Introducing Risk Management into Cloud Computing

J. Oriol Fitó and Jordi Guitart
Barcelona Supercomputing Center and Technical University of Catalonia
Barcelona, Spain
{josep.oriol, jordi.guitart}@bsc.es

***Abstract*—The Cloud computing paradigm is offering an innovative and promising vision concerning the Information and Communications Technology (ICT). Actually, it gives the possibility of improving IT systems management and is changing the way in which hardware and software are designed and purchased. Notwithstanding, the use of Cloud resources, which usually are external assets to their consumers, implies risk issues that must be taken into account.**

In this paper, we propose the involvement of risk management procedures into Cloud computing. In this sense, we present a Cloud computing risk management approach aware of Business-Level Objectives (BLOs) of a given Cloud organization. More to the point, we propose an innovatory SEmi-quantitative BLO-driven Cloud Risk Assessment (SEBCRA) as the core subprocess of this Cloud risk management approach.

In addition, we present, as a use case, a Cloud Service Provider (CSP) that is able to improve the achievement of a given BLO, i.e. profit maximization, by managing, assessing, and treating Cloud-related risks. As demonstrated in the experimentation, this provider maximizes its profit by transferring the risks of provisioning its private Cloud, either under- or over-provisioning, to third-party Cloud Infrastructure Providers (CIPs).

***Index Terms*—Cloud computing, Semi-quantitative Risk Management, Business-Level Objectives, Cloud Service Provider**

I. Introduction

Nowadays, Cloud computing is recognized as the most promising computing paradigm of the last several years [1]. It is already a reality and will be a sea change for the Information and Communications Technology (ICT), by modifying the way in which software and hardware are conceived and purchased. This popularity has mainly come due to two of its key capabilities: all computing needs are offered as a service (commonly expressed as *Everything-as-a-Service*) and the ability of dynamically provision computational resources.

Up to now, there are primarily two types of Cloud providers: *Cloud Service Providers (CSP)* or either *SaaS or PaaS providers*, e.g. Google App Engine [2], which offer Cloud services over the Internet; and *Cloud Infrastructure Providers (CIP)* or *IaaS providers*, e.g. Amazon EC2 [3], which provide Cloud infrastructures (typically virtualized execution environments) as a service and, thus, serve as the foundation layer for Cloud systems. Actually, a lot of Cloud computing models have arisen, each one offering different characteristics or services, at different degrees of flexibility and involving distinct risks. Given the fact that Cloud computing encompasses new technologies such as virtualization, there are both new risks to be determined and old risks to be re-evaluated. For these reasons, it is stringently necessary to introduce risk management processes into the whole Cloud computing domain. Generally, treatment of risks in Cloud environments must be performed at service, data, and infrastructure layers.

Even beyond all these considerations, note that day-to-day interactions between Cloud users and providers, as well as between providers themselves, imply several trust and risk issues, which must be addressed by the Cloud community to ensure a successful growing of the paradigm. Actually, Cloud providers and its users will always be exposed to hazard events which can greatly reduce all Cloud computing benefits, unless Cloud-related risks are addressed. The involvement of risk-aware methods into Cloud systems is needed in order to minimize unsung expenditures. Moreover, we also consider the other side of the issue: risks that may result in a benefit or positive impact for Cloud organizations. In this sense, a remarkable tradeoff appears when considering the best action to carry out for each risk.

Going further, and considering a scenario composed by a CSP which interoperates with underlying CIPs in order to consume resources from their public Clouds, risks due to Cloud resources outsourcing (to these third-party CIPs) are significant and they cannot be belittled. Even more, the inclusion of risk management into CSP's operation will lead to an improvement in achieving its Business-Level Objectives (BLOs), such as maximizing both its profit and the energy efficiency of its private Cloud, among many others.

In this work, we contribute to the inclusion of risk management into the Cloud computing paradigm. In particular, we propose a Cloud risk management process led by BLOs, and a SEmi-quantitative BLO-driven Cloud Risk Assessment (SEBCRA) as its core subprocess. All risk-aware actions within these processes are oriented to address the impact of Cloud-specific risks on these BLOs. Basically, they allow any Cloud organization to be aware of Cloud risks and align its low-level management decisions according to high-level (business) objectives. Furthermore, we demonstrate, as a use case, that a Cloud provider (i.e. CSP) is able to improve the achievement of a significant variety of BLOs, by managing and assessing those Cloud-related risks.

II. Background

A. Risk Management and Assessment

According to the risk management standard of the Institute of Risk Management (IRM) [4], a *risk* can be defined as

the combination of the probability of an event and its consequences (whether positive or negative) (based on ISO/IEC Guide 73 [5]). In general, in all types of businesses there are events which represent opportunities for benefit or threats to success, i.e. positive and negative aspects of risks, respectively. Thus, and in contrast to traditional risk avoidance strategies, accepting some risks leads to obtain remarkable benefits.

The *Risk Management*, governed by principles and generic guidelines established in ISO 31000:2009 [6], is the process whereby organizations treat, in a methodical way, risks related with their activities. The main goal is to obtain benefits and sustainable values within each activity and across all of them. Actually, it is a fundamental part of any organization's strategic management.

