and service-based financial performance management software. The membership of the network is highly dynamic, i.e. people join and leave on a very short-notice basis. The software components are used by this organization via the Cloud, i.e. without knowing where these services actually run. So not only are the consumers of the services “nomadic”, but the services that this organization consumes have no physical location themselves either.

There is a significant difference in the above scenarios when compared to current services and infrastructures. The extended scope, the scale of deployment, the complexity of interdependence, and the dynamism of service interaction all pose problems from a security and privacy point of view. In addition, a core idea behind Cloud computing is to postulate a “*share and you will get more*” mentality as being of higher value for people than their actual safety and protection. Consequently, services and infrastructures for the Cloud are being developed with maximum openness in mind. This of course contradicts current business and societal values such as (individual) privacy and (corporate) information protection, per se. To address this, we would need to investigate: business incentives for security in the Cloud: whether we could actually provide “Cheap Cloud Security”: and, how far security could actually be functionalized? (From a non-functional, infrastructure aspect towards a functional, service-driven added value).

Another ground-breaking aspect of Cloud computing is the presumed volatility of relationships, meaning that relationships between people-and-people or people-and-organizations may change with ever increasing frequency, until there is no long-term relationship to be managed. To some extent, this contradicts the purpose of today’s identity management architectures, where the organizational belonging of an individual is the basis for the whole idea of identity management. There is a need for a new type of identity management framework that respects these ideas of

volatility e.g. by using web service-based short-term assertions, with long-term, pseudonymous accountability properties, such as an “assertion infrastructure”. Authentication and authorization could then be realized by context-based access control (such as using proof-of-ability types of techniques) instead of identity-centric access control. Since confidentiality cannot, once lost, be recovered, there also exists the need for a “*bad-guy*” tolerant infrastructure.

In which legal context will the different services run? Ideally, there would be a single, common legal environment for Cloud computing services. This may exist at some stage in the future, if at all, and thus, there must be mechanisms to provide maximum accountability across the different technical layers along the service invocation chain. This means that track records and protocols are available and legally binding, according to the set of legal frameworks encompassing these services. Such accountability service properties would themselves be subject to complex security, privacy and trust properties. Therefore, there is a need to carefully add security and privacy to SLA automation, and to map compliance requirements to technical policy controls, in addition to run-time accountability, up-front automated processing of standardized clauses.

From a user side, there will be an explosion of identities and devices. There will be a need for identity aggregation on devices, as well as analyzing the impact of location and sensor technology that provides environment transparency and/or disclosure. This aspect is addressed in detail in the “Nomadicity” use case.

It is important to develop solutions to break up the “*Security and Privacy is a zero-sum game paradigm*”. Only then can public security and individual freedom be maintained simultaneously. This requires new technical protocols and innovative solutions that are able to address multi-party security requirements according to predefined policies (supported/coordinated with corresponding policy action). Special attention should be given to supporting privacy rights without infringing on the ability to react for the sake of public security - and vice versa.

To achieve trust, there is a requirement for maximum transparency, along with the ability to perform upfront and *continuous* risk assessments, both for businesses and end-users. To do this, security measurability techniques, tools and processes will be necessary for software and services to be deployed in the cloud.

Summing up, the previous considerations could be grouped into the following substantial themes:

1. Automated Policy Governance
2. Privacy Aware Accountability Infrastructure
3. Measurability and Assessment
4. Cross Topic: Standards, Interoperability and Open Platforms across all entities

There is a chance that such scenarios will be implemented without taking multi-party security, privacy and trust requirements into account; i.e. the business value chain and the public security processes will be implemented without “security in mind”. In other words, as a society, we could start running before we can walk, with long-standing cultural and societal values being places at risk. It is therefore crucial to start the discussions now and to explain to cloud computing innovators the role of the non-functional security and privacy requirements.

3. Roadmap

To achieve a trustworthy Cloud, there is a need to both provide the right services from a security engineering perspective, as well as to educate and mandate qualitative approaches for the standard IT engineering and innovation processes. Consider:

- 1) The necessary services consist of building blocks that can easily be used by non-security savvy IT engineers to realize security properties. These cover, for example, a common, open architecture, frameworks, libraries and free code snippets for multi-party security and privacy policy implementations. There is a need for supporting “standard” IT engineers, since even with the best program for secure software engineering, the ability to develop secure IT services will strongly depend on the availability of frameworks and source code for standard security tools and enhancements. Therefore, research programs are required which aim at developing and offering the necessary frameworks and toolsets for security in the Cloud, including specification, construction and verification methods and tools for the Cloud. A test scenario might be to migrate Cloud enterprise applications in real time.
- 2) Attaching policy information to data: Data-centric access control and security management technologies have to be explored as a means to serve user-centric security approaches. This covers the application of cryptographic techniques, such as DRM/IRM technologies and blinding techniques. The management of data security, based on personal credentials and system security, will not scale in a Cloud environment: the data must be protected “on-the-go”. Therefore, encryption, rights management and authorization mechanisms must be tied to the information itself, while being able to model user-centric security policies. In addition, there is a need for the development of new concepts where existing protocols are not sufficient (too difficult to implement or too costly to use). Finally, these technologies must be able to implement multi-party security requirements, including privacy, especially in an environment of flexible and configurable service delivery. Most important though is the interoperability of data formats and platform independence.
- 3) A pervasive, privacy-aware identity management infrastructure is needed. Given the fact that existing identity management systems will need to co-exist, a meta-approach for assuring integration and interoperability is needed. In addition, the privacy and volatility requirements should be modelled and implemented into that meta-structure, allowing for the assurance of accountability, yet also providing adequate privacy. Dynamic federation is needed between arbitrary technologies and sources of trust, as well as developing one-time-credential technology and assertion mechanisms without disclosing personal/identity information.

For integrating security into IT innovation and engineering processes, awareness, education and measurability are important:

- 4) Policy makers and experts must be educated about the technical capabilities to overcome the zero-sum game paradigm that is currently predominant in these communities.
- 5) A critical component is to be able to measure and/or assess progress in security, so that one can demonstrate the value of integrating security measures. Besides established models for justifying security spending, measurement techniques are required to demonstrate incremental improvements.
- 6) A process driven approach would ensure multi-party security requirements during IT innovation. IT innovation needs to become professionalized, so that non-functional requirements can be taken into account, even if they are not in the focus of the research activities. The best approach would be to establish a standardized, best-practice process framework for projects, recognized both by security and engineering methodology experts, to make sure non-functional requirements such as security, privacy and accountability are met. A good model to start from might be the PRISE privacy process approach².

² <http://www.prise.oeaw.ac.at/>