



HAL
open science

Cloud Computing: New Research Perspectives for Computer & Law

Daniele Bourcier, Primavera de Filippi

► **To cite this version:**

Daniele Bourcier, Primavera de Filippi. Cloud Computing: New Research Perspectives for Computer & Law. 13th International Conference of Artificial Intelligence & Law, Jun 2011, United States. pp.79. hal-00713405

HAL Id: hal-00713405

<https://hal.science/hal-00713405>

Submitted on 6 Nov 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Cloud Computing: New Research perspectives for Computers and Law

Daniele Bourcier¹, Primavera De Filippi¹

¹CERSA - CNRS - Universite Paris II

(daniele.bourcier, primavera.de-filippi)@cersa.cnrs.fr

Abstract. Cloud computing represents a new business paradigm whereby a series of computing resources are offered as a service, available on-demand, on a pay-per-use basis, over the Internet. In this paper, we propose a hypothesis of how cloud computing can be described as a complex system and we describe the various risks and opportunities connected with the current implementation cloud computing. We then present a preliminary model for the implementation an automated system of certification based upon the formalization of contractual rules and consumers' preferences.


Keywords. Cloud Computing, automated agents, contractual negotiations


1. Cloud Computing: A Novel Paradigm of Complexity?

1.1 Definition

The Cloud consists of a distributed infrastructure that is made of a collection of interconnected computers, whose resources are pooled together into a virtual machine that maintains and manages itself. As opposed to other distributed architectures, the particularity of the Cloud is that its architecture is completely independent from the physical infrastructure it relies upon. This allows for extreme flexibility, as resources can be dynamically added or removed according to actual needs.

Although not a significant breakthrough in terms of technology (most of the technologies employed in this model of computing were already available), Cloud computing has revolutionized the way in which technology is being employed. A new business paradigm has emerged where every application or resource is offered as a service, available on-demand, on a pay-per-use basis, over the Internet.





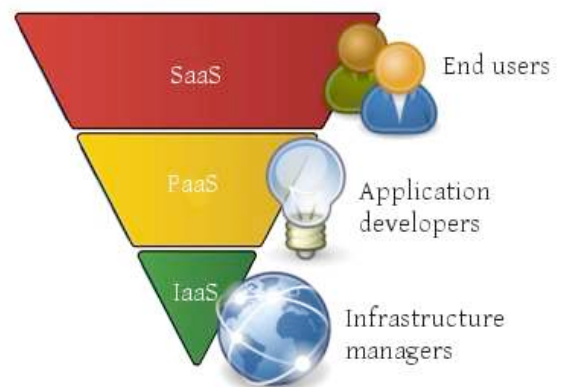
Virtualization, in particular, is a key technology for the implementation of Cloud computing environments. The idea is to pool together different physical resources into a single virtualized environment by means of specific virtualization software (such as Vmware, Xen, etc). The objective is to create a series of logical (virtual) machines with a dynamic set of resources. Virtualization permits a more efficient utilization of available resources. Indeed, thanks to this technology, the computing resources assigned to every virtual machine are not directly related to the underlying physical infrastructure, but are rather assigned dynamically according to the actual needs of the moment. Although a necessary attribute of Cloud computing, virtualization is not sufficient as such. It is the automated and self-provisioning aspect of Cloud computing that distinguishes it

from former technologies in virtualized environments. Human intervention is no longer required in the case of Cloud computing, as resources are able to manage and re-organize themselves according to temporal and contextual contingencies.

Cloud computing is often broken down into three different categories: Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS). In the case of IaaS, the provider offers basic computing resources, such as computer networking, load balancing, data storage, and virtual operating systems. The client can benefit from the use of physical (hardware) resources, without the problems associated with the management thereof, and with the advantage that they can be dynamically resized according to the client's needs. PaaS provides a platform for users to develop software applications. It consists of a series of interactive tools, such as database management and application development, to support the making of powerful and flexible applications that will run in the underlying infrastructure of the Cloud. Finally, SaaS provides end-users with a software solution delivered over the Internet. It aggregates IaaS and PaaS together into a software application, which represents the service that end-users actually interact with.

Even though they qualify as different services, there is a definite overlap between IaaS, PaaS and SaaS, as the latter cannot exist without the former two. It is, however, often the case that these different services are provided by different service providers.

From an operational perspective, the participation of separate but interconnected operators in the provision of one service is the main factor differentiating Cloud computing from former models of service delivery over the Internet. As opposed to previous business models, where one actor was responsible for the provision of one service, in its entirety, to a variety of clients, in the case of Cloud computing, the provision of a service requires the integration of a variety of services (infrastructure, platform, software, etc) provided by a variety of actors. This necessarily requires a complex network of contractual relationships amongst every actor involved in the provision of one integrated service to end-users.



Cloud computing is already widely deployed in the private sector and is, nowadays, also acquiring popularity in the public sector. Since government agencies must operate within a limited budget, Cloud computing can be used to decrease the costs and increase the efficiency of public administration, as well as to promote new services and initiatives that provide additional value to citizens. Given the central role it is starting to play in society, more and more lawyers and researchers are investigating the legal aspects of Cloud computing and debating strategies for the new challenges it engenders.

The Cloud's economic benefits are clear. Use of the Cloud enables both businesses and casual users to maintain as much or as little electronic data as they wish on a third party's mainframes without having to buy or maintain their own hardware systems. However, Clouds can be a legal minefield for companies and their lawyers. Data breaches, hosting of illegal content and inaccessibility of critical business information are just a few examples of difficult situations Cloud users can face.

