Risk Based SLA Management in Clouds

A legal perspective

Tom Kirkham, Karim Djemame, Mariam Kiran, Ming Jiang, Django Armstrong
University of Leeds
United Kingdom
T.Kirkham@leeds.ac.uk

George Kousiouris, George Vafiadis, Athanasia Evangelinou
National Technical University Athens
Greece
Gkousiou.mail.ntua.gr

***Abstract*— Increasing emphasis in data protection legislation on the provision of user consent and privacy protection poses a challenge to automated Cloud Computing applications.. Here, the infrastructure is removed from the user and therefore no direct link for consent or user control exists. In such environments assurance and best effort from the Cloud Provider to carry out the users wishes and protect his or her privacy must be achieved. This paper introduces such an approach using legally focused risk assessment management of Service Level Agreements between Cloud Service Providers and Cloud Infrastructure Providers.**

Keywords-component; *Cloud Computing, Service Level Aggrements, Legal, Data Privacy, Negotiation*

I. INTRODUCTION

From a legal perspective in a typical Cloud Computing environment there are three main parties. These are, the user (application requestor / data owner), the Service Provider (Cloud Provider (CP)) and the Infrastructure Provider (Virtual Machine (VM) provider). From an Infrastructure provider perspective data management services are supplied by the Cloud Provider to co-ordinate and execute infrastructure in terms of data processing and quality of service.

From a Service Provider perspective the data management is the point at which Cloud Environments are formed and infrastructure is managed in terms of the deployed application. Typically the user's interaction stops at the Cloud Provider but all parties are bound by legal commitments to the user's original request and possible data sharing.

In the implementation described in this paper the Data Manager ([6]) links the Service Provider to the Cloud Infrastructure containing Virtual Machines that are managed and maintained by different sets of vendors. Therefore the Data Manager acts as a gateway to the infrastructure and is therefore the first port of call for communication for requests related to data processing and quality of service within the Cloud Infrastructure.

In order to manage the Cloud infrastructure resources the Data Manager monitors them against specific Service Level Agreements (SLA). The SLA are specific to the Cloud Applications deployed by the Data Manager on the infrastructure and are the first legal point at which the user's rights are expressed. This SLA consists of specific requirements from the Service Provider (adapted from the user request) which can include Data Management specific technical options such as replication rate in the Infrastructure to more user driven requirements like data privacy requirements.

In legal terms the actions of the Data Manager in controlling the infrastructure is done within the confines of the SLA and the wider legal environment, this helps ensure Cloud Application compliance with both the law and specific legal contractual terms linked to the SLA. The live negotiation and management of a deployed SLA in terms of Legal and Data Management is the focus of this paper.

Emerging revisions to the EU data Protection Directive place greater emphasis on user control and consent [1]. In automated Cloud environments this is not always possible and the approach taken here focuses on risk management as a possible way to give greater control in automated Cloud Computing environments. This risk management covers all aspects of the legal and Data Management relationship in terms of Cloud Service deployment and execution management.

II. SLA NEGOTIATION

Live deployment and execution of services on Cloud Infrastructure has to dynamically take into account the requirements of the service and infrastructure provider along with the user. The typical model being that the user relies on the Service Provider to deploy data / services in a Cloud environment according to specific requirements.

These requirements as previously described can be expressed using a SLA. This ensures that the user requirements are bound into a documented agreement that all parties agree upon. Normally the user when setting a SLA with a Service Provider does so in terms of agreeing to terms and conditions and setting preferences which are translated into a typical SLA template the Service Provider uses. For example in other forms of business such as home broadband provision the SLA can be seen in the contract for service that is shared by the customer and provider.

In the Cloud the concept is very similar but as the agreements are managed automatically by the infrastructure they are expressed in a XML standard called WS-Agreement

[2]. Typically the SLA is formed by automated negotiation between the Service Provider and Infrastructure Provider. The Service Provider sends the Infrastructure Provider a proposed SLA and the Infrastructure provider returns this with an offer. The offer can be a modified version of the SLA if the infrastructure provider wishes to make changes; the process is repeated until a final agreement is reached.

