A Security Risk Assessment Model for Business Process Deployment in the Cloud

Elio Goettelmann, Karim Dahman, Benjamin Gateau, Eric Dubois, Claude Godart

► To cite this version:

Elio Goettelmann, Karim Dahman, Benjamin Gateau, Eric Dubois, Claude Godart. A Security Risk Assessment Model for Business Process Deployment in the Cloud. IEEE International Conference on Services Computing, Jun 2014, Anchorage, AK, United States. pp.307 - 314, 10.1109/SCC.2014.48 . hal-01095874

HAL Id: hal-01095874

<https://hal.inria.fr/hal-01095874>

Submitted on 16 Dec 2014

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

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

A Security Risk Assessment Model for Business Process Deployment in the Cloud

Elio Goettelmann^{1,2}, Karim Dahman¹, Benjamin Gateau², Eric Dubois² and Claude Godart¹

¹LORIA - INRIA Grand Est

²CRP Henri Tudor

Université de Lorraine, Nancy, France L-1855 Luxembourg-Kirchberg

{*elio.goettelmann, karim.dahman, claude.godart*}@loria.fr, {*benjamin.gateau, eric.dubois*}@tudor.lu

Abstract—Managing security risks on information systems is essential to guarantee their security while handling costs. However, the complexity of risk assessments is greatly increased when data is spread on multiple environments. In this paper we present a security risk assessment model for distributing business processes in a multi-cloud environment. We aim at offering the full power of cloud computing to composite applications while shielding companies from the complexity related to security risk assessments in the Cloud. We also want to give them the capability to automatically generate secure and cost-effective applications across multiple clouds. Our approach is based on existing risk assessment methodologies, while using the industry recognized IT standards.

Index Terms—Business Process; Security Risk Management; Cloud

I. INTRODUCTION

Cloud computing avoids upfront infrastructure costs, and helps organizations to focus on their core business activities, instead of their system infrastructure [1]. It enables to rapidly adjust resources to meet unpredictable demand and transforms investment costs in operating expenses through a pay-as-you-go approach. So, cloud computing can be of great benefit, especially in times where cost reduction plays a vital role.

But recent events, like the NSA spying scandal, underscore the threats of such services. Using off-premise and shared infrastructures exposes the information systems of companies to new kind of security risks. A taxonomy of these risks are published by CSA [2] and ENISA [3]. New techniques and solutions, to prevent harms that can adversely affect companies business activities, must be defined before broader adoption.

One technique, proposed by Jensen et al. [4] or by AlZain et al. [5], consists in spreading applications over different locations to increase the complexity for a third party to gather sensitive business information. In previous work [6], we presented how to implement automatically such a solution by working on business processes, which formally structure and describe the activities of a company. However, it does not obviate the realization of **security risk assessments** [7] to understand all of the security risks. And current assessment methods, if applicable for on-premises architectures, need to be adapted for multi-cloud environments.

In this sense, our goal is to reduce the complexity of security risk assessment in a cloud context, in order to generate secure and cost-effective applications on multi-cloud platforms. In fact, existing cloud solutions offer different types

of services regarding security, whereas the business processes have their own security requirements. As a consequence, two important questions arise when outsourcing processes in a cloud environment. How to compare different cloud offers when considering security? How to align process requirements with existing cloud offers? To answer these questions, we present an approach for assessing security risks in a cloud context before distributing a business process execution across multiple clouds.

The paper is organized as follows. Section II presents some background about security risk assessment applied to cloud computing, a motivating example and an overview of our framework. Section III formally defines our concepts and presents our approach on security risk assessment in a multi-cloud environment. In Section IV, we demonstrate a proof-of-concepts implementation and experimental evaluations. Then, we situate our work in the related research in Section V, and we finally present future extensions in Section VI.

II. BACKGROUND, MOTIVATING EXAMPLE AND OVERVIEW

Here we give some definitions to explain why current risk assessment methodologies need to be adapted for cloud computing. We also present a motivating example used all along this paper and overview our approach.

A. Security Risk Assessment Background and Cloud Computing

Roughly speaking, a **risk** is defined as the combination of the *probability* that an *event* occurs and its *consequences* [8]. In the IT security context, where IT components (ex. hardware, network, etc.) support business assets (ex. information, processes, etc.), the **security risk** is defined in a more fine-grained fashion. The event is usually seen as a **threat** which uses one or more **vulnerabilities** of the IT components in order to create a negative **impact** (ex. destruction, alteration, theft, etc.) on the business assets [7]. For instance, an attacker steals customers' information (i.e., threat) through a compromised interface (i.e., vulnerability) which leads to the business reputation loss (i.e., impact).