There currently are a large number of initiatives, events, and international conferences on this topic taking place all over the world,¹ with this paper, we intend to launch the debate in the AI & Law community. We will describe, firstly, a hypothesis of how cloud computing can be described as a complex system, and, secondly, the risks and opportunities connected with the current implementation cloud computing. We will then present a preliminary model for the implementation an automated system of certification based upon the formalization of contractual rules and consumers' preferences.

1.2 Cloud Computing as a Complex System

Complex systems are characterized by (1) a large number of components, (2) partial knowledge of the relationships between them, and (3) limited predictability over the system evolution and dynamics due to the large number of actors involved. In the legal domain, the modeling of complex systems has been employed in a variety of fields (e.g. in France, in order to better understand and represent the network of articles referenced within and across legal codes²). In view of the large number of actors involved in the provision of Cloud applications, the complexity and lack of transparency that characterizes the network of contractual relationships, together with the ubiquitous and transient character of these relationships that must be dynamically updated or modified, the Cloud could ultimately be regarded as a complex system. As a consequence, the costs and complexity of managing and setting up the various aspects of an IT infrastructure (which can be significantly reduced thanks to the deployment of Cloud Computing), have been replaced by a new type of complexity, related to the management and coordination of the complex network of actors involved in the system.

Multiplicity of Actors

The network of contractual relationships in the context of Cloud Computing is becoming increasingly intricate and complex because a large variety of actors are generally involved in the provisions of one service. This creates a situation characterized by considerable **contractual complexity**.

¹ See, for instance, the State of the Art Analysis at <http://crossroad.epu.ntua.gr/>

² MAZZEGA P., BOURCIER D., BOURGINE P., NADAH N., BOULET R., A Complex-System Approach : Legal Knowledge, Ontology, Information and Networks, in *Approaches to Legal Ontologies*, Theories, Domains, Methodologies Series : Law, Governance and Technology Series, Vol. 1 Sartor, G. ; Casanovas, P. ; Biasiotti, M. ; Fernández-Barrera, M. (Eds.) 1st Edition., Springer, Heidelberg 2011, XIII, 279 p, Chap 7

Cloud computing modifies the relationship that subsists between users and service providers, but also amongst service providers themselves. While many operators are in charge of providing infrastructure and platform development, an increasing number of operators are offering services that rely on the infrastructure provided by others. Many Cloud applications involve a large number of actors that provide one or more services overlaid on top of the infrastructures tying them together (vertical integration) or based on the aggregation of services offered by others (horizontal aggregation). The higher is the number of services integrated together, the higher will the value of these services be to the users.

Even though, by exporting their resources in the Cloud, clients are moving away from the complexity of managing and coordinating the technological infrastructure, this complexity is being replaced by the necessity to coordinate the activities of different actors and to manage a complex network of contractual relationships.

Since they are purchasing a service rather than a product, it is important for clients to describe the service that they are willing to purchase. This is achieved by means of Service Level Agreements (SLA) – standard agreements intended to establish a common understanding between the clients and the Cloud provider with regard to the priorities and responsibilities of each party. Given that these agreements stipulate the minimum level of service that must be delivered by the Cloud provider, they constitute the basis upon which clients can build up their expectations in terms of quality of service, infrastructure (uptime, response time, etc), security, privacy, responsibilities and potential liability of the service providers.

The problem is that clients generally enter into a direct contractual relationship only with one actor (i.e. the Cloud broker), but are generally affected by the choices of a large number of actors (i.e. different service providers) whose activities are critical to the provision of the service to which they have subscribed. While the Cloud broker might be aware of the client's expectations, there is no guarantee that the other actors involved will properly understand those expectations and actually fulfill them.

Opacity of the network

One of the main advantages of Cloud Computing is the reduction in costs resulting from an increased flexibility and scalability of resources. This has, however, to be counterbalanced with the higher costs that must be incurred to ensure the quality of the service. As the internal operation of the Cloud is inherently opaque, users inevitably lose control not only over the way in which they can access their own data, but also over the manner in which all data stored in the Cloud can eventually be exploited either by the Cloud provider or by third parties.

Nested Contractual Relationships

Before they enter into a contractual relationship with the Cloud provider, it is extremely important that users properly understand the terms of service. However, end-users are often reluctant to read the terms and conditions of the contract they agreed to because Service Level Agreements are often extremely complex and confusing. In addition, many end-user agreements are likely to change over time without any notice being given to end-users, who have already agreed to be automatically bound to the new terms and conditions.

The problem is further complicated by the fact that users usually enter into a contractual agreement only with the last actor in the supply chain (the Cloud broker) and are thus left without any recourse against the other actors involved in the actual provision of the service, who are not necessarily informed of the terms and conditions of the end-user agreement. Since the internal structure and operations of the Cloud provider or broker are generally not disclosed to the public, it becomes increasingly difficult for users and organizations to understand the actual scope of their contracts, and, in particular, to identify the terms and conditions that are not an integral part of the main contract.

Lack of Transparency

Whenever they move into the Cloud, users or institutions must export their data into the hand of a third party service provider. By doing so, they lose control over the way in which their data is being used, stored and processed by Cloud providers, as well as the manner in which the service will be delivered, as they have no knowledge nor control over the internal operations of the Cloud.

The terms of service can be defined by contractual means, by means of Service Level Agreements between end-users and providers, which have become a key aspect of Cloud Computing. Due to the dynamic nature of the Cloud, ensuring that every provision of the SLA has been properly implemented and is still being respected requires however an active and continuous task of monitoring the Quality of Service (QoS) – and this is especially important in the case of enterprise customers that may outsource critical data. In particular, due to the raising concerns for privacy and data security, consumers may be hesitant to disclose certain details to cloud providers.