Risk Assessments can be used at both ends of the negotiation as the SLA is formed. As each statement in the agreement has to be assessed by each party a risk assessment approach can carry out this assessment. For example the SLA may contain penalties in the event that specific things are not done by either provider. The risk of the penalties occurring in the case of probability and impact is one element of the evaluation on the SLA during negotiation. Other factors in the negotiation could not depend on risk but be apply a more rule based approach depending on specific factors in the implementation.

III. IMPLEMENTATION

The implementation in this paper has used a six stage process for SLA negotiation between a Service Provider and Infrastructure Provider. In each stage one of the six stage process there is a risk assessment.

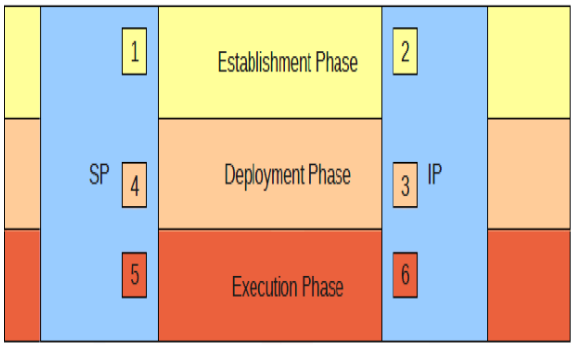

Figure 1. SLA Lifecycle

As Fig 1 illustrates the 6 stages of risk assessment in the SLA lifecycle can be broken down into three phases, establishment phase concerned with the initial SLA offer, deployment phase concerned with SLA deployment and execution phase which monitors service execution with respect to the deployed SLA.

A. Step 1: SP Assessment and IP Selection

Step 1 is the initial risk assessment done on the Service Provider (SP) side regarding data release to the Infrastructure Provider (IP). This risk assessment has to take into account the requirements of the user who has made the initial request to the SP. This request is likely to contain data from the user which the Cloud Service Provider is legally obligated to offer levels of protection to. These levels can be defined by the user in terms of explicit consent for certain forms of processing, or in other cases the permitted actions can also be defined by local data protection legislation.

Thus the legal risk at this stage is based upon the level of compliance the IP has with the user preferences held by the SP, and the impacts (where present) of non compliance from the IP. The risk assessment uses this figure determined by matching the capabilities of the IP with the user request and multiplying this by the impact of non compliance in terms of risk set by the user in terms of requirements. The user specifies impact in the implementation in terms of importance 10 being the highest and 1 being the lowest. The root of the impact is determined by expert opinion and can be quantified in simple terms such as cost.

As Fig 2 illustrates this calculation with i being the number of capabilities not fulfilled (c) and their corresponding impacts (n). The risk is the sum of these calculations.

$$\text{Combined Risk} = \sum_{i=1}^{n} c_1 * i_1$$

Figure 2. Legal Risk Calculation

The compliance data from the IP is retrieved from the Data Manager in the form of a XML document outlining the characteristics of the IPs Cloud Deployment (Figure 2).