In this sense, a **security risk assessment** consists usually in evaluating the following formula: **risk** = **vulnerability** × **threat** × **impact** ([9], [8]). The goal is to estimate security risks in a quantitative and/or qualitative manner, to select

Fig. 1. Business Process Motivation Example: Insurance Claim Recovery Chain.

those that need to be reduced and to develop countermeasures. Developing countermeasures involves the implementation of **security controls** by constraining technical solutions and by reducing vulnerabilities on the business settings. Security controls are management, operational, or technical safeguards prescribed for an information system to protect the confidentiality, integrity, and availability of the system and its information. Examples are *firewalls* or *intrusion detection systems*.

In a Cloud context, the major problem of security risk assessments is that parts of the IT are outsourced. Thus, **Cloud consumers** have no control over the equipment hosting their assets and must integrate with an architecture defined by the **Cloud provider**. So, the identification of vulnerabilities becomes a complex task, especially as Cloud providers may be tempted to conceal the vulnerabilities of their technical solutions. However, we are convinced that in a mature market, Cloud consumers will be able to select Cloud providers with favorable terms, especially, secure Cloud offers at suitable prices. In such a direction, standards like CSA Cloud Control Matrix [10] or ISO 27017 [11] help Cloud providers to secure their installations and service offers.

Nevertheless, the impact of given cloud security risks, i.e. the potential financial losses, are still related to Cloud consumers, as they hold the value of their business assets. So, security risk assessments cannot be delegated to Cloud providers. Thus, in order to leverage the global quality of the security risk assessments, some **Cloud brokers** have emerged. Cloud brokers are entities managing the use, performance and delivery of Cloud services [12]. They can help customers to choose adequate Cloud providers according to their security requirements. They are responsible for evaluating the likelihood of risks by analyzing the security controls that are implemented by Cloud providers.

B. Motivating Example - Insurance Recovery Chain

To illustrate the goal of our proposal, we define an *insurance claim recovery chain* [13] as a business process model using BPMN notation (depicted in Fig. 1). In this business scenario, a *beneficiary* initiates the claim recovery chain by sending a declaration to the *insurance* company. First, to handle an incident declaration, the *insurance* invokes the *emergency* service to obtain a statement on the incident details. Second, with

the details provided by the *emergency* service, the *insurance* collects in parallel the different reports from the *hospital* and the *police*. They are needed by the *expert* to decide on the reimbursement amount. Finally, the *bank* is optionally requested for a transfer and a reimbursement decision is sent to the *beneficiary* in order to complete the claim recovery chain.

The data associations (dotted arrows) between tasks show the data dependencies between the service-component that implement the processes. For example, the *Initiate Claim* activity transfers the *Claim* document to the *Obtain Incident Details* activity.

Now, suppose this company wants to outsource this process to the cloud to save money. The problem is, that the company has not necessarily the competencies to do this on its own. First of all, it does not know the cloud specific security risks and how to evaluate them for this process handling private data about its clients. Then, the selection of adapted cloud providers can be really time consuming, as there are many providers and it is not always easy to compare their offers. Finally, the company knows that for cost-effectiveness and security reasons it would be better to distribute the process on multiple providers, but it does not know how the decomposition for this purpose is done. Thus, the company needs some support.

Note that it can also have additional requirements about the distribution of a process. As example, the dotted boxes in Fig. 1 describe some design decisions (modeled as constraining groups on tasks) made for this *insurance* service. The *inclusion* means that the tasks have to enact in the same context. The *exclusion* means that the tasks must be distributed among different contexts. See [14] for more about that.

C. Overview of the approach

To support such requirements for security risk assessment and secure deployment of a business process in the Cloud, we have defined the following approach. It consists of a design-time framework for generating secure and cost-effective process-centered applications in cloud environments.

First, a *Cloud consumer* (e.g. the *insurance* company) models a **business process** extended with non-functional requirements (**security needs**). He transmits them to a *Cloud broker* who analyzes the **cloud offers** coming from different *Cloud providers*. The broker realizes a **risk assessment** by

Fig. 2. Our design-time framework for multi-cloud business process deployment.

evaluating the risk of the different available offers for the given needs.

Then, the broker conducts a **cloud selection** based on the functional requirements and the costs of the available offers. At this point there are two possibilities. Either the risk is taken as a constraint (the offers where the risk is too important are excluded) and the selection optimizes the cost. Or the cost is taken as a constraint and the selection is done in order to minimize the risk.

Finally, the broker undertakes the **process decomposition** to split down the process into sub-processes. As each task can be assigned to a different cloud provider, the structure of the process needs to be adapted. The process is automatically decomposed, and the different fragments are assigned to their respective environments. The assignment of fragments to clouds is what we call a **configuration**. This part is based on a tool we developed in the context of composite service deployment. It is not the objective of the paper to enter in the details of the algorithm for splitting a process. A detailed description of the process decomposition and service-component configurations can be found in [15] and [6].