Numerous other factors must be taken into consideration in order to assess the reliability and trustworthiness of a Cloud provider. These include, but are not limited to, the type of services provided, the overall accessibility and availability of these services, the formats, standards, and interoperability of the system, but also the respective roles and responsibilities of each party involved. Since different actors are likely to have different preferences and different adversities to risk, every contract must be carefully analyzed and assessed.

The higher is the number of parties, the harder it is to perform a proper assessment. The complex nature of the Cloud is therefore likely to introduce a series of challenges related to the **protection of privacy, the enforcement of intellectual property, the security and confidentiality of data, and, most importantly, the problem of liabilities and responsibilities** involved with the enforcement of various rights and obligations assigned to different actors, either corporate or consumers.

Unpredictability of relationships

Cloud Computing is used to provide flexible solutions that can automatically be adjusted to the changing needs of users. In a dynamic environment, relationships between actors need to be constantly changing or evolving. Clients are increasingly attracted to Cloud solutions because of the lower initial costs it entails, but mainly because of the possibility to pay only for the resources that they effectively use at any given period of time. This is the concept of utility computing, a new model of business whereby computing resources are no longer a product to purchase, but rather a service to subscribe to. This naturally requires a higher degree of elasticity with regard to the infrastructure, services and the different actors involved. Due to the dynamic character of the architecture of the Cloud, and to the temporary or transient character of every contractual relationship it made of, it becomes almost impossible to predict the way in which the Cloud will evolve over time.

Volatility of Actors

Cloud Computing has disrupted the traditional value chain of service provision. A cloud service is delivered by a variety of actors, whose identity can change over time without necessarily changing the nature or the quality of the service.

Even though they appear as infinite to end-users, the amount of resources available in the Cloud are of course limited to the resources provided by the various actors in the Cloud. Perfect elasticity requires the Cloud broker to be able to contract a new service provider whenever the need arises, and resource optimization requires that one service provider be replaced by another whenever the service of the latter is more valuable and/or less costly than that of the former.

As a result of virtualization, actors can come and go in and out of the Cloud in a completely transparent way. The identity of any actor whose role is to provide a particular kind of resources is ultimately irrelevant, provided that the resources it provides are actually interchangeable with each other. Users are not directly affected by the shift from one service provider to another, because most of the resources they provide are simple commodities, which have been gathered together into a virtual infrastructure that is completely independent from the underlying infrastructure.

The advantage is that, given that they are in a contractual relationship only with the last player in the supply chain, changing the identity of service providers does not have any impact on the usability of the system as a whole. Hence, users do not need to be informed of any change that is performed within the internal structure of the Cloud.

Dynamic Revision of Contractual terms

A dynamic revision of contractual provisions is necessary in order to allow for a better re-organization of resources. Given that it has been designed to support unpredictable workloads, the architecture of the Cloud cannot itself be predicted. At any moment, clients' needs might either drop or drastically increase for a very short period of time. Clients might also decide to upgrade their subscription with a particular Cloud provider - in order to benefit from a broader range of services or resources - or even to subscribe to a completely new or different service, perhaps with a new service provider.

Because of the pace at which these revisions happen, terminating and re-creating a new contract each time would prove to be extremely tedious and inefficient. The solution is to integrate within the contract itself the possibility for the client (and sometimes even the service provider) to change the terms and conditions regulating the provision of the service. While this higher degree of flexibility significantly reduces the costs and complexity of contractual negotiations, it however considerably increases the level of complexity in the system, thus making it even harder to predict the manner in which the Cloud environment is likely to evolve in the future.

Transnationality

The widespread deployment of Cloud Computing is likely to have a significant impact on the legal system as a whole, which traditionally relies upon the **concepts of jurisdiction**, national boundaries and territoriality.

Cloud computing services generally extend over several jurisdictions with a large number of data centers globally distributed around the world. In order to ensure a fast and reliable service at minimum costs, data is often replicated in several data centers and may end up distributed across multiple jurisdictions. Cloud computing technologies are designed for data to move around from one data center to the other according to the actual and expected utilization of available computing resources, but also depending upon the current level of congestion of the network. Minimum latency (i.e. the time required to access the data when requested) can be obtained by storing data simultaneously in multiple locations, whereas maximum storage and computing capacity requires a constant flow and transfer of data across different data centers. All these algorithms are unlikely to take national boundaries into account. Although certain service providers allow their clients to specify the country or region in which their data must be stored and/or processed, this is generally the exception rather than the rule, given that the geographical location of data is often difficult to determine ex-ante.

Overall, the issue can be traced to the fact that the fluidity and volatility of data stored in the cloud is in conflict with the more static and deterministic character of the law. National boundaries are irrelevant in the context of Cloud Computing, whose infrastructure exclusively depends on the architecture of the Internet and on the performance of the network. It becomes therefore very difficult to identify the applicable law, and, in the case of litigation, to determine who should ultimately be held liable for what.

2. Risks and Opportunities of Cloud Computing

2.1 Costs

The most obvious advantage of Cloud Computing relates to its costs. Huge economies of scale make it possible for large service providers to offer their services at only a fraction of the costs that their clients would otherwise have to incur in order to set up an analogous infrastructure by themselves. Virtualization allows for a more efficient repartition of resources by separating the logical infrastructure from the technical and hardware architecture. A dynamic configuration of resources based on actual needs promotes a more efficient allocation of resources and reduces the risk of entering into a situation characterized by under/over utilization of resources. From the perspective of providers, this can be extremely valuable because it reduces the sunk costs that have been previously incurred in order to set up their underlying infrastructure. Virtually any resource that is not currently being used by the Cloud provider can be temporally assigned to one of its client. From the perspective of the clients, this can be very convenient because it completely eliminates the initial investment necessary to purchase hardware or software resources and properly setting them up. Most of the fixed costs (in terms facilities, hardware and software resources, technical management and engineering, etc) are fundamentally converted into variable costs.