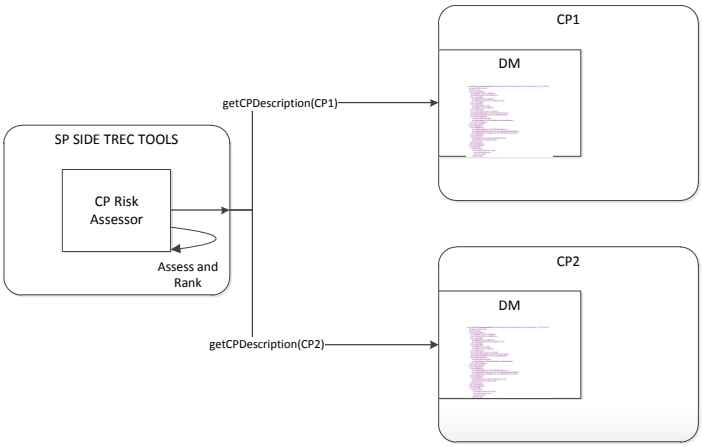

Figure 2: SPRAT Assessment of IP DM

The information in the XML Provider Description contains critical information regarding the setup and management of the IP. This includes aspects of the DM([6]) like availability to configure the replication factor dynamically, the location of the storage resources, the supported security protocols, the legal terms under which the provider handles the data (corporate governance rules applied to the governance of the VMs amongst other factors [7]). The results of the risk calculation are ranked using Euclidean Distance. This calculation takes the user requirements as the target risk by creating a figure of perfect compliance according to all user set preferences. The IP risk scores are then compared against

this to determine ranking, the closest to the rank gets selected in step 1 of the process.

The Euclidean distance calculation can be seen in Fig 3 with i being the number of providers and q representing ideal score with p the actual IP score.

$$\sqrt{\sum_{i=1}^{n}(q_i - p_i)^2}.$$

Figure 3. Ranking Calculation

In terms of legal provision this ranking is important. Not only does it help ensure that the user requirements and therefore consent for specific data processing is abided by, it also illustrates that the SP has provided assurance that these preferences set by the user are respected. The ability to get consent is not possible often in automated Cloud operations such as this and literature often has cited proven assurance as one alternative to this in terms of legal compliance [3,4].

Steps 2, 3 and 4: IP Assessment and SLA deployment

In Step 2 the IP receives the SLA offer and then carries out a risk assessment from the IP perspective on the infrastructure. As previously described this risk assessment checks the capability of the IP to fulfil the requirements sent from the SP. Once Step 2 is complete the IP performs a risk assessment at Step 3. This assessment is done on the SP by the IP to ensure the risks associated of working with the SP from the IP perspective are not too high. In this case the IP will use the SP side provider XML to investigate business orientated risk including Credit Rating and more technical data again including things like legal jurisdiction of the SP.

If the assessment is passed at Step3 the accepted or modified SLA is sent back to the SP and Step 4 commences. This is a re-evaluation of the SLA risk in case the IP has changed it during its risk assessment. If happy the SLA is then implemented on both sides.

B. Steps 5, 6: Execution Phase

Steps 5 and 6 concern service operation from the IP and SP. On both sides risk assessments continually happen to ensure no changes in the SP or IP are in breach of the previously set SLA. This monitoring is done using a risk assessment service that acts as a neutral third party primed with both the requirements from the user / SP and the IP along with the SLA.

Legally data is monitored from the Data Manager to ensure the characteristics of the infrastructure do not pose a too high risk. This is done by a report from the Data Manager that produces a XML view on the current status of the infrastructure. A key area of legal compliance is the location of the VMs as this can often determine what legal jurisdiction they are bound to in terms of local data protection laws. In the case of any change in infrastructure and the provision of extra VMs the location of VMs could change as new VMs are collected by the Data Manager from a new provider. The Risk Assessor should check the location of the VMs to ensure that nay new VMs don't breach the SLA or local (user/ SP) data protection legislative requirements.

During operation, the DM exposes two interfaces for the IPRAT to perform assessments and suggestions. The first one consists of a periodically updated XML report that is used to provide configuration and location information for the storage resources (Figure). This report contains the used replication factor, utilization and location information on the storage nodes. The IPRAT collects this information and calculates the according risk level of the data infrastructure.