III. FORMAL FRAMEWORK

This section specifies the concepts of our risk assessment model and formalizes it. Its structure is depicted in Fig. 3. As our approach is based on the duality between risk management and cost optimization, the section is arranged in two parts. First we discuss our risk assessment model (Section III-A), our main target. Then we present a synthesized cost management model (Section III-B).

Fig. 3. Risk Assessment Model for multi-Cloud environments

A. Cloud Security Risk Assessment Model

Definition 1 (Cloud Consumer): An entity which browses the offers from Cloud providers (or Cloud brokers) and uses the service adapted to its functional and non-functional requirements. It is specified as follows:

Asset = $Process \cup Task \cup Data \cup Message$, the set of business assets of the consumer information system,

Objective = $\{Authenticity, Availability, Confidentiality, Integrity, Non-repudiation\}$, a set of security objectives,

Need : $Asset \times Objective \rightarrow Level$, defines security needs,

Threshold : $(\{Low, Medium, High \mid Low < Medium < High\})$, defines a global acceptable risk level,

ExecutionCost : $Task \rightarrow \mathbb{N}^*$, defines task execution costs (in seconds),

SizeCost : $Data \rightarrow \mathbb{N}^*$, defines data size costs (in bytes),

A security **objective** helps to determine the impact of a security risk. It is frequently defined in terms of *confidentiality, integrity, availability, non-repudiation* and *authenticity* (CIANA)¹. These objectives must be fulfilled to ensure the security of a technical solution and its data [8].

Different levels of security **needs** are usually expressed on process data objects ([16],[17]). As business processes are mainly task-centric (typically a task represents a web service call), we need to relate these data-centric security needs to the tasks of the business process. This is achieved by applying a simple Bell-LaPadula model as in [18]. Thus, we define the security **need** of a task as the highest needed value of the data it handles. Formally, for a task ta , in- and out-going data elements d of ta , and an objective o , the security need is:

$$Need(ta, o) = \max_{d_i \in Data(ta)} Need(d_i, o) \quad (1)$$

The **threshold** is a global value used for the final Cloud provider selection. It indicates the acceptable level of risk, below which the risk will be retained. It will determine which Cloud provider can be used and which not. This is further illustrated in Section IV-B.

Definition 2 (Cloud Provider): An entity which can hold multiple Cloud offers. It is responsible for the implementation of security controls (typically controls defined by the CSA [10]

¹Legality, robustness, minimal performance or scalability can also be listed [8].

$$Coverage(p, t) = \frac{\sum_{c \in Control} (Implementation(p, c) \times Mitigation(t, c))}{\sum_{c \in Control} Mitigation(t, c)} \quad (2)$$

or the ISO 27017 [11]), to protect its offers from attacks and to comply with regulations. It is specified as follows:

Offer, as a set of Cloud offers,

Control, is a set of controls,

Implementation : $Offer \times Control \rightarrow \{0,1\}$, defines if a control is implemented or not by an offer.

UsageCost, StorageCost, TransferCost : $Offer \rightarrow \mathbb{N}^*$, defines offers' costs in terms of usage (in \$/second), storage (in \$/byte) and transfer (in \$/byte)

Instead of focusing on vulnerabilities, as in [19], our risk assessment approach focuses on security *controls*. As example, the CSA defines 197 controls in its Cloud Control Matrix [10]. Due to a lack of space, we do not show an exhaustive list of these controls, but examples are:

- **IS-19.4** Do you maintain key management procedures?
- **RI-01.1** Is your organization insured by a 3rd party for losses?
- **SA-02.7** Do you allow tenants to use third party identity assurance services?

We indicate the **implementation** of each control by a provider offer with a boolean function. If the offer implements the given security control, the function returns 1, otherwise 0. This choice is recommended by the CSA in [10].

Definition 3 (Cloud Broker): An entity which can provide three types of services [12]: enhance existing services (service intermediation) like security, combine multiple services (aggregation) or measure different providers and select the best (arbitrage). It is specified as follows:

Threat, as a set of security threats

Control, the same set of controls as the Cloud provider

Objective, the same set of objectives as the Cloud consumer

Mitigation : $Threat \times Control \rightarrow \{0,1\}$ indicates if the control mitigates the threat or not,

Consequence: $Threat \times Objective \rightarrow \{0,1\}$ indicates if the threat has a consequence on the objective or not,

As explained previously, there are different cloud-specific security **threats** (ex. *data breaches*, *data losses*, *malicious insiders*, etc.). Each threat has a **consequence** on one or more **security objectives**. We use the CIANA list of objectives and the CSA list of threats, where each threat is related to multiple objectives of a business asset (see [2]). For example, *data breaches* have only a consequence on *confidentiality*. Whereas, *data losses* have a consequence on *availability* and *non-repudiation*. Also, *malicious insiders* have a consequence on all five objectives. This binary relation can be represented in a matrix relating the threats to the objectives.