Yet, the complexity of the Cloud introduces a variety of new costs related to complexity management. On the one hand, as the number of actors involved in the provision of a service increases, contractual negotiations becomes increasingly complex and costly. On the other hand, as the level of control over the infrastructure and the resources decreases, identifying a breach of the contract can become very difficult. The costs of monitoring the operations of the Cloud in order to ensure compliance with the agreed terms and conditions are likely to be very high whenever there is more than one actor involved in the provision of a service. This is even more critical when the identity of the actors responsible for providing the service has not been previously established or is likely to change over time. In an international context, the costs of enforcing contractual provisions can eventually overcome the benefits derived from an increased elasticity and scalability of resources.

2.2 Security

Cloud computing can either improve or reduce the security of a system. While most security mechanisms provided by Cloud providers are likely to be more robust and effective than those set up by end-users, the centralization of data into the hand of a few can make those players more prone to be attacked.³ If Cloud Computing is characterized by the virtualization of a common pool of shared resources, every service provider must have a mechanism to control and access a variety of resources (e.g. storage, processing power, memory, bandwidth) from a centralized interface (the “hypervisor”) in charge of re-organizing and re-allocating these resources according to the specific needs of the moment. To the extent that they are accessible through the Internet and that they provide access to a much larger quantity of data and resources, Cloud-based services constitute a more attractive target for attacks than more traditional servers.

The shift from traditional on-premise storage and operations to Cloud-based solutions can greatly reduce the costs for clients to set-up and secure their own infrastructure, which can generally be done in a more efficient and securely

³ See the ENISA report (2009) on Cloud Computing: Benefits, risks and recommendations for information security, which identifies the main risks of Cloud Computing in terms of information security as being due to loss of governance and user lock-in; isolation failure and compliance risks; management interface compromise; improper data protection; incomplete or insecure data deletion; and malicious insiders.

manner by a professional team of system administrators. However, this reduction in costs is to be compensated by the additional costs to be incurred to ensure that the security mechanisms adopted by every player in the Cloud are actually in line with the security requirements of each individual client. Clients often require their service providers to follow good security practices as an attempt to decrease the risks of attack and to diminish the consequences thereof. Yet, several actors are usually involved in the provision of a Cloud-based service. The greater is the number of actors involved, the higher are the risks that something will eventually go wrong. Besides, most of the service providers that clients communicate with are often unable (or unwilling) to provide everything on their own terms. They frequently aggregate a series of third-party services under a common framework, which - although presented as a single integrated service - is actually made up of a variety of services administered by a variety of actors with their own individual policies and security practices. Regardless of the degree of protection promised by the cloud provider, the security of information is ultimately determined by the weakest link in the chain. Insofar as data is transferred through several intermediaries, only one of them needs to be violated for any malicious user to obtain the relevant information. Hence, the chances for inadvertent exposure increase substantially with every new intermediary and with every new layer of abstraction.

2.3 Privacy and Confidentiality of Information

There is an inherent security risk in the use of the Internet to transfer sensible information and personal data, but that risk has been considerably increased with the deployment of Cloud Computing. The transfer and processing of personal information in the Cloud need to be carefully monitored in order to ensure that the privacy of end-users has not been infringed. The reason is that information stored in the infrastructure of a third party has weaker protection than information that remains in possession of the data subject.

To begin with, the laws of certain countries oblige certain service providers to communicate to the authorities any information that constitutes evidence of criminal activities. This means that government agencies can, under certain circumstances, require the disclosure of personal or confidential information.⁴ The international character of Cloud Computing introduces an additional layer of complexity, given that information stored in the Cloud can be subject to a variety of different laws depending on the location where it is being stored or transmitted. Cloud providers might avail themselves of the services of other Cloud providers located in different jurisdictions, or they might distribute their data amongst multiple data-centers according to economic and/or legal incentives (i.e. forum-shopping). The difficulty for users to know with certainty which law applies to the information published into the Cloud raises a series of critical concerns in terms of privacy and confidentiality of information. Finally, while users generally disclose information voluntarily on the Internet (by means of e.g. through blogs, forums, newsgroups, mailing lists, search engines), problems would arise if the information given to separate (and apparently independent) services were actually aggregated together by one single entity (either because it is the common provider of said services, or because it has acquired the data from third parties). If one single entity were to provide a large variety of services and the data collected through all of these services were to be processed into an integrated framework of analysis, that entity would fundamentally be able to know much more about its user-base than what has been voluntarily disclosed by each individual user. This is problematic because, even though information had been voluntarily provided by users, aggregated data might provide further information about users, which they did not necessarily want to disclose.

⁴ For instance, in the USA, although the Electronic Communications Privacy Act (ECPA) provides a series of protections against the access by governmental agencies to personal information held by third parties (18 U.S.C. § 2510-2522 and § 2701-2712), these protections have been subsequently weakened by the USA PATRIOT Act, which entitles the FBI to compel, following a court order, the disclosure by Cloud providers of any record stored on their servers (50 U.S.C. § 1862).

2.4 Liability and Responsibilities

In front of such a large number of actors and such a diversity of regulations around the world, the traditional role of the law is getting less and less relevant and contractual relationships are assuming an increasingly important role.