The second interface refers to the ability of the IPRAT to suggest actions towards DM for the reduction of the aforementioned risk level. The suggestions include a number of available actions:

- enablement of federation
- addition/removal of a storage resource
- increase/decrease in the number of resources
- increase of replication

The proposed actions may be grouped in two categories, mandatory and voluntary ones. The purpose behind this setup is to give the DM the ability to reject actions that are considered as illegal or less cost efficient, and select the optimal ones. The mandatory actions should of course be followed, however they have to do with internal management decisions. The option to federate for example should be considered in the voluntary set, since a service owner may have explicitly stated in the SLA that they do not wish their data to be externalized under any circumstances.

```
- <service>
  - <id name="progModelService">
      <replication>3</replication>
    - <storagevm>
        <vm_id>34xbe</vm_id>
        <ip>10.8.0.1</ip>
        <location>ESP</location>
        <type>namenode</type>
        <federation>false</federation>
        <cpu>0.12</cpu>
        <memory>0.78</memory>
        <disk>0.82</disk>
      </storagevm>
    - <storagevm>
        <vm_id>35xbe</vm_id>
        <ip>10.8.0.6</ip>
        <location>ESP</location>
        <type>datanode</type>
        <federation>false</federation>
        <cpu>0.72</cpu>
        <memory>0.35</memory>
        <disk>0.25</disk>
      </storagevm>
    - <storagevm>
        <vm_id>36xbe</vm_id>
        <ip>10.8.0.14</ip>
        <location>SWE</location>
        <type>datanode</type>
        <federation>true</federation>
        <cpu>0.45</cpu>
        <memory>0.12</memory>
        <disk>0.68</disk>
      </storagevm>
    </id>
  </service>
```

Figure 4. DM status reporting

IV. IMPLEMENTATION

The steps described above were implemented as parts of research work in the EU Framework 7 project OPTIMIS. The projects main focus being the optimization of Cloud Service use and this includes the provision of optimized management and execution of services.

The implementation architecture involved the management of virtual machines for data storage in a application scenario based on Cloud storage of data. The Data Manager service provided the main interface to the Cloud Infrastructure of Virtual Machines providing distributed files storage (DFS). The main services can be seen in Fig 5.

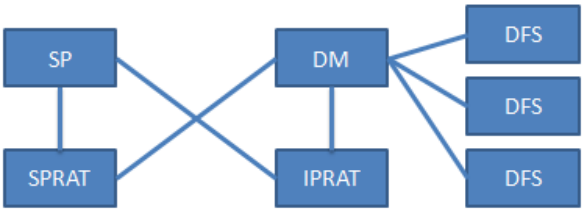

Figure 5. Legal Risk Demo

The Service Provider Risk Assessment Tool (SPRAT) and Infrastructure Provider Risk Assessment Tool (IPRAT) in Fig 5 are core OPTIMIS components. These components perform the previously described legal risk assessments based on data from the Service Provider (SP) and Data Manager (DM).

The deployment was done using Cloud Facilities provided by the University of Umea and ATOS Origin who are OPTIMIS partners. The main flow of the demonstrator was the adaption of the services due to legal requirements in the SLA. In this case the combination of the DM configuration in terms of replication at the IP level and the SLA required a corrective action in terms of DFS setup. This need was picked up by the IPRAT monitoring of the DM during stages 5-6 as described in the previous section.

The resulting action was to increase replication on the DFS or attempt to federate. The federation was discarded as an option due to the fact that the service owner had explicitly stated in the SLA that no federation should be performed. Thus the DM followed the suggestion to increase the replication factor, thus reducing the host failure risk level.

V. RELATED WORK

Risk assessments can be found in workplace health and safety documents to financial institutions lending procedures. The research area of automated risk assessments in distributed computing environments emerged from work in Grid computing. Data standards to describe Cloud capabilities in terms of risk assessments are a new area of research. However, the Cloud Security Alliance is working towards a Cloud Compliance standard in the Governance, Risk Management and Compliance (GRC) group [5]. But as of the time of writing no standard is available to test.

In terms of privacy, research on data security in Clouds can be seen as closely related to other work on data protection in the distributed computing community. These approaches are less with the Cloud environment and more with authentication and authorization around data access. Enabling technology to secure data shared in distributed environments includes Digital Rights Management (DRM) to the concept of sticky policies [8,9].