The threats are **mitigated** by one or multiple **security controls**. We use the matrix from [2] to relate controls to threats. For example, *data breaches* are mitigated by the controls **IS-19.4** and **SA-02.7**. Control **IS-19.4** also mitigates *malicious insiders*, whereas control **RI-01.1** mitigates *insufficient due diligence*.

In order to quantify the risk, we combine the previously defined concepts for calculating two different values for each threat: the **coverage** of a provider, and the **harm** on an asset, as defined below.

Definition 4 (Coverage): A score calculated for a given provider and a given threat. It is calculated with the security controls implemented by the provider and their mitigations on the threat. Formally, for a threat t by a provider p , the **coverage** is defined in Formula.2.

Mostly, a provider who implements many controls will be more secure than one who implements fewer controls. In our approach this score allows us to compare the response of a provider to a specific threat, it corresponds to a percentage of implemented controls. In [2] for example, the CSA shows for each provider, their implemented controls. So, if the CSA gave 10 controls to mitigate a threat, and a provider implemented 5 of them, he gets a coverage of 50%.

Definition 5 (Harm): A level calculated for a given task and a given threat. It is obtained by combining the security needs of the task and the corresponding consequences of the threat. Formally, the **harm** of a threat t on a task ta is:

$$Harm(t, ta) = \sum_{o \in Objective} (Consequence(t, o) \times Need(ta, o)) \quad (3)$$

Basically, some threats will be more important for some assets than others. Thus, we sum the needs of the asset, when there is a consequence on the objective. It corresponds to the exposition of a task to a specific threat. A *denial of service*, for example, has only a consequence on the *availability*. Whereas in our motivating example, the *process reimbursement* task has a *sparse* need in *availability*. So the *denial of service* will not really be harmful for the *process reimbursement* task, indeed it does not prevent the reimbursement itself and could be done later.

Definition 6 (Risk Level): A level calculated for a given threat, a given task and a given provider. It is the sum of the harm of the threat on the task, and the vulnerability of the provider to this threat. The vulnerability is obtained by using the coverage of the provider and a score denoted $Covg_{max}$, which represents the maximum possible value for covering a threat². Formally, the **risk** of a threat t , a task ta and a provider p is:

$$Risk(t, ta, p) = Harm(t, ta) + (Covg_{max} - Coverage(p, t)) \quad (4)$$

With the risk level we are able to compare providers. It classifies the providers towards their risk for deploying a specific task (e.g. Table. III in Section IV). It is similar to the risk formula stated in Section II-A. However, we use a sum to better reflect the independence of the vulnerabilities from the impacts in the cloud context (as the consumer cannot reduce

²Basically, the maximum coverage is reached by implementing all controls that mitigate the threat.

the vulnerabilities of the provider, and the provider cannot be involved for the evaluation of the impact).

B. Cost Model

As the main goal of risk management is to balance the risk against the costs, we define the following cost model. It takes into account three types of costs: *Usage Costs*, *Storage Costs* and *Transfer Costs*. We find that these three types are the most characteristic of the Cloud business model and we can generally map each pricing model to these three attributes.

Usage Costs correspond to the CPU power needed to execute the process. Generally, on IaaS Cloud offers, the consumer can select which computer power he needs (often in terms of GigaHertz³). In our case, we have decided to express these costs in Dollars per GigaHertz on a monthly basis, as many providers still remains to this type of pricing scheme. So, we annotate each activity of our business process model with the need in terms of CPU power (per month), called the *ExecutionCost*.

Storage Costs correspond to the storage space needed by the data used by the process. This is expressed in our model in Dollars per GigaByte (per month). Each data of the process is annotated with its estimated size, the *SizeCost*.

Transfer Costs correspond to the amount of incoming and outgoing data. Generally, Cloud providers bill their consumers according to the transferred Gigabytes of data. These costs are usually expressed in Dollars per GigaByte. We calculate this amount using the *SizeCost* of the data exchanged between the process fragments.

Table IV gives examples of such costs for offers used in our experimentations.

IV. EXPERIMENTATION AND EVALUATION

To demonstrate the feasibility and the interest of our approach, we have developed a use case, based on our motivating example and a particular cloud offering. The objective is to study the balance between security risk and cost, and to see how much more expensive it is to use a less risky configuration.