Given the complexity of Cloud Computing, particular attention should however be given to the specific rights and obligations assigned to each party to the contractual relationship. The dynamic character of the Cloud is such that any service provider could decide at any given time to out-source part of its infrastructure and operations to third-party providers, without ultimately informing the other parties to the contract. Although the operation is generally not visible to end-users, it might nonetheless affect the quality and reliability of the service as a whole. In order to preclude any responsibility in the eventuality of failure, most of the services provided to end-users (SaaS) are offered under specific Service Level Agreements that stipulate that the service provider cannot be held responsible or liable for the activities performed by third-party contractors.

This raises a series of legal challenges, which have still to be properly addressed. If service providers disclaim any form of liability towards end-users, what kind of recourse is available to users? Do they have a legitimate cause of action against the subcontractors who actually caused the damage, even though they are not in direct contractual relationship with them? If there is no recourse, who should be held responsible for a breach in the system? Who should be held liable for the improper transfer or illegitimate processing of data in the Cloud? Most importantly, if the players involved in the provision of a services have not been previously determined and are likely to change over time, how can users ensure that the level of service will remain the same? What are the legal consequences of any change in control? These questions have thus far not been addressed by the majority of Service Level Agreements. Given the strong asymmetries of information and the difference in bargaining power, not only is it very difficult for users to ensure that the service complies with the terms and conditions of the contractual agreement, but it is even harder to enforce these terms upon every actor involved in the provision of that service.

3. Formalization of Contractual Rules: towards an automated System of Certification?

As Cloud Computing is being adopted by an increasingly larger number of businesses and individuals, the underlying technology and infrastructure is continuously evolving, but the law does not seem to follow the pace. Public regulation (such as intellectual property law, privacy law, and consumer protection law) is being superseded by **private regulation**.⁵ Today, private parties - rather than legislators - are determining the rules of the game. What can or cannot be done is no longer a matter of law, but more a matter of what has been previously agreed upon between a variety of private entities. The problem is that if everything is to be regulated by contracts, the number and the complexity of contractual agreements will constantly keep increasing.

This complex and fluctuating system of contractual relationships requires more sophisticated means of management and enforcement, in order to embrace - rather than resist - the dynamic nature of the system. In this regard, we believe that semantic rules combined with Artificial Intelligence (AI) could reveal themselves useful, not only in

⁵ “as Facebook extracts commercially-valuable information from the aggregation and correlation of millions of users” in Gillian Hadfield, “Legal infrastructure and new economy” *USC CLEO Research paper* n° C10-7, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1567712

order to simplify the work of lawyers in elaborating new contracts, but also in order to counter some of the concerns generated by use of these new technologies by way of technology itself.

3.1 Automated contracts

The size and complexity of contractual relationships in a Cloud environment highlights the need for electronic support in every aspect of contractual activities. More precisely, the formalization of contractual rules can reduce the complexity associated with Service Level Agreements at the level of (1) the negotiation, by simplifying the procedure of identifying a common ground of agreement between each client and the different service providers involved in the provision of a Cloud service; (2) the formation of the contract, by allowing for the drafting of a contract to be performed automatically according to the specific criteria which have been previously agreed upon during the process of negotiation; (3) the performance, by providing a more efficient way of identifying the various rights and obligations assigned to the relevant actors; and (3) the enforcement, by providing a benchmark against which to compare the levels of performance of the services obtained from monitoring.


There is an unlimited range of possible tools that could be deployed to support the formation, performance and enforcement of contractual agreements in a Cloud environment, let us analyze a few.

Contractual Negotiation

The automatic negotiation of SLAs requires that every Cloud provider specifies in advance the terms of service that it will abide to - in terms of the resources provided (i.e. hardware infrastructure, software applications, network bandwidth, etc) and the manner in which these resources will be provided (i.e. costs, up-time, security, privacy, conditions, liabilities, responsibilities, etc); and that users expressly communicate the minimum level of service that they are willing to accept - in terms of the resources they want (e.g. storage, services and applications) and the way they want it (e.g. speed, up-time, security and privacy level, etc).

Through the formalization of the preferences of each party into a language that can be understood by a machine, it becomes possible to implement a mechanism whereby an automated system can autonomously determine whether the service offered by a provider actually complies with the individual needs and requirements each individual users (or other service providers) by merely comparing the formalized terms of the service provider with the formalized preferences of each client. The procedure can be repeated as many times as necessary, according to the number of actors involved in the provision of the Cloud service, and must be reiterated every time new service providers are

incorporated into the Cloud, or whenever they change or update their terms of service.



For the purpose of clarification, we provide an illustration on how formalizing the preferences of the various actors involved can simplify the process of contractual negotiation. Let us take as an example the different services (A, B, and C) offered by different companies. Each offer is characterized by a series of attributes or guarantees that each service provider is willing to provide (e.g. computing resources, uptime, response time, degree of security and respect for privacy) and the various criteria or

conditions at which it is willing to provide them (e.g. costs, liabilities, responsibilities). Offer A is very costly, but it is also extremely secure and is guaranteed to be working 99.9% of the time. It does not, however, guarantee a very high standard of privacy. On the contrary, Offer C is very concerned with the privacy of end-users, but does not however care too much about security. Finally, Offer B is much cheaper than the other two, but - although it is averagely good - the service does not actually excel at anything. Now let us now consider the preferences of user John. John does not need a high level of privacy, but is rather concerned with the availability and security of the system. He is not willing to pay more than \$80 for such a system. Provided that all those terms and conditions have been formalized into a machine-understandable language, John can rely on an automated system in order to identify the offer that best fits its criteria - according to the weight that has been assigned by John to every one of his preferences. In the case under analysis, B is the only offer that actually satisfies the four criteria stipulated by John, and is therefore the one that will be ultimately selected by the system.