Work has applied data policy mechanisms such as approaches expressed through languages such as XACML [10] to exiting distributed authentication and authorization technology. A good example of work here focuses on the identity of the accessing party using frameworks for distributed identity such as the Liberty Alliance [11].

I. CONCLUSION

Cloud based data processing relies on automated decisions based on pre defined criteria. Legally this is set in a negotiated SLA between parties. The automated creation and management of this SLA depends on finer management of the services present to ensure the SLA is both effective for all parties and is not breached. In this paper we have illustrated that a Risk Assessment approach can be used to support the SLA management process during Cloud formation, deployment and execution.

II. FUTURE WORK

An area of future work is the development for a supporting infrastructure for the exchange and publication of compliance XML. We are working towards the adoption of a similar model of SSL certificates or federation membership in distributed communities such as implemented by approaches like ZXID [12].

We are also looking to develop the users expression in the application model and are focused on improving anifest representation of user requirements. It is the aim that such requirements that translate through to instances of Virtual Machines can support data privacy policies such as expressed by XACML policies. A natural step for this will also be to incorporate such privacy expression in the WS-Agreement statements of SLA.

In terms of risk we are looking to make the assessments richer in terms of what factors we include in legal risk and how risk is expressed. Currently risk is at the six stages of SLA management and is based on various metrics assessing the IP or SP. Merger of these metrics into unifying values such as cost which uses can better understand is one future approach. Another is to build the risk around business processes. By using a standard such as BPEL it is possible to chain services in applications. Risk Assessments can be applied to this chaining structure to ensure added protection of data when executed.

III. REFERENCES

1. DPD 1995, *Data Protection Directive 95/46/EC of the European Parliament and the Council* (Amended 2003), http://www.dataprotection.ie/viewdoc.asp?docid=89. Access Date September 2012.

2.WS-Aggrement-Negotiation version 1 2011 http://www.gridforum.org/Public_Comment_Docs/Documents/2011-03/WS-Agreement-Negotiation+v1.0.pdf Access Date September 2012.

3. TAS3 Official Deliverable "Privacy Governance and Contractual Options" 2011 http://vds1628.sivit.org/tas3/content/deliverables/TAS3_D6p1_v2_TRequirements:_Privacy,_governance_and_contractual_options.pdf Access Date September 2012.

4. T Kirkham et al, "Assuring Data Privacy in Cloud Transformations" IEEE TRUSTCOM 2012

5. Cloud Security Alliance GRC Stack https://cloudsecurityalliance.org/research/grc-stack/ Access Date September 2012

6. G Kousiouris, G Vafiadis, T Varvarigou. "A Front-end, Hadoop-based Data Management Service for Efficient Federated Clouds." In Proceedings of the 2011 IEEE Third International Conference on Cloud Computing Technology and Science (CLOUDCOM '11). IEEE Computer Society, Washington, DC, USA, 511-516.DOI=10.1109/CloudCom.2011.76

7. B Barnitzke, W Ziegler, G Vafiadis, S Nair, G Kousiouris, M Corrales, O Waldrich, N Forgo and T Varvarigou, "Legal Restraints and Security Requirements on Personal Data and Their Technical Implementation in Clouds", in Proceedings of eChallenges 2011, 26-28 Oct. 2011, Florence, Italy

8. Q. Liu, R. Safavi-Naini, and N. Sheppard. Digital Rights Management for Content Distribution. In Proc. Australasian Information Security Workshop, pages 49–58, 2003.

9. M. Casassa-Mont, S. Pearson and P. Bramhall, "Towards Accountable Management of Identity and Privacy: Sticky Policies and Enforceable Tracing Services", Proc. DEXA 2003, IEEE Computer Society, 2003, pp. 377-382.

10. T. Moses. eXtensible Access Control Markup Language (XACML) version 1.0. Technical report, OASIS, Feb. 2003

11. Liberty Alliance Project homepage www.projectliberty.org Access Date September 2012.

12. ZXID Project homepage www.zxid.org Access Date September 2012.