A. Risk Assessment

We start by applying our risk assesment tool to our use case in three stages: security needs definition, risk evaluation for each provider, too risky cloud provider exclusion.

1) *Security needs definition*: To assess the risk, first we define the **security needs** for our motivating example (Fig. 1). We define three levels for each CIANA objective and annotate the data objects of the process with them (see Table I). Each level is mapped to a value between 0 and 2 (for example $\{Public = 0, Restricted = 1, Secret = 2\}$ for *Confidentiality* and $\{Passable = 0, Alterable = 1, Fixed = 2\}$ for *Integrity*). We assume that the terms used are self-documenting.

As explained in Section III-A (Formula. 1) we calculate with these values the task-centric security needs. For

TABLE I
SECURITY NEEDS ON THE MOTIVATION EXAMPLE.

Data Associations	Confidentiality	Integrity	Availability	Non-repudiation	Authenticity
Claim	Public	Passable	Continuous	Tolerable	Irrelevant
Hospital det.	Public	Alterable	Usual	Futile	Common
Incident det.	Public	Alterable	Usual	Futile	Common
Hospital rep.	Restricted	Fixed	Usual	Trusted	Verified
Police rep.	Restricted	Fixed	Usual	Trusted	Verified
Medical rep.	Secret	Fixed	Usual	Trusted	Verified
Expert rep.	Restricted	Fixed	Usual	Trusted	Verified
Account	Secret	Fixed	Sparse	Futile	Verified
Amount	Restricted	Fixed	Sparse	Trusted	Verified

example, *Obtain Incident Details* is associated to the data objects *Claim*, *Incident Details* and *Hospital Details*. So, by taking the maximum of each data object's need we get the need of *Obtain Incident Details*: $\{0 (Public), 1 (Alterable), 2 (Continuous), 1 (Tolerable), 1 (Common)\}$ for respectively $\{Confidentiality, Integrity, Availability, Non-Repudiation, Authenticity\}$.

2) *Risk evaluation for each provider*: Then, we use the **consequences** to get the **harm** of each task (Formula. 3). For example, the *Malicious Insider* threat (t) has the following harm on *Obtain Incident Details* (ta):

$Harm(t, ta) = (1 \times 0) + (1 \times 1) + (1 \times 2) + (1 \times 1) + (1 \times 1) = 5$. The first value of each bracket is the consequence (always 1 because the *Malicious Insider* threat has a consequence on all five CIANA objectives), the second value is the need obtained previously.

Since consequences are binary values, and the five CIANA objectives are defined on levels from 0 to 2, the harm is a value between 0 and 10.

For the **security controls** and the **mitigations**, we use the STAR Registry and the matrix defined by the CSA [20], [10].

In order to perform realistic tests we selected five Cloud offers from the CSA STAR Registry [20]. Examples are, for the previously given controls (III-A) and the five providers:

- *CloudSigma* AG⁴ and *SHI International*⁵ declare implementing **IS-19.4**, **SA-02.7** and **RI-01.1**.
- *Terremark*⁶ implements **IS-19.4** and **RI-01.1**.
- *FireHost*⁷ and *Softlayer*⁸ only implement **RI-01.1**.

To equally take into account the **harm** and the **coverage**, we brought the percentage value of the coverage on a scale of 0 to 10 (and thus $Covg_{max} = 10$). So each risk is a value between 0 and 20.

For example, in the case of *Obtain Incident Detail* (ta) on *CloudSigma* (p_2), it is the *Malicious Insider* (t) threat which has the greatest risk value (Formula.4):

$$Risk(t, ta, p_2) = Harm(t, ta) + (Covg_{max} - Covg(p, t)) = 5 + (10 - 9) = 6$$

⁴<http://www.cloudsigma.com>

⁵<http://www.shi.com>

⁶<http://www.terremark.com>

⁷<http://www.firehost.com>

⁸<http://www.softlayer.com>

³Amazon proposes another type of measure, Elastic Compute Unit (ECU)

TABLE II
OUTPUT FOR DIFFERENT RUNS

	First run					Second run					Third run					Fourth run				
	Softlayer	CloudSigma	FireHost	SHI Int.	Terremark	Softlayer	CloudSigma	FireHost	SHI Int.	Terremark	Softlayer	CloudSigma	FireHost	SHI Int.	Terremark	Softlayer	CloudSigma	FireHost	SHI Int.	Terremark
Initiate claim				x					x			x					x			
Obt. incident det.				x					x	x					x					x
Obt. hospital rep.				x					x					x			x			
Obt. police rep.					x					x				x			x			
Req. claim docs.					x				x								x			
Obt. expert rep.					x					x		x			x		x			
Send reimb. dec.					x				x			x					x			
Process reimb.					x				x					x			x			
Risk			7					6					5					5		
Cost (\$/mo)			101.81					231.32					284.08					477.73		

According to [20], *CloudSigma* has a coverage for this threat of 10, and the harm of this threat on *Obtain Incident Detail* is equal to 5. The risk values for the other threats (as *Data Loss* or *Data Breach*) are below this value. We can not show the details for all the threats, as it would be extraneous for this paper.