The advantage of this approach is that every actor independently declares the minimum level of service that it is willing to provide or accept. The client enters into a contractual relationship only if the service as a whole (in aggregated form) fulfills all of the predetermined criteria. Not only can this significantly reduce the complexity involved in contractual negotiations, but this is also likely to increase the satisfaction of users who no longer have to commit to a standard-form agreement, but can actually obtain a service that specifically complies with the terms of the service to which they have subscribed.

Contract Formation

Once the best offer has been identified, it becomes possible to formulate a contract automatically without further negotiations, since all the relevant elements of the contracts have already been determined by the parties beforehand. A contract is an organized collection of concepts; a collection of rights, obligations, permissions, entitlements, and so on. It is also a collection of procedures that specify the operative aspects of the contract, e.g. how a particular exchange is to be conducted in practice, and a collection of parameters, such as the parties involved, the product of trade, the price of that product, and so on.⁶

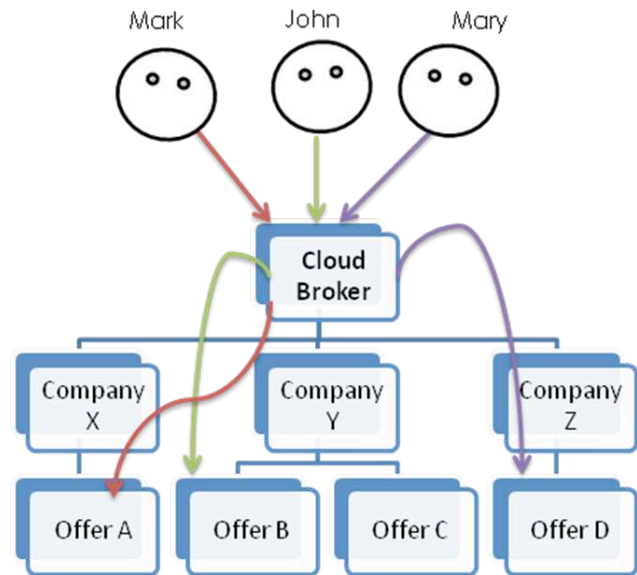
Most importantly, a contract can be regarded as a collection of separate but interrelated sub-agreements. If contractual negotiations were to be guided by a formalized set of rules and constraints, contract formation could be supported by automated tools that understand the ways in which the contract is to be constructed in all of its components and sub-components (and where compliance with the rules and constraints of every part of the contract is a necessary requirement for the coherence of the contract as a whole). Provided that every user's criteria and every provider's condition can be linked to the corresponding contractual provisions it refers to (in the form of a template), once negotiations are over, an automated system could subsequently proceed to the "composition" of the contractual agreement (as opposed to the drafting thereof). This is achieved by gathering together the relevant sections of the contract (i.e. a series of template provisions) and filling them up with the values that represent the common grounds of agreement between every service provider and client.

Performance

The formalization of contractual rules could strongly facilitate the exercise or the performance of contractual rights and obligations within a Cloud environment. Given that every individual user has entered into a different contractual agreement with different service providers, the proper execution of these contracts ultimately depends both on the identity of users and the distinctive characteristics of the service that every service provider has committed to give them. and could support the verification of the extent to which performance actually complies with the contractual provisions.

⁶ These notions have already been studied extensively in legal theory, namely, in the field of Artificial Intelligence, see : A. Daskalopulu & MJ Sergot, The representation of legal Contracts, AI & Society, 11, Nos 1/2, pp. 6-17

In a recent paper, Pankesh Patel, Ajith Ranabahu, and Amit Sheth⁷ propose a mechanism for managing SLAs in a Cloud Computing environment using the Web Service Level Agreement framework, developed for SLA monitoring and SLA enforcement in a Service Oriented Architecture (SOA). The authors suggest that all tasks performed within the Cloud can be defined by logical operators or functions. If this were to be the case, the contractual provisions of every SLA could be formally represented according to a series of logical standards in order to come up with custom logic-based tools capable of understanding and potentially even enforcing these contractual rules.



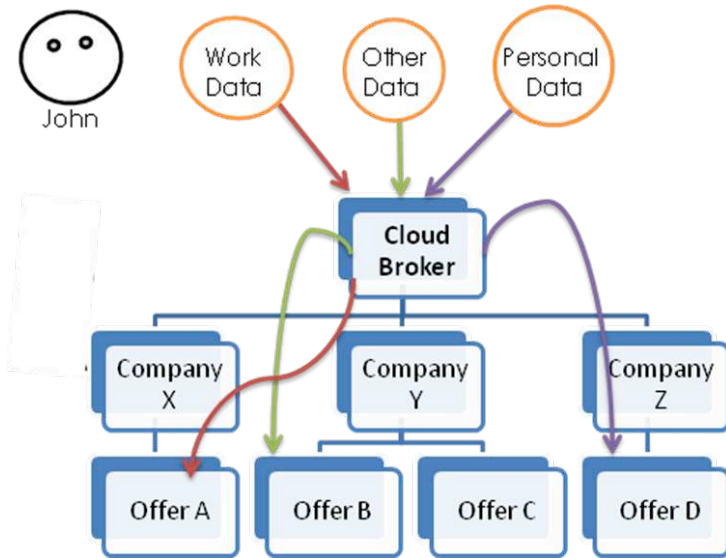
One particularity of the Cloud is that, in general, clients do not enter into a direct contractual relationship with every actor of the supply chain, but only with one particular agent that assumes the function of an intermediary between the clients and different actors involved in the provision of a service. This is the role of the Cloud Broker, who is ultimately in charge of gathering together a large number of services offered by a variety providers and reorganizing them into a single integrated service that is offered to end-users.