Entering in detail in its core subprocess, i.e. risk assessment, there are three primary methods according to [7]: *qualitative*, which uses simple calculations and thus it is not needed to determine the numerical value of all assets at risk and threat frequencies; *quantitative*, which assigns numerical values to both impact and likelihood of risks; *semi-quantitative (or hybrid)*, which is less numerically intensive than quantitative methods and classifies (prioritizes) risks according to consequences and foreseen probabilities. Particularly, quantitative risk assessments have been criticized for being overly reductive and divert attention from preventive actions. In addition, they ignore important qualitative differences among risks. Although calculations involved are tedious and include a strong element of arbitrariness, their main advantage is that they provide accurate measurements of impacts' magnitude. However, these quantitative impacts may be unclear, thus requiring to be interpreted in a qualitative way. Contrariwise, the main advantage of a qualitative assessment is that it prioritizes risks and identifies the most important areas for improvement. Even so, it does not provide enough quantifiable measurements concerning probabilities and impacts of risks. As a result, semi-quantitative methods basically take profit of both aforesaid advantages and, therefore, provide risk prioritizations and useful quantifiable impacts analysis.

B. Cloud Computing Risks

Together with the surging of the Cloud computing paradigm, several new risks have appeared. Within these new risks we find specific issues imposed by law or regulations and a lot of operational risks imposed due to using an external provider or service. The Cloud community has to clearly identify Cloud-specific risks and re-evaluate conventional ones. Further, Cloud services, and not just providers, should be the subject of risk management and assessment. Risks in Cloud environments must be considered at service, data, and infrastructure layers. Notice that the level of risk will, in many cases, vary significantly with the type of Cloud architecture being considered. Moreover, it is possible for Cloud customers to transfer some risks to external Cloud providers. In any case, we believe the operation of Cloud services should be consistent with both risk management strategies and BLOs.

According to [8], [9] and [10], the most important risks introduced by Cloud computing are: SLAs breaches, ability to adequately assess risks of a Cloud provider, responsibility to protect sensitive data, virtualization-related risks, loss of direct control of resources and software, compliance risks, and decreased reliability since service providers may go out of business, among others. On the contrary, there are some traditional risks that must be re-evaluated. For instance, the risk of network breaks is now more critical for Cloud organizations since they are totally based on the network. Furthermore, other risks, such as natural disasters, must be considered in a different way (because of the constantly use of external resources) for ensuring high-availability of Cloud services.

In short, there is appearing a wide range of Cloud-specific risks that must be well-identified and managed by the Cloud community in the coming years. Hence, we need to include risk management processes into Clouds and develop risk-aware strategies, policies and heuristics required to face both those Cloud risks and traditional ones that have changed.

III. Related Work

The risk itself and its management have been considered in a great amount of research fields, e.g. statistics, biology, engineering and systems analysis, since many years ago.

Aven has made pioneer contributions for risk analysis and management, such as [11]. Additionally, he has introduced the concept of considering a risk as an event where the result is uncertain, either positive and negative, [12]; and has stated that semi-quantitative risk analysis replaces very well tedious quantitative approaches [13].

Somewhat similar to our work, Yeo and Buyya [14] pose the problem whether a resource management policy implemented in the commercial computing service is capable of meeting the required objectives or not. For that purpose, they develop two evaluation methods to validate the effectiveness of resource management policies in achieving the required objectives: separate and integrated risk analysis. The experimentation performed demonstrates the applicability and success of their risk-based methods. Furthermore, in [15] and [16] the risk of paying penalties to compensate service providers' users is minimized and, thus, they are able to increase the profit of those providers.

Moreover, there are many research contributions toward risk assessment that differ in the level of development of methodology items. They are all grouped under the three risk assessment methods: quantitative [17], which most of them are based on Bayes' theorem [18]; qualitative [19], and semi-quantitative [13].

In addition, the risk has been widely tackled in other successful distributed computing paradigm, i.e. Grid computing. In this field, Djemame et al. have contributed with a lot of research works (see for instance [20]) within the context of the AssessGrid project [21]. As we propose in this work focused on the Cloud, they have widely addressed the inclusion and implementation of risk management and assessment methods into Grid environments. In particular, they treat risk-aware negotiations and SLAs, risk-based decision-support for infrastructure management, and calculation of risk-indicators for provider ranking and competition, among others.

However, and as stated in the above Section II-B, risk management and assessment methods are yet to be introduced into the Cloud computing field of study and there are a lot of open research issues.

IV. BLO-driven Cloud Risk Management

In this section, we introduce a novel Cloud risk management process driven by organization's interests (i.e. BLOs). In essence, it is designed to address impacts and consequences of Cloud-specific risks into well-defined BLOs of a given Cloud organization. In fact, it has the main goal of increasing the probability of success and, thus, decreasing both the chance to failure and the uncertainty in achieving organization's BLOs. In this direction, Cloud organization's core operations will be dynamically adapted by means of risk-aware scheduling and policies.

As illustrated in Figure 1, our risk management proposal is governed by organization's BLOs and strategic objectives, and is split in the following processes (based on the FERMA's Risk Management Standard [22]): *SEBCRA (Risk Assessment)*, which is basically the overall process of risk analysis and evaluation (see next Section IV-A for a detailed explanation); *Risk Reporting* and communication; *Risk Treatment*, which implements and selects risk-aware policies, as well as measures, actions, and controls to amend consequences of risks; and *Risk Monitoring* where all the above steps are reviewed.

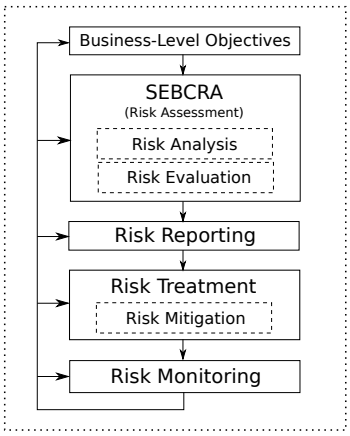

Fig. 1. BLO-driven Cloud risk management steps.