Table III shows the maximum acceptable risk value for the nine CSA threats, for the tasks of our motivating example, of the five providers.

TABLE III
MAXIMUM RISK VALUE OF THE ASSETS FOR EACH PROVIDER

	Softlayer	CloudSigma	FireHost	SHI Int.	Terremark
Initiate claim	12	9	13	10	13
Obt. incident det.	8	6	9	6	9
Obt. hospital rep.	12	9	13	10	13
Obt. police rep.	11	8	12	9	12
Req. claim docs.	6	2	7	3	4
Obt. expert rep.	11	8	12	9	12
Send reimb. dec.	12	9	13	10	13
Process reimb.	11	8	12	9	12

Note that, as our approach is model driven, it can also work with different type of objectives (not only CIANA) and other kind of threats (as the 35 ENISA threats).

3) *Too risky provider exclusion*: In accordance with the consumer, the broker defines the level of acceptable risk (referred to as **threshold**). For a given asset, this threshold is used to exclude providers above this value. This eliminates all deployment options where the risk is too high.

Back to our example, supposing that the threshold is set to 9 by designers, the cells of eliminated providers are grayed out in Table. III. Respectively a white cell means that the asset can be deployed on the provider.

B. Cloud Selection

Here we select the best cloud in two stages: different configurations evaluation and final clouds selection.

1) *Configurations evaluation*: The problem is how to assign n different tasks to p different Cloud offers. Thus, we are

facing a Quadratic Assignment Problem (QAP) [13]. Therefore we have proposed a heuristic approach to find a good configuration in an acceptable time. We do not show details about our algorithm, as it is not the primary objective of our proposal, but the principle is as follows. First, we build an initial solution (so-called *Greedy* solution) in an iterative way for all process tasks. The result is then improved by a second algorithm (*Tabu* search). Details about these algorithms are available in [13].

We have tested our algorithms applied to our use case in four different ways to compare the different outputs. It is based on the pricing models of our 5 Cloud offers available in Table IV. The results can be seen in Table II.

- 1) The first run does not contain any constraint relative to the risk. We take only into account the inclusion and exclusion constraints and the costs. This is done in order to get a "cheap" solution by neglecting the security risks.
- 2) For the second run we constrain the configuration with a global threshold for the risk of 6. This means that the first configuration is no longer possible. It results in an output where some tasks are moved to another more "secured", but also more expensive offer.
- 3) For the third run we decrease the global risk threshold to 5, in order to get a even more "secured" solution than the previous one. This time, the tasks are located on three different offers. The configuration is even more expensive, but with a more acceptable global risk level.
- 4) In the last run we check if moving a maximum of tasks on the most "secured" offer has some significant changes on the risk or the costs. For our example this has no impact on the risk value, whereas a very significant impact on the costs.

TABLE IV
COSTS OF 5 CLOUD OFFERS

	UsageCost (\$/GHz/mo)	StorageCost (\$/GB/mo)	TransferCost (\$/GB)
Softlayer	\$20.00	\$0.10	\$0.10
CloudSigma AG	\$13.86	\$0.18	\$0.06
FireHost	\$25.70	\$2.78	\$0.50
SHI International, Corp.	\$11.56	\$0.29	\$0.01
Terremark	\$3.60	\$0.25	\$0.17

It is possible that no configuration is found if the threshold risk value is too low. In this case, either the user increases the threshold value, or he considers the Cloud context as too risky and decides not to move the process to the Cloud.

For our motivating example we constrained the risks and optimized the costs. But our algorithm also allows constraining the cost while optimizing the risk, which could be interesting in some other cases.

2) *Final configuration selection*: At this point we have a set of possible configurations (Table II). By making a balance between risk level and costs, a risk manager can make a choice.

Suppose that finally he chooses the *Third run*. The problem is now to deploy our example process on the selected clouds.

C. Process deployment in the cloud

Process deployment is done in two parts: first we decompose the initial process into fragments and then we deploy these fragments on the corresponding clouds.

1) *Process decomposition*: The decomposition transforms a single process, into a new collaboration between separate process fragments (a fragment groups activities executed on the same cloud). The combined behavior of those fragments provides the same execution behavior as the initial process. Each resulting fragment represents an autonomous business process with additional synchronization tasks. These synchronization tasks are added to the process for guaranteeing the control flow of the initial process. They send messages to notify remote fragments about their execution status. More details about process decomposition can be found in [21] and in [6].

