

Toward Business-driven Risk Management for Cloud Computing

J. Oriol Fitó, Mario Macías and Jordi Guitart

Barcelona Supercomputing Center and Technical University of Catalonia

Barcelona, Spain

{josep.oriol, mario.macias, jordi.guitart}@bsc.es

Abstract—The Cloud computing paradigm is offering an innovative and promising vision concerning the Information and Communications Technology (ICT). 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 present a Cloud computing risk management approach aware of the Business-Level Objectives (BLOs) of a given Cloud organization. More to the point, we propose an innovative SEMI-quantitative BLO-driven Cloud Risk Assessment (SEBCRA) as its core subprocess.

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

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]. 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, many Cloud computing models have emerged at different degrees of flexibility and involve 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, the treatment of risks in Cloud environments must be performed at service, data, and infrastructure layers. In addition, and entering in detail in the core subprocess of managing risks, i.e. risk assessment, there are three primary methods according to [4]: quantitative, qualitative and semi-quantitative (or hybrid). Quantitative risk assessments have been criticized for being overly reductive and divert attention from preventive actions. 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 assessments is that they prioritize risks and identify the most important areas for improvement. Even so, they don't provide enough quantifiable measurements concerning probabilities and impacts of risks. As a result, semi-quantitative risk assessments replace very well tedious quantitative approaches [5], and incomplete qualitative methods.

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. 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 CSP which interoperates with CIPs in order to consume resources from their public Clouds, risks due to these outsourcing operations 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).

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. Basically, they allow any Cloud organization to be aware of Cloud risks and align their 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 Cloud-related risks.

II. BLO-DRIVEN CLOUD RISK MANAGEMENT

In this section, we introduce a novel Cloud risk management process driven by organization's high-level interests. In essence, it is designed to address impacts and consequences of Cloud-specific risks into BLOs of a given Cloud organization.

In fact, its main goal is to increase the probability of success and, thus, decrease both the chance to failure and the uncertainty in achieving those objectives. 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 [6]): *SEBCRA (Risk Assessment)*, which is essentially the overall process of risk analysis and evaluation; *Risk Reporting* and communication; *Risk Treatment*, which implements and selects risk-aware policies, as well as measures, actions, and controls to face with organization’s risks; and *Risk Monitoring* where all the above steps are reviewed.

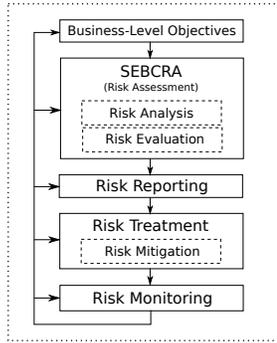


Fig. 1. BLO-driven Cloud risk management steps.

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 best efforts on dealing with threatening risks. In this sense, we present a new information security risk assessment model, i.e. SEBCRA, which has the known purpose of ranking 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 (B_i) affected. Generally, the whole assessment method is subdivided into the risk analysis and its evaluation. The *risk analysis* 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 (based on ISO/IEC 27005:2008 [7]). We can divide risk analysis in three stages: *Risk identification*, which establishes and defines organization’s potential risks; *Risk description*, that guarantees a comprehensive risk assessment method; and *Risk estimation*, which figures out the likelihood of occurrence and the estimated impact on BLOs of each risk previously recognized. Those impacts are considered in terms of threats (downside risks) and opportunities (upside

risks) and are usually evaluated using 3x3, 4x4 or 5x5 risk-level matrices, depending on the granularity of risk assessment desired. 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): 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 on a given BLO ($I_i(B_i)$), either a threat, a benefit, or both, semi-quantified between very high (-100/100, for negative and positive impact, respectively), high (-75/75), medium (-50/50), low (-25/25) and very low (-10/10); and the *Risk Level Estimation* for each BLO ($RLE_i(B_i)$), which is proportional to the *probability* of a given risk and its *impact* on the BLO in question, 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 $-100 \leq RLE_i(B_i) \leq -50$; *unacceptable* if $-50 < RLE_i(B_i) < -10$; *negligible* if $-10 \leq RLE_i(B_i) \leq 10$; *profitable* if $10 < RLE_i(B_i) < 50$; and *high profitable* if $50 \leq RLE_i(B_i) \leq 100$. Therefore, we have to avoid risks with a RLE within the critical or unacceptable ranges, and take advantage of those that lead to an improvement in achieving the BLOs considered. For a better understanding, in Table I we illustrate all the possibilities concerning risk level estimations for a given BLO in terms of semi-quantitative ranges.

Once risks has been assessed, the *Risk Treatment* subprocess defines 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.

III. USE CASE: RISK MANAGEMENT IN A CSP

A. 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 and risk level estimations on BLOs as described in Table II. On one hand, the risk of over-provisioning can appear at any moment ($P_i = \text{‘frequent’}$) and its risk level estimations are: *critical* for *hazard events minimization (HazMin)*, *energy efficiency maximization (EnEffMax)*, and *profit maximization (ProfMax)*, because the exposure to threat risks increases, the provider is consuming more energy than the strictly needed and it pays for more resources than necessary, respectively; and *negligible* for *reliability maximization (RelMax)*, *reputation maximization (RepMax)*, *trust maximization (TrustMax)*, *Quality of Service*

		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	High 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 SEMI-QUANTITATIVE RANGES) OF SEBCRA ON BLOS.

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
IMPACTS ON BLOS OF THE RISKS OF PROVISIONING A PRIVATE CLOUD.

maximization (*QoSMax*), and maximization of customers' satisfaction (*SatMax*), because this risk has almost no impact on these BLOs. On the other hand, the risk of under-provisioning has the same probability (frequent), but dissimilar *RLEs*: 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, it does not incur significant improvements in terms of energy efficiency.