A. SEmi-quantitative BLO-oriented Cloud Risk Assessment

Risk management literature commonly specifies the need to rank and prioritize risks in order to identify areas for immediate improvement and, thus, focus the best efforts on threatening risks. In this sense, we present a new information security risk assessment model, i.e. a SEmi-quantitative BLO-oriented Cloud Risk Assessment (SEBCRA), which has the known purpose of generating a ranking of Cloud risks. Moreover, the main difference with other risk assessment models is that it evaluates the impact of Cloud-related risks on BLOs considered, instead of considering effects on the whole Cloud organization.

In fact, it is the core process of the BLO-driven Cloud risk management and has Risk Level Estimations (RLEs) as outputs, which are individually specified for each risk and BLO affected. Generally, the whole assessment method is subdivided into the risk analysis and its evaluation.

1) Risk Analysis: This is the step in which the probability of risks and the magnitude of their consequences are determined. We propose a semi-quantitative risk analysis, which uses a standard risk level matrix in order to bring out risk level estimations. Note that we have based these estimations on ISO/IEC 27005:2008 [23]. We can divide risk analysis in three stages: *Risk identification*, which establishes and defines organization's potential risks. It is perhaps the most difficult aspect of managing risks, because more risks will happen in the future than can be predicted today, i.e. the problem of "unk-unks" (unknown unknowns) [24]; *Risk description*, that is an essential step to guarantee a comprehensive risk assessment method. Its objective is to show the previously identified risks in a structured format; and *Risk estimation*, which figures out the likelihood of occurrence and the estimated impact on BLOs of each risk previously recognized. In particular, the impact is considered in terms of threats (downside risks) and opportunities (upside risks) and is usually evaluated using a 3x3, 4x4 or 5x5 risk-level matrices, depending on the granularity of risk assessment desired. Notice that we use a 10x5 matrix because we are considering five possibilities either for positive and negative impacts, while standard matrices (e.g. 5x5) only consider the negative side.

Going into detail, we establish the following semi-quantitative classifications: the *probability* of occurrence of risk (P_i), expressed by means of: very unlikely - 0.1 (e.g. once in 1000 years), unlikely - 0.25 (1 in 10 years), possible - 0.5 (yearly), likely - 0.75 (monthly or weekly), and frequent - 1.0 (e.g. at any moment); the *impact* of risk (I_i), either a threat, a benefit, or both, semi-quantified between very high (-100 or 100, for negative and positive impact, respectively), high (-75/75), medium (-50/50), low (-25/25) and very low (-10/10); the *BLO affected* by risk (B_i); and the *Risk Level Estimation* (RLE), which is proportional to the *probability* of a given risk and its *impact* on the *BLO* in question ($I_i(B_i)$), resulting in the following equation:

$$RLE_i(B_i) = P_i \cdot I_i(B_i)$$

Notice that five levels of RLE are defined: *critical* if $-100 \leq RLE \leq -50$; *unacceptable* if $-50 < RLE < -10$; *negligible* if $-10 \leq RLE \leq 10$; *profitable* if $10 < RLE < 50$; and *high profitable* if $50 \leq RLE \leq 100$. Therefore, we have to avoid risks with a RLE that is within the critical or unacceptable range, and take advantage of risks that lead to an improvement in the achievement of the BLOs considered. For a better understanding, in Table I and Figure 2 we illustrate all the possibilities concerning risk level estimations for a given BLO in terms of ranges and numeric values, respectively. Note that values given to P_i and $I_i(B_i)$ are useful for assessing resulting $RLE_i(B_i)$.

2) Risk Evaluation: In this step, it is needed to compare risks levels estimated against a risk acceptance criterion, which establish a threshold that determines the acceptability of risks.

			Probability P_i				
			Very unlikely (0.1)	Unlikely (0.25)	Possible (0.5)	Likely (0.75)	Frequent (1.0)
Impact $I_i(B_i)$	Benefit	Very high (100)	Negligible	Profitable	High profitable	High profitable	High profitable
		High (75)	Negligible	Profitable	Profitable	High profitable	High profitable
		Medium (50)	Negligible	Profitable	Profitable	Profitable	High profitable
		Low (25)	Negligible	Negligible	Profitable	Profitable	Profitable
		Very low (10)	Negligible	Negligible	Negligible	Negligible	Negligible
	Threat	Very low (-10)	Negligible	Negligible	Negligible	Negligible	Negligible
		Low (-25)	Negligible	Negligible	Unacceptable	Unacceptable	Unacceptable
		Medium (-50)	Negligible	Unacceptable	Unacceptable	Unacceptable	Critical
		High (-75)	Negligible	Unacceptable	Unacceptable	Critical	Critical
		Very high (-100)	Negligible	Unacceptable	Critical	Critical	Critical

TABLE I
RISK-LEVEL MATRIX (IN QUALITATIVE RANGES) OF SEBCRA ON BLOS.

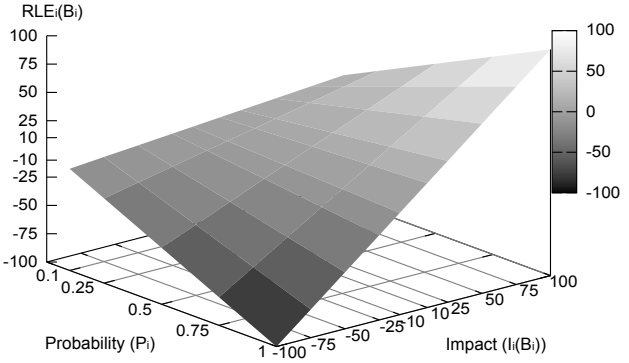

Fig. 2. Risk level estimations (in numeric values) of SEBCRA on BLOs.