Fig. 4 depicts the decomposition of our motivating example in the selected configuration (grey activities are synchronization tasks).

2) *Deployment in clouds*: The last step of our approach consists in deploying this configuration on the selected Cloud offers. In [6] we presented how we deploy such service composition as BPEL programs on remote service orchestration engines (e.g., Apache ODE). But as current Cloud offers do not support this kind of services, we selected offers providing infrastructure services and used them to deploy our own execution engines. We are confident that such type of platforms will rapidly emerge in the future, WSO2 Stratoslive⁹ is such a PaaS even if still not available through web services. For some type of processes, tasks could be mapped to existing offers at the SaaS layer of the Cloud. For others, not service based, a development phase can be required before deploying them on the Cloud.

V. RELATED WORK

In [22] authors present abstract design methods to split an application such that various parts of it can be moved to different target Cloud environments. This split can be done manually or use optimization algorithms. One key feature of these methods is that the moving-to-Cloud problem is

considered in the early stages of business process and service modelling. However, these related work do not propose an approach to directly relate partitioned processes to the existing Clouds. Also they do not consider security aspects.

Managing security risks in business processes is currently a hot topic [23]. But most of the current researches mainly focus on process execution, as in [24] where manual decisions are supported by a risk-prediction algorithm. At the opposite, we focus on automating decisions at design-time, before deployment. These approaches are complementary, but also focus on different kind of risks.

Other related work as [25] try to secure business processes by modifying them through patterns after a risk assessment. In [26] the authors even work on processes in a cloud context; they propose a tool for analysing processes and decide if their outsourcing is possible. But none of these methods explain how the risk is calculated and the process is broken down.

Some models like the Common Vulnerability Scoring System¹⁰ or this presented in [19] try to evaluate security vulnerabilities for SOA. However, they present some limitations in the Cloud context as vulnerabilities are more difficult to establish and need to be verified. Emerging standards as [11] or [10] seem to indicate that our coverage approach using security controls is more adapted.

Another important aspect of our proposal is the ability to combine risk assessment with other constraints (cost, quality of service, etc.), which are rarely considered together in similar approaches. In [18] a cost model is presented and used to decompose processes among a set of cloud providers. But the security model is limited as it defines arbitrary security levels for each provider. Moreover, the author considers all deployment options, whereas when working in a realistic cloud environment with many different providers, this is not always feasible. Our previous work in [6], [21] is on this vein, but limited to simple security constraints without risk measurement.

More generally, cloud computing costs are investigated in many studies, like the the Total Cost of Ownership Approach presented in [27]. Even if they are defined in a more detailed fashion and are considering many other aspects, they are not yet adapted for an automated treatment, and especially not considered in the business process context.

VI. CONCLUSION AND FUTURE WORK

In this paper, we have proposed a technique for assessing security risks of business processes before deploying them in a multi-Cloud environment. This technique relies on two main aspects, on the one side business process security needs and on the other side Cloud providers guideline conformance. By combining the impact evaluation on the Cloud consumer and the vulnerability assessment of the Cloud provider, a Cloud security broker can help companies to deploy securely their applications on a multi-Cloud environment. We also included a

⁹WSO2 Business Process Server, <http://wso2.com/products/business-process-server>

¹⁰<http://nvd.nist.gov/cvss.cfm>

Fig. 4. A partitioned orchestration process.

cost model as a representation to confront security dimensions with other QoS properties.

To illustrate our approach, we have used the example of an insurance company. We have defined the security needs on the five CIANA objectives, and we used the CSA security controls to evaluate the risk levels of five Cloud providers taken from an industry-recognized registry: STAR [20]. But remember that our approach is model driven, so it can be extended to other sets of controls or objectives.

The approach has some limitations which will be addressed in future works. The first one is the shortage of empirical evaluation on real case studies. This is planned in the future with domain experts and industrial partners. Another one is that the risk assessment and the provider selection is done at design-time, it would be interesting to extend our approach for a configuration at run-time. As the Cloud is a very dynamic context, taking into account the change of either the business process or the Cloud providers would be an interesting improvement. We also argue that for the presented example, the binary values for the implementations and mitigations are sufficient for comparing the different solutions. But an extension could be to add a weighting. For instance, some controls could be rated as "better" mitigations than others, and a provider could also implement "partially" a control.

