A General Approach for a Trusted Deployment of a Business Process in Clouds

Elio Goettelmann*†, Nicolas Mayer*, Claude Godart†
*CRP Henri Tudor
L-1855 Luxembourg-Kirchberg
{elio.goettelmann, nicolas.mayer}@tudor.lu
†Université de Lorraine
LORIA - INRIA Nancy - Grand Est
F-54500 Vandœuvre-lès-Nancy, France
claude.godart@loria.fr

ABSTRACT

It is recognized that the most important obstacle to the development of the cloud is the variety of new security threats which requests new methods and mechanisms. This is even truer for those who want to deploy business processes, because of the critical knowledge they encapsulate in terms of know-how and data. This paper proposes an approach combining modeling techniques and cloud selection for a trusted deployment of a security risk-aware business process in security constrained clouds.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous; D.2.9 [**Software Engineering**]: Management

General Terms

Management, Security, Design

Keywords

business process, cloud, security

1. INTRODUCTION

The business model of cloud computing can help organizations to reduce costs by outsourcing their information system (IS) assets. But related security issues are still preventing its broader adoption, until new techniques and solutions will be defined. Especially, existing business processes (BP) cannot be deployed in the cloud as they are because the new threats are endangering their critical knowledge more than ever. This paper proposes a comprehensive approach for a trusted deployment of business processes in clouds in three steps: requirements definition, BP remodeling and cloud selection. The goal of our approach is to categorize good practices in order to implement some of these techniques dynamically in a given business process.

MEDES'13 October 29-31, 2013, Neumünster Abbey, Luxembourg

It is organized as follows. Section 2 discusses our motivation; it starts with a motivating example, then situates our approach taking into account the cloud context with regards to traditional risk management in business process (BP) modeling, and finally overviews the steps of our methodology. Section 3.2 exposes the first step of the methodology, i.e. the BP security requirements engineering process. Section 4 explains the deployment of a business process in the cloud in two steps, i.e. first generating the cloud risk-aware BP logic and the security constraints compliant with the initial BP model and the security requirements, and then, based on these inputs, configuring the BP by assigning BP fragments to clouds. Section 5 illustrates our methodology on our motivating example. Section 6 discusses implementation issues and overviews our current deployment framework. Section 7 discusses the state of the art. Section 8 concludes.

2. MOTIVATIONS, APPROACH AND METHODOLOGY

This section introduces our work through a motivating example, gives the main principle of our approach and finally overviews the proposed methodology.

2.1 Motivating example: the shipping company

Let us take the case of a shipping company who has acquired a great reputation in the shipping of sensitive objects, with special characteristics such as high value, huge volume and/or dangerousness ... These goods characteristics are the criteria used by the company to organize the shipping (selection of shippers and/or of paths). The initial company process is depicted in figure 4A.

Now, this company wants to outsource this process on the cloud to save money. However, it is a little mistrustful because it fears revealing to cloud providers its knowhow, which is its real business value.

In the same way, as the company manages private data about its clients, it wants high level guarantees regarding data confidentiality. Especially, it wants that when data is supposed to be deleted, it is effectively, and cannot be found afterwards.

Finally, as it is very suspicious, even if cloud providers give guarantees, it wants also some capabilities to verify if the selected clouds deserve the trust placed by the company, and operate accordingly with promises.

Knowhow preservation, *Data confidentiality* and *Trust verification* are security objectives of the company.

One can claim that these objectives are general objectives

not so specific to BP and cloud. This is in some way true but the cloud context intensifies the need of *Knowhow preservation* and of *Trust verification*.

We use this example in the following to illustrate our methodology: the re-engineering of this BP to support its security objectives in the cloud is developed in section 5.

2.2 Approach

This section overviews our approach to manage the cloud characteristics extending traditional risk management for supporting the deployment of a BP in the cloud.

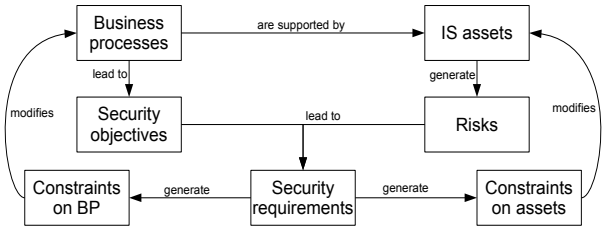

Figure 1: Classic risk management process

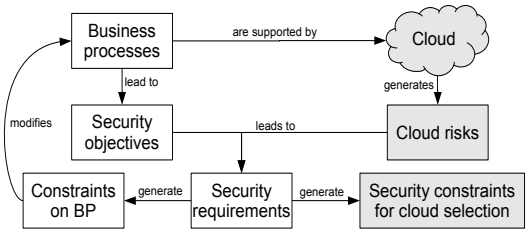

Figure 2: Risk management process with cloud computing

In a classic infrastructure, where the business processes of a company run locally (on-premises), the implementation of security objectives is yet complex but globally well understood. As shown in Fig. 1, business processes are supported by the resources controlled by the company (IS assets). These assets can have various vulnerabilities which, combined with different threats, lead to security risks for the company. From security objectives and risks, security requirements are deduced for defining how to manage these security risks.

Basically, the treatment of risks is achieved in four ways: retention, transfer/sharing, reduction or avoidance [20]. Traditionally, these are mainly implemented, on the one hand by constraints on the BP model (BP evolution for supporting security requirements), and on the other hand by constraints on the assets.