In fact, the action of moving up and down this threshold (i.e. risk acceptance criterion) is a serviceable practice to determine the organization's risking level. Really, this step determines risk management priorities according to semi-quantitative relationships between risks and their impacts.

B. Risk Reporting

Thereafter the SEBCRA method, the risk management approach has to report, in a comprehensive manner for Cloud organization's executives, all results obtained in the assessment subprocess. Its main goal is to bring risk analysis and evaluation back to the business core. Afterward, this information is useful for rethinking high-level objectives according to the constantly improved knowledge around risks.

C. Risk Treatment and Mitigation

Once risks has been identified, evaluated, and reported, the *Risk Treatment* subprocess takes place. This involves the definition of potential risk-aware actions, controls, and policies to conduct an appropriate *Risk Mitigation* methodology, which aims to move risks on the negligible or profitable levels. In this sense, there are four possible responses to effectively deal with each risk. *Avoid* the risk, by eliminating its cause(s). *Reduce* the risk by taking steps to cut down its probability, its impact, or both. *Accept* the risk and its related consequences. *Transfer* or delegate the risk to external organizations.

Obviously, risks with a positive risk level will always be 'accepted'. On the contrary, for negative risks, the organization has the other three alternative responses. Future risks' impacts will depend vastly on these decisions taken.

D. Risk Monitoring

Finally, *Risk Level Estimations predicted (RLE_p)*, that the organization expects to have after carrying out the risk response decided in the previous step, are also reported. Differences between RLE and RLE_p are the motivation of performing the whole risk management process and both are inputs for the *Risk Monitoring* subprocess. This is the last part of the whole management process and helps the organization to know if the actions previously decided are correct or not in managing each risk. Its main goal is to adapt any of the previous subprocesses according to information gathered from monitoring methods. In fact, monitoring information is very important for a risk-aware self-management of Cloud services driven by business-level objectives.

V. USE CASE: RISK MANAGEMENT IN A CLOUD SERVICE PROVIDER

In this section, we present a Cloud Service Provider (CSP) as a use case of the risk management procedure explained above in Section IV.

A. CSP in Hybrid Cloud Scenario

We consider a scenario in which a CSP interoperates with underlying CIPs. In this context, the CSP is the responsible for operating, maintaining, and managing its *private Cloud* (which is composed by its in-house virtualized resources) and is able to outsource Cloud resources to *public Clouds* managed by third-party CIPs. This outsourcing operation takes place when the capacity of those in-house resources (i.e. the private Cloud) of the CSP, are insufficient for Cloud services' time-varying demands. This environment composed by both public and private Clouds is known as *hybrid Cloud*.

B. Risks and BLOs of the CSP

Firstly, we identify the most important risks to which the CSP is exposed in the scenario considered.In particular, we classify these risks into the following groups, depending on the source of them:

1) *Cloud capacity provisioning*. We distinguish between two types of risk, which are over- and under-provisioning a given Cloud. Regarding the CSP, this risk affects assets constituting its own private Cloud (see Section V-E for further information).
2) *Service Level Agreements*, such as the risks of accepting new Cloud service's SLA, its violations due to poor performance and service disruptions, etc.
3) *Virtualization*, i.e. those related with the underlying technology of Clouds, such as the risks of virtual machine isolation or virtualization performance overhead.
4) *Cloud applications data*. In this case we consider the risks of data integrity loss, destruction of data, etc.
5) *Cloud resources outsourcing*. Risks associated with the loss of governance and hidden costs, among others.
6) *Others*. Here we group the risks of power loss to the IT systems, natural disasters, fire, etc.

Secondly, since the SEBCRA procedure is oriented to the impact of risks on BLOs, we specify the CSP's BLOs considered at this moment: profit maximization (*ProfMax*), Quality of Service maximization (*QoSMax*), energy efficiency maximization (*EnEffMax*), maximization of customers' satisfaction (*SatMax*), hazard events (threat risks) minimization (*HazMin*), trust maximization (*TrustMax*), reliability maximization (*RelMax*), and reputation maximization (*RepMax*). Evidently, some other high-level objectives can be determined.

C. SEBCRA for Risk Assessment in a CSP

We want to demonstrate the feasibility of the SEBCRA procedure to be used in a CSP. For this reason, we present an example showing how different risks have distinct impacts on the BLOs considered. For instance, the risks concerning the provisioning of the CSP's private Cloud have the impacts on BLOs as described in Table II.

Risk i	B_i	$I_i(B_i)$		$RLE_i(B_i)$
		Benefit	Threat	
Over-prov.	HazMin	0	Very high	Critical
	EnEffMax	0	Very high	Critical
	ProfMax	0	Medium	Critical
	RelMax	Very Low	0	Negligible
	RepMax	Very Low	0	Negligible
	TrustMax	Very Low	0	Negligible
	QoSMax	Very Low	0	Negligible
	SatMax	Very Low	0	Negligible
Under-prov.	ProfMax	0	Very high	Critical
	HazMin	0	Very high	Critical
	RelMax	0	High	Critical
	RepMax	0	High	Critical
	TrustMax	0	High	Critical
	QoSMax	0	High	Critical
	SatMax	0	High	Critical
	EnEffMax	Very low	0	Negligible

TABLE II
THE IMPACTS ON CSP'S BLOS OF THE RISKS OF PROVISIONING ITS PRIVATE CLOUD.