The problem is that different clients might have different preferences, criteria, or expectations. Every user request needs therefore be processed by the Cloud Broker before it can be forwarded to the actual service providers. Whether or not the request will be passed on to a particular service provider ultimately depends upon whether or not the service it provides is actually compliant with the terms and conditions incorporated within the SLA of the specific user in question. The same applies at deeper levels of analysis - e.g. if certain service providers decide to outsource part or all of their services to one or more third parties. Users' requests will only be forwarded to those service providers who can guarantee that the service provided by the external contractors is line with each and every user's preferences or requirements. The distinctive characteristics and attributes of each aggregated service (in terms of quality of service, security, privacy, etc) will therefore be ultimately determined by the least valuable or trustworthy of the services it aggregates.

In this respect, the role of the Cloud Broker is to aggregate different service providers under a common framework, while ensuring compliance between each user's criteria and the terms of each service provided. Given that the internal operations of the Cloud are invisible to end-users, the Cloud appears to end-users as one comprehensive service, regardless of the number of actors involved in the actual provision thereof. Who is in charge of providing that service is ultimately irrelevant to end-users, who are only concerned with ensuring that they are actually getting a service that satisfy their criteria. This means that, provided that they all guarantee the minimum standard of service requested by a user, it is theoretically possible for the Cloud broke to shift from one service provider to the other without affecting the interests of end-users, nor infringing any contractual provision. Although this might be a very challenging task, the formalization of user preferences and service specifications into a formal language that can be understood by a machine could drastically reduce the complexity of identifying the routing assigned to each user requests, by allowing for every user's criteria to be assessed against the technical specification of alternative services.

The formalization of contractual rules and user preferences could even go further and extend to data itself. Indeed, it is often the time that one single user has different requirements for different types of data which has to be exported into the Cloud. For instance, while many users are likely to request that their personal data be subject to a higher

⁷ See Pankesh Patel, Ajith Ranabahu, Amit Sheth, "Service Level Agreement in Cloud Computing" Cloud Workshops at OOPSLA09, 2009.



standard of privacy, they might rather give more importance to speed, uptime and security when it comes to the storage or processing data that use on a daily basis. Temporary data of no or little importance could instead be assigned to a different service provider who does not guarantee much privacy or security, but whose cost is much lower than competing services.

By means of metadata, data could be tagged in such a way as to automatically communicate to the system the various locations where it can be stored and the way in which it can be processed, e.g. 'this is personal data that must be treated according to UK law'. With such kind of information, the Cloud broker is able to

determine, without human intervention, how to properly route the data in compliance with a series of criteria. All that the Cloud broker has to know is that the technical specification of the service provider comply the predefined requirements which have been contractually determined by the parties during contractual negotiations. The advantage of encode information directly into the data, rather than into the SLA, is that the conditions becomes inherently linked with the data itself, as opposed to the identity of the users. This allows for data to travel from one Cloud provider to the other without the necessity of entering into a new contract each time.

Contract Enforcement

Not only can the formalization of contractual provisions simplify and eventually enhance the performance of many SLAs, but it can also facilitate the procedure of enforcement. With the tools provided by recent developments in defeasible logic,⁸ it is in fact possible to formalize the specific damages or reparation obligations that must be executed by one party whenever a right has been infringed or an obligation has not been properly fulfilled. This enables all parties to precisely understand the consequences of their acts and the compensation they can expect from the breach of any contractual provision. Whenever a particular event is triggered as a result of an action or non-action by one party, another party will be granted a new right, which generally constitutes an obligation to be fulfilled by the counterpart. With the representation of these rules into a formal and logical language, an automated system can communicate with the interested parties in order to inform them that a new obligation has emerged resulting from the breach or the improper performance of another right or obligation, and, to the extent that it is practically possible, this new obligation can be automatically enforced from within the system.

The problem with said mechanism is that it fundamentally qualifies as a mere mechanism of **auto-certification**, based on the formalization of individual preferences or criteria, on the one hand, and the formalization of the service's provider terms of services on the other. The problem is that, while users have no incentive to lie about their own preferences, there is a strong incentive for service providers to commit to a much higher standard of service than what they are actually able or willing to provide, in particular because there is no way for users to actually find out whether or not their commitment has been properly or entirely fulfilled. Most cloud services are offered as a

⁸ The activation of certain obligations in case of other obligations being violated is referred to as contrary-to-duty obligations (CTD) or reparation obligations (e.g. damages). These obligations are in force only when normative violations occur and are meant to 'repair' violations of primary obligations. See Governatori, Sadiq (2008), The Journey to Business Process Compliance

black-box and provided to users without knowledge or visibility over the operational aspects thereof. Hence, any system of auto-certification ultimately depends upon the credibility and reputation of operators. Even if an operator is genuinely offering a service that purports to comply with certain standards or criteria, users can never be sure that it will actually succeed in fulfilling the prescribed standard of service. While it is always possible to introduce a system of liabilities and compensation in case of failure, there is no way for users to find out whether there has been a breach in certain provisions of the SLA (e.g. whether the proper level of security has been secured or whether the proper standard of privacy has been respected) before the situation gets out of hand.

In spite of the advantages provided by such a mechanism of auto-certification, the system is inherently flawed in that there is no guarantee that the terms of service stipulated by every Cloud provider will be respected, and, most importantly, there is no way to find out whether these providers are actually implementing the policies to which they promised to abide. The lack of transparency that is characteristic of every Cloud environment requires therefore the introduction of a new actor, whose function is to monitor and analyze the internal operations of the Cloud.