B. SEBCRA for Profit Maximization in a CSP

Now we exemplify how the SEBCRA procedure can be used to improve a given BLO. Notice that all risks can have impact on many BLOs, but in this case we only present the consequences on the *ProfMax* BLO. 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 that 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. Indeed, these outsourcing operations are very suitable in this scenario in order to maximize the total profit of the CSP. This fact is demonstrated in the evaluation presented at subsection IV-B.

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.

C. Transferring Risks of Private Cloud Provisioning to CIPs

An over-provisioning strategy implies that servers are under-utilized in low demand situations, with the corresponding 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, some unexpected demands due to sudden events could appear (e.g. slashdot effect).

After assessing risks, now the risk treatment subprocess takes place. In this sense, outsourcing operations to public Clouds allow the CSP to transfer these critical risks, i.e. the risks of provisioning its private Cloud, to third-party CIPs. Actually, outsourcing operations are performed implicitly by the Cloud elasticity method. It comes to play when the CSP needs to scale up the Cloud infrastructure. Thus, 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. 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.

IV. EXPERIMENTATION

A. Experimental Environment

We use Apache Tomcat v5.5 as the back-end web servers of the CSP, with the SPECweb2009 [8] banking web application deployed on them; and EMOTIVE [9] as the third-party CIP. 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. Furthermore, we use the following SLA economic parameters: a *Price* of 1€/h; *Cost* of 0.18€/h and 0.15€/h for in-house and outsourced Cloud resources, respectively; and changeable *Penalties* for the provider, which depend on both the degree of SLA violations and their magnitude (see [10] for further information).

B. CSP Profit Maximization

Figure 2 shows, from top to bottom, the variable input load pattern, the number of back-end web servers used, and the instantaneous profit earned by three different CSPs: the *risk-aware*, which uses the SEBCRA procedure; and the two possible cases without using any assessment of risks, i.e. *under-provisioned* and *over-provisioned* private Cloud.

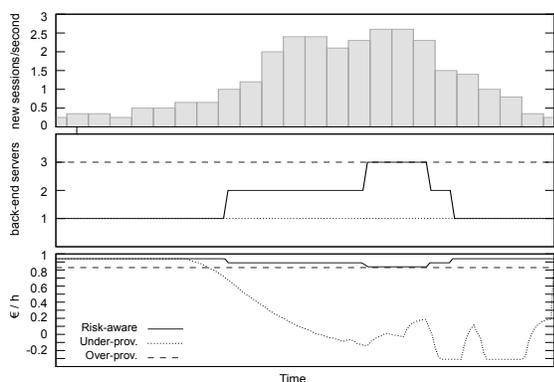


Fig. 2. CSP profit with a one-day typical workload.

The final profits in these three cases are the following: 21.84€/day for the risk-aware, 8.03€/day for the under-provisioned strategy and 19.68€/day for the over-provisioned. So, the real economic loss due to each risk is 63.23% and 9.89% for the under- and over-provisioned strategies, respectively. 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 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. On the contrary, the risk-aware CSP dynamically adapts the number of back-end servers employed. In this sense, the first server is running on in-house resources, while the other ones, needed to attend service's peak demands, are outsourced. As a result, the CSP is able to achieve the maximum profit (91% of the price paid by clients) by transferring 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. Note that the consideration of possible threats that appear implicitly with outsourcing operations are out of the scope of this paper.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we have introduced risk management into the Cloud computing. We have presented a Cloud-specific risks management procedure oriented to determine risks impacts (either positive or negative) on BLOs. In addition, we have proposed the SEBCRA procedure, which has the main goal to prioritize Cloud risks according to their impact on different BLOs. It introduces a methodology that brings transparency to making-decision processes that are based on risk.

As a use case, we have presented a CSP that is able to improve the achievement of a BLO, i.e. profit maximization. The results obtained from the experimentation have confirmed that the CSP is able to maximize its profit by transferring private Cloud's provisioning risks to third-party CIPs.

Our future work includes the completion of the BLO-driven Cloud risk management introduced herein. Its integration into a Cloud management framework needs an autonomic risk-aware scheduler, which will be based on business-driven policies and heuristics that help the CSP to improve its reliability. Moreover, we will tackle scenarios where multiple BLOs are defined by organizations. In these cases, several trade-offs come up, therefore, complex business-driven management policies need to be developed. Finally, we will extensively address all the other Cloud-related risks named in [11].

ACKNOWLEDGMENT

This work is supported by the Ministry of Science and Technology of Spain and the European Union (FEDER funds) under contract TIN2007-60625, by the Generalitat de Catalunya under contract 2009-SGR-980, and by the European Commission under FP7-ICT-2009-5 contract 257115 (OPTIMIS).

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.ine.
- [3] "Amazon EC2," Website, 2010, http://aws.amazon.com/ec2.
- [4] D. Macdonald, *Practical Machinery Safety*. Newnes, 2004.
- [5] 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.
- [6] "FERMA's Risk Management Standard," Website, 2002, Available at http://www.ferma.eu/Portals/2/documents/RMS/RMS-UK(2).pdf.
- [7] "ISO/IEC 27005:2008," *Information technology - Security Techniques - Information security risk management*, 2008, http://www.iso.org/iso/catalogue_detail?csnumber=42107.
- [8] "SPECweb2009," Website, 2010, http://www.spec.org/web2009/.
- [9] "EMOTIVE Cloud," Website, 2010, http://www.emotivecloud.net.
- [10] 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.
- [11] J. O. Fitó and J. Guitart, "Introducing Risk Management into Cloud Computing," Computer Architecture Department, Technical University of Catalonia, Tech. Rep. UPC-DAC-RR-2010-33, 2010.