On one hand, the risk of over-provisioning can appear at any moment (probability of occurrence = frequent) and its risk levels estimations are: *critical* for *HazMin, EnEffMax and ProfMax*, because the exposure to hazard events increases, the provider is consuming more energy than the strictly needed and it pays for more resources than necessary, respectively; and *negligible* for *RelMax, RepMax, TrustMax, QoSMax and SatMax*, because this risk has almost no impact on these BLOs.

On the other hand, the risk of under-provisioning have the same probability (frequent), but dissimilar impacts and BLOs affected: *critical* for *ProfMax, HazMin, RelMax, RepMax, TrustMax, QoSMax and SatMax* because the provider is not able to meet with the QoS agreed in the SLA and thus, clients' satisfaction, QoS offered, and its reliability, reputation, trust and total gain are clearly diminished; and *negligible* for *EnEffMax* because, although the energy consumption is many times less than the required by Cloud applications, the under-provisioning technique does not incur significant improvements in terms of energy efficiency.

D. SEBCRA for Profit Maximization in a CSP

In this section, we exemplify how the SEBCRA procedure can be used to improve a given BLO, which is the profit maximization of the provider in question. Notice that all risks can have impact on many BLOs (as shown in Table II), but in this case (Table III) we only present the consequences on the *ProfMax* BLO. Moreover, all the key parameters of the SEBCRA procedure are also identified in Table III. After completing the table of probabilities and impacts, the SEBCRA procedure helps the CSP in the tasks of categorizing and prioritizing risks according to their importance to a given BLO. In this sense, the CSP is able to put its best efforts for addressing risks that may incur more benefit to the *ProfMax* BLO in this case: the risks concerning the provisioning of its private Cloud. Indeed, this SEBCRA procedure focused on the profit maximization indicates that the CSP will be able to move the RLEs of these risks from the 'critical' range to the 'high profitable', by transferring them to third-party CIPs.

It is noteworthy that this transference of risks is carried out by outsourcing Cloud resources to public Clouds owned by CIPs. In fact, these outsourcing operations are very suitable in this scenario in order to maximize the total profit of the CSP, but also may have important side effects that cannot be overlooked. For instance, the probability of hazard events during these outsourcing operations may be increased (for an example see Section V-F).

As a result, this innovative SEBCRA procedure is very convenient to be used by the CSP driven by BLOs. Basically, the provider in question will be able to better align its BLOs with the implemented resource management and policies aware of risk probabilities, impacts, and level estimations. Risk management strategies could be largely used to deal with critical and unacceptable levels of risk. For instance, risk avoidance for rejecting a new Cloud service and risk reduction, by executing redundantly an application on different Cloud resources, for minimizing the negative impact due to SLA violations, service disruptions, and performance losses.

Going further, and depending on the priority of the provider on each BLO, actions to be performed in order to address risks will be different and, thus, the impact of those (either a benefit and/or threat) will also be dissimilar.

Risk$_i$	Probability (P$_i$)	Impact		RLE$_i$(B$_i$)	Action(s)	Consequences		RLE$_p$(B$_i$)
		Benefit	Threat			Benefit	Threat	
Under-provisioning	Frequent	0	Very High	Critical	Transfer	Very high	0	High profit.
Over-provisioning	Frequent	0	Medium	Critical	Transfer	Very high	0	High profit.
Accept new Cloud Service	Likely	High	Low	Profitable	Accept / Avoid	Very high	0	High profit.
SLA violations	Likely	Very low	Medium	Unacceptable	Reduce / Avoid	0	Very low	Negligible
Service disruptions	Possible	0	Medium	Unacceptable	Reduce / Avoid	0	Very low	Negligible
Performance loss	Possible	0	Medium	Unacceptable	Reduce / Avoid	0	Very low	Negligible
Outsourcing hidden costs	Unlikely	0	High	Unacceptable	Avoid	0	Very low	Negligible
VM isolation	Very unlikely	0	High	Negligible	Accept	0	High	Negligible
Virt. performance overhead	Very unlikely	0	High	Negligible	Accept	0	High	Negligible
Data integrity loss	Very unlikely	0	Medium	Negligible	Accept	0	Medium	Negligible
Destruction of data	Very unlikely	0	Medium	Negligible	Accept	0	Medium	Negligible
Loss of governance	Very unlikely	0	Medium	Negligible	Accept	0	Medium	Negligible
Power loss of IT systems	Very unlikely	0	Very low	Negligible	Accept	0	Very low	Negligible
Natural disasters, fire, etc.	Very unlikely	0	Very low	Negligible	Accept	0	Very low	Negligible

TABLE III
USE OF SEBCRA PROCEDURE FOR PROFIT MAXIMIZATION (PROFMAX) OF THE CLOUD SERVICE PROVIDER.

E. Transferring the Risks of Private Cloud Provisioning to CIPs

As shown in Figure 3, an over-provisioning strategy implies that servers are underutilized in low demand situations, with the corresponding energetic, economic, and administration expenditures. On the other side, an under-provisioned datacenter, the provider will not pay so much for these costs. However, it will lose part of clients as it is not able to attend peak demands. In addition, although one thinks that has all the needed resources capacity for the estimated peak demand, a wrong estimation can be made or some unexpected demands due to sudden events could appear (e.g. slashdot effect).

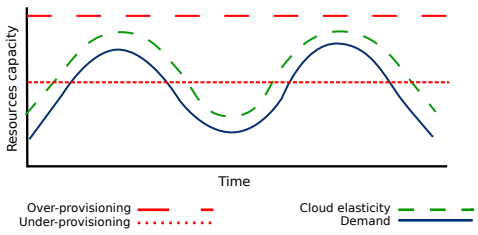

Fig. 3. Cloud provisioning strategies.