REFERENCES

- [1] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "A view of cloud computing," *Commun. ACM*, vol. 53, no. 4, pp. 50–58, Apr. 2010.
- [2] Cloud Security Alliance, "The Notorious Nine - Cloud Computing Top Threats in 2013," Tech. Rep., 2013.
- [3] European Network and Information Security Agency, "Benefits, risks and recommendations for information security," Tech. Rep., 2009.
- [4] M. Jensen, J. Schwenk, J. Bohli, N. Gruschka, and L. Iacono, "Security prospects through cloud computing by adopting multiple clouds," in *CLOUD'11*, 2011, pp. 565–572.
- [5] M. AlZain, E. Pardede, B. Soh, and J. Thom, "Cloud computing security: From single to multi-clouds," in *HICSS'12*, 2012, pp. 5490–5499.
- [6] E. Goettlmann, W. Fdhila, and C. Godart, "Partitioning and cloud deployment of composite web services under security constraints," in *IC2E'13*, 2013. [Online]. Available: <http://eprints.cs.univie.ac.at/3565/>
- [7] N. Mayer, "Model-based Management of Information System Security Risk," Ph.D. dissertation, University of Namur, Apr. 2009. [Online]. Available: <http://tel.archives-ouvertes.fr/tel-00402996>
- [8] National Institute of Standards and Technology, "Information Security - Guide for Conducting Risk Assessments," 2002.
- [9] "AS/NZS 4360 SET Risk Management, Australian/New Zealand Standards," 2004.
- [10] Cloud Security Alliance, "Cloud Control Matrix," <https://cloudsecurityalliance.org/download/cloud-controls-matrix-v3/>, Tech. Rep., 2014.
- [11] "ISO/IEC 27017, Information tech., Security techniques, Code of practice for information security controls for cloud computing services based on ISO/IEC 27002."
- [12] National Institute of Standards and Technology, "Cloud Computing Reference Architecture," 2011.
- [13] W. Fdhila, M. Dumas, and C. Godart, "Optimized decentralization of composite web services," in *CollaborateCom'10*, 2010, pp. 1–10.
- [14] E. Goettlmann, N. Mayer, and C. Godart, "A general approach for a trusted deployment of a business process in clouds," in *MEDES'13*, 2013.
- [15] K. Dahman, F. Charoy, and C. Godart, "Alignment and change propagation between business processes and service-oriented architectures," in *SCC'13*, Santa Clara, CA, United States, Jun. 2013.
- [16] T. Grandison, M. Bilger, L. O'Connor, M. Graf, M. Swimmer, M. Schunter, A. Wespi, and N. Zunic, "Elevating the discussion on security management: The data centric paradigm," in *BDIM '07*, 2007, pp. 84–93.
- [17] W. Zhou, M. Sherr, W. R. Marczak, Z. Zhang, T. Tao, B. T. Loo, and I. Lee, "Towards a data-centric view of cloud security," in *CloudDB'10*, 2010, pp. 25–32. [Online]. Available: <http://doi.acm.org/10.1145/1871929.1871934>
- [18] P. Watson, "A multi-level security model for partitioning workflows over federated clouds," in *CloudCom*, 2011, pp. 180–188.
- [19] S. Sackmann, L. Lowis, and K. Kittel, "A risk based approach for selecting services in business process execution," in *Wirtschaftsinformatik (1)*, 2009, pp. 357–366.
- [20] Cloud Security Alliance, "Security, Trust and Assurance Registry," <https://cloudsecurityalliance.org/star/>, Tech. Rep., 2014.
- [21] W. Fdhila, U. Yildiz, and C. Godart, "A flexible approach for automatic process decentralization using dependency tables," in *ICWS '09*. Washington, DC, USA: IEEE Computer Society, 09, pp. 847–855.
- [22] F. Leymann, C. Fehling, R. Mietzner, A. Nowak, and S. Dustdar, "Moving applications to the cloud: an approach based on application model enrichment," *IJCIS*, vol. 20, no. 3, pp. 307–356, 2011.
- [23] N. Ahmed and R. Matulevicius, "A taxonomy for assessing security in business process modelling," in *RCIS*, 2013, pp. 1–10.
- [24] R. Conforti, M. de Leoni, M. L. Rosa, and W. M. van der Aalst, "Supporting risk-informed decisions during business process execution," in *CAiSE'13*, Valencia, Spain, 2013, pp. 116–132.
- [25] O. Althuhova, R. Matulevicius, and N. Ahmed, "Towards definition of secure business processes," in *CAiSE Workshops*, 2012, pp. 1–15.
- [26] S. Wenzel, C. Wessel, T. Humbert, and J. Jürjens, "Securing processes for outsourcing into the cloud," in *CLOSER*, 2012, pp. 675–680.
- [27] B. Martens, M. Walterbusch, and F. Teuteberg, "Costing of cloud computing services: A total cost of ownership approach," in *ICSS'12*, 2012, pp. 1563–1572.