3.2 Third party certification

If a system of auto-certification is not able to ensure accurate and transparent disclosure of information, the process of certification must be delegated to a trusted third party.

In this respect, the introduction of a new actor - the Cloud auditor - could further simplify the process of contractual negotiations by decreasing the costs of acquiring information and by reducing the risks of false or inaccurate declarations. Auditing Cloud-based services can however be quite challenging, not only due to the lack of transparency on the part of Cloud providers, but also because services are often deployed across different Cloud providers, each with their own distinctive attributes and characteristics.

In addition to the mechanism of auto-certification, a complementary mechanism of certification could therefore be adopted, whereby each Cloud provider whose services actually satisfy a particular set of requirements would be granted a particular certificate by a third party certification authority. The duty of each certification authority is to investigate the internal operations of the Cloud and to issue a certificate whenever certain criteria are met. Certificates could theoretically refer to any aspect of the Cloud (e.g. the Certificate of Privacy, the Certificate of Security, etc) and could eventually be subdivided in different categories (e.g. level 1: minimum security, level 10: very high security) to precisely convey the range of minimum requirements that every service provider actually complies with.

In order to further facilitate contractual negotiations between service providers and end-users, these certificates could also be encoded into a format that can be understood by a machine, so as to make it possible for service providers to incorporate a certificate directly into their terms of services (i.e. to convey that the service complies with the requirements of that particular certificate), and for users to incorporate it into their own set of criteria (i.e. to convey that they are only willing to subscribe to a service to which that particular certificate has been granted).

To a certain degree, this mechanism of certification could be regarded as a preliminary system of standardization, given that each certificate can be regarded as a “tag” or “label” – acting as a guarantee that a service provider complies with a certain standard of service, regardless of the way in which the service is actually being implemented at the operative or technical level. For instance, to the extent that they all achieve a similar level of security, several service providers could be granted the same Certificate of Security regardless of the technology they use to actually secure their system. To the extent that users are only required to understand what does the certificate implies, rather than having to understand the pro and cons of the underlying technologies used by every service provider, each certificate can be regarded as a short-cut which has the potential of significantly reducing the costs for end-users to select the offer that best suit their needs.

One problem is due to the flexibility of Cloud Computing and the inherent difficulty to predict the way in which the Cloud will evolve over time, since elasticity might require new services or resources to be delivered in real time. Auditing the infrastructure of a Cloud is therefore a process that must be performed on an on-going basis - with the inevitable risk is that a certificate which has already been granted must subsequently be revoked. Certain service providers might no longer comply with the minimum set of requirements that had been previously satisfied, either because they have changed their policy over time, or because they have outsourced their services to other providers which are unable to guarantee the same standard of quality as before.

In that context, the transparency of the certification system and public disclosure of information by the service providers will be an important requirement for traceability. Before any certificate can be issued, the Cloud auditor must ensure that all relevant information necessary to assess the quality of a service has been disclosed and that this information is true. In the case of Cloud Computing, given the inherent opacity of the system, information can be obtained either by mandatory disclosure, i.e. by requiring that all relevant data logs be disclosed to the relevant certification authorities, or by internally monitoring the operations of the Cloud by means of automated software designed to assess compliance with every user's SLA.

In both scenarios, while the goal is to ensure a fair and transparent process of certification, the question is how to make sure that the certification authority will not be tempted to deceive the public in order to increase its profits. The issue arises from the fact that there is a conflict of interest given that the Cloud auditor is providing a service to the public at large, but is actually being remunerated by the service providers which it has been requested to certify.

Third party regulatory control might potentially help avoiding or reducing bias and corruption, although it does not really change the nature of the problem, but merely moves it at a different level. The fundamental question remains as to who is in charge of controlling the controlling authority.

We believe that this problem is however only a marginal one, given that natural market mechanisms might be able to resolve the issue without the need for any governmental intervention. Different certificates could be issued by different certification authorities according to different standards or criteria. Cloud auditors will not be tempted to deceive the public with a distorted system of certification, because their reputation is directly connected to the reliability of the certificates they have issued. The higher the trust of the public in a particular certification scheme, the greater the number of service providers who will request to be certified, and the higher the value of these certificates will be. Assuming that it is possible to preserve competition in the market for certifications, there would be no incentives for any Cloud auditor to provide false or erroneous information, because it would otherwise be immediately taken over by competition. Openness and transparency in the process of certification will instead be rewarded by a higher level of trust from the public. The result is likely to be an increased level of transparency in the private sector - in line with the various Open Data initiatives that are currently emerging in the public sector.

4. Conclusion: Legal and Technical Issues

Although still an evolving paradigm, Cloud Computing has already been extensively deployed in the past few years and is already at the center of attention in many fields of business, industry and academia.

At the technical level, a large number of research projects are exploring and investigating the use of Cloud Computing and ICT for Governance and policy modeling. RESERVOIR, for instance, is an EU FP7 funded project that purports to enable massive scale deployment and management of complex IT services across various administrative domains and governmental services.⁹

⁹ For more details, see www.reservoir-fp7.eu

From a more socio-economic and legal perspective, some of the issues discussed are currently being explored in the French project ADAM on Distributed Architecture and Multiple Multimedia Applications (2010-2013), which is partially being undertaken at the CERSA (CNRS). In the next two years, our interdisciplinary team plans to investigate the specificities of Cloud Computing, the social impact of this new paradigm of business, together with the new legal challenges it engenders. In this paper we launch the debate at the first step of this research. Given the current state of the art of Computers & Law in the context of Cloud Computing, the objective of this paper is to propose a series of ideas that could eventually be implemented into practical solutions as an attempt to address the new legal challenges faced by different actors in the Cloud.