The afore-mentioned outsourcing operation to public Clouds allows the CSP to transfer these critical risks, i.e. the risks of provisioning the private Cloud of the CSP, to third-party CIPs. Actually, the outsourcing operation is performed implicitly by the Cloud elasticity method and takes place when a CSP needs to scale up the Cloud infrastructure. It is carried out when Cloud services' demands overcome resources capacity of the private Cloud managed by the CSP itself.

As a matter of fact, the CSP is able to obtain remarkable benefits by transferring the risks of provisioning its private Cloud, as well as using the Cloud elasticity capability and the Cloud resources outsourcing. Within economic benefits we observe the direct consequence of maximizing its profit. Moreover we can highlight, for instance, the maximization of customers' satisfaction and CSP's reputation. Generally, the provider does not have to necessarily make a great investment for its private datacenter or do any capacity planning.

F. Risk of Service Disruptions

The risk of service disruptions increases implicitly when performing outsourcing operations to external CIPs. Heretofore, Cloud providers, either of infrastructures (resources) or services, exhibit a lack of proper and suitable Service Level Agreements (SLAs) or, simply, only low-level parameters are included in them, such as the availability of the corresponding Cloud service. For instance, Amazon Web Services, the most successful CIP of nowadays, only provides simple guarantees in the SLAs for its offered services, that are Amazon Elastic Compute Cloud (Amazon EC2) [25] and Amazon Simple Storage Service (Amazon S3) [26]. Particularly, only the availability (expressed as the annual uptime percentage) of these Cloud services is defined in those SLAs: 99.95% and 99.9% for EC2 and S3 services, respectively. Notice that Amazon pays a penalty equal to 10% of the monthly bill for the EC2 service if the availability is less than the agreed in the SLA, and 10% or 25% of the monthly bill for the S3 service if the availability is equal or greater than 99% but less than 99.9%, or if it is less than 99%, respectively.

Using these SLAs, any Cloud provider is able to manage the risk of service disruptions and consider its acceptability by using both the availability parameter and associated penalties determined in SLAs. Nevertheless, and due to the absence of complete SLAs, the CSP is not still able to cope with many other risks involved in the hybrid scenario considered (see Section V-B for a description of them).

Actually, we are currently working on the integration of the CSP presented herein with different CIPs. The future combination of both this integration with multiple CIPs and suitable SLAs definition from those CIPs will allow the CSP to use the risk management process in order to implement policies that assess in the decision of which is the best CIP to outsource for each Cloud service offered. Probably, the 'best' CIP will be simply the one with lower risk or another one with an assumable risk and offering better other significant parameters like ecological efficiency, prices, reputation, loyalty, and security. Notice that these additional parameters must be considered throughout the risk assessment methodology. In addition, historical information about previous agreements carried out with external CIPs will be very useful.

VI. Experimentation

In this section, we present the experimentation which demonstrates the benefits achieved by the CSP by incorporating the BLO-driven risk management into its operation.

A. Experimental Environment

We use Apache Tomcat v5.5 with the hybrid architecture [27] as the back-end web servers of the CSP. All of them are encapsulated into virtualized and isolated Cloud resources with one processor unit and 512MB of memory available for the server. Moreover, we use Squid [28] as the proxy server for load balancing among the available web servers (either local or outsourced), and EMOTIVE [29] as the third-party CIP. We have deployed the SPECweb2009 [30] banking web application. It is based on Internet personal banking and, consequently, all clients' requests are through the SSL protocol. All the machines used are connected through 1 Gbps Ethernet and run Xen 3.3.1 over Linux kernel 2.6.18.

1) Workload: The workload pattern was obtained from a European ISP (its name cannot be disclosed) and is the typical received by current web applications during a whole day. It has been produced using Httperf [31] (instead SPECweb2009 client emulator) because it allows to make more configurable and variable tests. Workload requests generated by Httperf were extracted from an exhaustive characterization of the SPECweb2009 client emulator. All tests lasted one day.

2) SLA economic parameters: For this experimentation, we have used the following economic parameters, which are specified in SLAs agreed with clients: a *Price* of 1€ per hour, which is the amount of money that the clients pays to the CSP for deploying and managing the associated web application; *Cost* of Cloud resources (with 1 processor unit and 512MB of memory) of 0.18€/h and 0.15€/h for in-house and outsourced resources, respectively; and changeable *Penalties* for the provider. Note that these penalties depend on two parameters that determine the degree of SLA violations, namely the total amount of time with an SLA violation, and the magnitude of this violation (see [32] for further information).

B. CSP Profit Maximization

Figure 4 shows, from top to bottom, the variable input load pattern, the number of back-end web servers, and the instantaneous profit earned by three different CSPs: *risk-aware*, which uses the SEBCRA procedure and, thus, considers the transference to a third-party CIP (EMOTIVE) of the risks of provisioning its private Cloud; and the two possible cases without using any risk assessment procedure, i.e. *under-provisioned* and *over-provisioned* private Cloud.

The final profits in these three cases are the following: 21.84€/day for the risk-aware, 8.0304€/day for the under-provisioned strategy and 19.68€/day for the over-provisioned. Notice that the loss in earnings in the over-provisioning case is due to the fact of paying at all time the maximum amount of Cloud resources needed for attending the highest peak demand (three in this case). In the under-provisioning case, the provider does not pay for so many resources (in fact, only for one), but

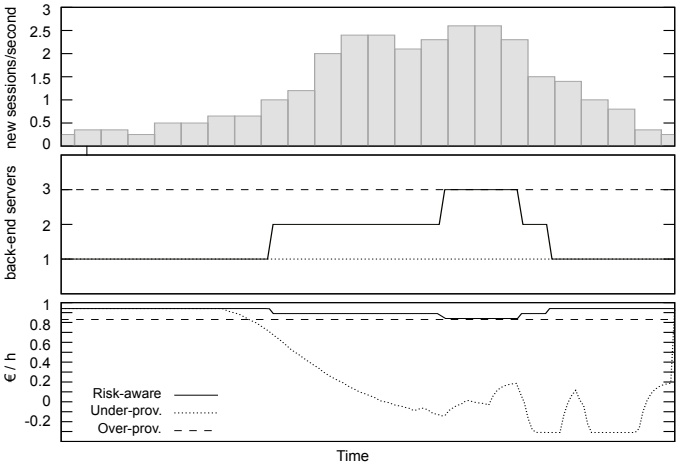

Fig. 4. CSP profit with a one-day typical workload.

the penalties due to SLA violations are very high because the amount of Cloud resources used is not aligned with the application's resources needs (shown in the second subfigure of Figure 4). On the contrary, the risk-aware CSP dynamically adapts the number of back-end servers used. In this sense, the first server is running on in-house resources, while the other ones, needed to attend service's peak demand, are outsourced.

Regarding the cost-benefit analysis (CBA) [33] of the transference of these risks to third-party CIPs, we can clearly determine that its cost is 0.15€ for each Cloud resource outsourced. Hence, the maximum possible gain is precisely the achieved by the risk-aware CSP (21.84€). The consideration of other threats and risks that appear implicitly when performing outsourcing operations are out of the scope of this paper. Therefore, we can conclude the real economic loss due to each risk: 63.23% and 9.89% for the under- and over-provisioned strategies, respectively. Furthermore, the RLEs of these risks to the *ProfMax* BLO provided by the SEBCRA procedure $\left(RLE_{under-prov.}\textit{(ProfMax) = very high threat}\right.$ and $\left.RLE_{over-prov.}\textit{(ProfMax) = medium threat}\right)$ (see Table III) are clearly representative.

Additionally, the fact of being unaware of these risks has some remarkable side effects. The under-provisioned strategy leads to lower customers' satisfaction due to the poor QoS offered in some periods of time, as well as diminished CSP's reliability and reputation. On the other side, the over-provisioned situation decreases greatly the energy efficiency.

As a result, the CSP is able to achieve the maximum profit (91% of the price paid by clients) by considering the transference of these risks to external CIPs, while achieving, at the same time, the maximization of other significant BLOs like energy efficiency, QoS offered, clients' satisfaction, and its reputation and reliability. However, outsourcing operations, needed to perform the transference of risks to outside CIPs, imply also the increasing of the exposure to risks due to this externalization, such as the loss of governance, the appearance of hidden costs and service disruptions from these third-party providers. For these reasons, and depending on provider's interests, is not always beneficial to perform these outsourcing operations.

VII. Conclusions

In this paper, we have introduced risk management into the Cloud computing paradigm. In this direction, we firstly expose the most important Cloud-specific risks and those for which a Cloud Service Provider is exposed. Afterward, we have presented a Cloud-specific risks management procedure oriented to determine the risk impact (either positive or negative) on BLOs. In addition, we have proposed the SEmi-quantitative BLO-oriented Cloud Risk Assessment (SEBCRA) procedure, which is the key process of the whole risk management process. Its main goal is to analyze Cloud risks and, consequently, prioritize them according to their impact on different BLOs. It introduces a methodology that brings transparency to making-decision processes that are based on risk.

Furthermore, we have demonstrated how each risk may have dissimilar repercussions for each BLO considered. As a use case, we have presented a CSP that is able to improve the achievement of a given BLO, i.e. profit maximization. After performing the SEBCRA method in this direction, we have observed that the most impacting risks on this BLO are those concerning the provisioning of the CSP's private Cloud. Finally, the results obtained from the experimentation conducted have confirmed that the CSP is able to maximize its profit by transferring these risks to third-party CIPs.

A. Future Work

Our future work includes the completion of the BLO-driven Cloud risk management introduced herein. Its integration into a Cloud management framework needs the implementation of an autonomic risk-aware scheduler, which will be based on business-driven policies and heuristics that help the CSP to improve its reliability. Moreover, it will assist in deciding which third-party CIP to choose depending on CSP's BLOs, Cloud service's requirements, and CIPs' dependability. In addition, we are pondering the inclusion of the ALARP (As Low As Reasonably Practicable) principle [34] into the risk analysis subprocess of the SEBCRA method. Basically, a risk is within the ALARP range if the cost needed to reduce it is greatly disproportionate to the benefit gained. In fact, this consideration involves a cost-benefit analysis (CBA) [33].

Another important issue will be the tackling of scenarios where multiple BLOs are defined by Cloud organizations. In this cases, several trade-offs appears and, therefore, complex business-driven management policies need to be developed.

Finally, we will carefully treat all the other Cloud-specific risks named in Section II-B. We will go toward the incorporation of risk-related parameters into SLAs. The CSP would perform risk estimations prior to SLA negotiation with the aim of well-establishing SLAs key attributes (i.e. price, penalties, etc.). Afterward, the Cloud service would only be deployed if those SLA's parameters and the corresponding risks are acceptable for the CSP.

Acknowledgment

This work is supported by the Ministry of Science and Technology of Spain and the European Union (FEDER funds) under contract TIN2007-60625 and by the Generalitat de Catalunya under contract 2009-SGR-980.

References

[1] R. Buyya, C. Yeo, and S. Venugopal, "Market-oriented cloud computing: Vision, hype, and reality for delivering it services as computing utilities," in *10th IEEE International Conference on High Performance Computing and Communications, 2008. HPCC'08*, 2008, pp. 5–13.
[2] "Google App Engine," Website, 2009, code.google.com/appengine.
[3] "Amazon EC2," Website, 2010, http://aws.amazon.com/ec2.
[4] "Risk Management Standard," *The Institute of Risk Management (IRM), London*, Website, 2002.
[5] "ISO Guide 73:2009," *Risk Management Vocabulary*, 2009, http://www.iso.org/iso/catalogue_detail?csnumber=44651.
[6] "ISO 31000:2009," *Risk management - Principles and guidelines*, 2009, http://www.iso.org/iso/catalogue_detail?csnumber=43170.
[7] D. Macdonald, *Practical Machinery Safety*. Newnes, 2004.
[8] Cloud Security Alliance, "Security Guidance for Critical Areas of Focus in Cloud Computing V2.1," *http://www.cloudsecurityalliance.org/csaguide.pdf*, 2009.
[9] "ENISA Cloud Computing Risk Assessment," *http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment*, 2009.
[10] H. Motahari-Nezhad, B. Stephenson, and S. Singhal, "Outsourcing Business to Cloud Computing Services: Opportunities and Challenges," *Report HPL-2009-23. HP Labs*, 2009.
[11] T. Aven, "Risk Analysis and Management. Basic Concepts and Principles," *Reliability & Risk Analysis: Theory & Applications*, vol. 2, 2009.
[12] T. Aven and O. Renn, "On Risk Defined as an Event where the Outcome is Uncertain," *Journal of Risk Research*, vol. 12, no. 1, pp. 1–11, 2009.
[13] T. Aven, "A Semi-quantitative Approach to Risk analysis, as an Alternative to QRAs," *Reliability Engineering and System Safety*, vol. 93, no. 6, pp. 790–797, 2008.
[14] C. Yeo and R. Buyya, "Integrated Risk Analysis for a Commercial Computing Service," in *IEEE International Parallel and Distributed Processing Symposium (IPDPS'07), 26-30 March*, 2007, pp. 1–10.
[15] D. Irwin, L. Grit, and J. Chase, "Balancing Risk and Reward in a Market-based Task Service," in *13th IEEE International Symposium on High Performance Distributed Computing, June 04–06*, 2004, pp. 160–169.
[16] F. Popovici and J. Wilkes, "Profitable Services in an Uncertain World," in *2005 ACM/IEEE conference on Supercomputing, November 12–18, 2005*, 2005, p. 36.
[17] A. McNeil *et al.*, *Quantitative Risk Management: Concepts, Techniques, and Tools*. Princeton Series in Finance, 2005.
[18] G. Apostolakis, "Bayesian Methods in Risk Assessment," *Advances in Nuclear Science and Technology*, vol. 13, pp. 415–465, 1981.
[19] P. Krause, J. Fox, P. Judson, and M. Patel, "Qualitative Risk Assessment Fulfils a Need," *Lecture Notes in Computer Science*, vol. 1455, pp. 138–156, 1998.
[20] K. Djemame, I. Gourlay, J. Padgett, K. Voss, and O. Kao, "Risk Management in Grids," *Market-Oriented Grid and Utility Computing*, p. 335, 2009.
[21] "AssessGrid - Advanced Risk Assessment & Management for Trustable Grids," Website, 2009, http://www.assessgrid.eu/.
[22] "FERMA's Risk Management Standard," Website, 2002, Available at http://www.ferma.eu/Portals/2/documents/RMS/RMS-UK(2).pdf.
[23] "ISO/IEC 27005:2008," *Information technology - Security Techniques - Information security risk management*, 2008, http://www.iso.org/iso/catalogue_detail?csnumber=42107.
[24] R. Wideman, *Project and Program Risk Management: a Guide to Managing Project Risks and Opportunities*. Project Management Institute, 1992.
[25] "Amazon EC2 SLA," Website, 2010, http://aws.amazon.com/ec2-sla/.
[26] "Amazon S3 SLA," Website, 2010, http://aws.amazon.com/s3-sla/.
[27] D. Carrera, V. Beltran, J. Torres, and E. Ayguade, "A Hybrid Web Server Architecture for e-Commerce Applications," in *11th International Conference on Parallel and Distributed Systems (ICPADS'05), July 20–22*, 2005, pp. 182–188.
[28] "Squid Proxy," Website, 2010, http://www.squid-cache.org.
[29] "EMOTIVE Cloud," Website, 2010, http://www.emotivecloud.net.
[30] "SPECweb2009," Website, 2010, http://www.spec.org/web2009/.
[31] "Httperf tool," Website, 2008, http://www.hpl.hp.com/research/linux/httperf/.
[32] J. O. Fitó, I. Goiri, and J. Guitart, "SLA-driven Elastic Cloud Hosting Provider," in *18th Euromicro Conference on Parallel, Distributed and Network-based Processing (PDP'10)*, Pisa, Italy, February 17–19 2010.
[33] E. Mishan and E. Quah, *Cost-Benefit Analysis*. Routledge, 2007.
[34] R. Melchers, "On the ALARP Approach to Risk Management," *Reliability Engineering and System Safety*, vol. 71, no. 2, pp. 201–208, 2001.