In a cloud-based infrastructure, where business processes can run remotely (off-premises), the implementation of security objectives is different, as depicted in figure 2. In fact, if security risks can still be avoided by changing the business process, the cloud cannot be constrained as easily (if even possible) as local assets, because it cannot be controlled by the company itself.

As a consequence, cloud specific security risks, as explained by ENISA [15], cannot be managed simply. However, if the customer company cannot control the cloud risks, it has the power to select the cloud providers which better fulfill its security requirements. Of course, this supposes that providers expose different and negotiable guarantees regarding security requirements (expressed in SLA), and that these guarantees can be compared based on reliable metrics, maybe provided by trusted third parties.

These principles are discussed in the methodology overviewed below.

2.3 Methodology

Our security management methodology, depicted in figure 3, consists in three main steps:

1. First, based on security objectives and cloud risks, a set of security requirements is defined.
2. Then, based on these security requirements, on the one hand the logical description of the process (BP model) is modified, and on the other hand, a set of security constraints on cloud properties is defined.
3. Based on this cloud risk-aware BP model and security constraints, the BP is configured, i.e. BP tasks/fragments are assigned to clouds.

This paper is mainly concerned with points 1) and 2) but 3) is partly discussed in section 6.

In fact, in line with the traditional BP process modeling methodologies, we have split security requirements in three parts[1]:

1. The first set addresses the BP logical level, making evolve the BP logic to integrate security objectives. This is deepened in section 3.2.1.
2. The second set addresses the BP organizational level, defining some constraints on the organization of clouds supporting the BP deployment. This is deepened in section 3.2.2.
3. The third set addresses the informational level, defining security constraints on cloud properties, clouds in which deploying the BP. This is deepened in section 3.2.3.

The first and second sets of requirements act on the BP model making it evolve (section 4.1) for avoiding some risks.

The second and third sets imposes constraints on cloud properties for supporting the selection of an optimized and secure cloud configuration (section 4.2) for reducing risks.

3. BP SECURITY REQUIREMENTS ENGINEERING

The first step of the methodology is to elucidate security requirements, i.e. to conceptually decide how to treat risks.

The security requirements engineering process rests on four main inputs: the initial business process model designed as if in a traditional local context, security objectives, cloud risks and cloud offerings: this section starts discussing them.

Then, the section describes the requirements engineering process.

3.1 Security requirements engineering inputs

3.1.1 Initial BP model

The initial BP model is a traditional BP model defined using a traditional notation (typically BPMN[2]) without the cloud perspective in mind. Such a model can contain special tasks specific to security, but designed in the context of a traditional execution setting as in [6].

[1]As in the definitions of the Workflow Management Coalition (`http://wfmc.org`)

[2]Business Process Model and Notation: http://www.bpmn.org

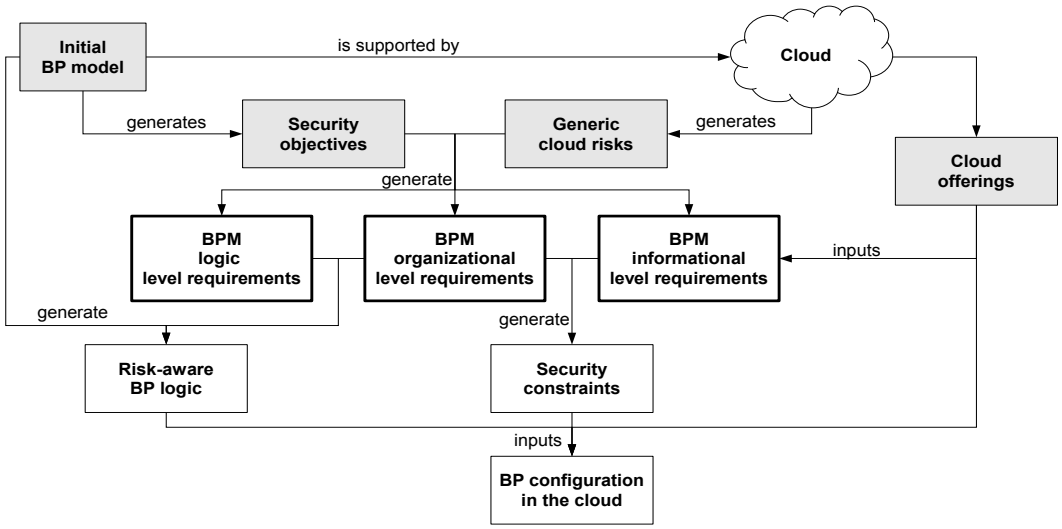

Figure 3: Methodology overview

Figure 4A is an example of such an initial BP model.

3.1.2 Security objectives

As introduced in section 2.3, security objectives of BP modelers are not so different when processes execute in the cloud than when executing on premises, and even less than when executing in a Web Service Oriented Architecture. However, in the cloud context:

- Some objectives are much more critical (for example *knowhow preservation*).
- And the way to reach these objectives in practice requests a new approach and security objectives implementation needs rethinking.

Traditionally, security objectives are defined in terms of Confidentiality, Integrity and Availability (the CIA triad). Then each can be refined to enlighten special aspects like accountability, authenticity, non-repudiation ... In the context of business processes, security objectives apply to BP assets, especially:

- the whole process itself
- fragments of the process
- tasks (single fragments)
- business objects (information, data)

For example, the *Knowhow preservation* security objective introduced above, which is very sensitive when outsourcing a BP in the cloud, can concern the whole process confidentiality, while a *Confidentiality of data* objective can concern only one business object.

To illustrate our discourse in the following, in relation with the example introduced in section 2.1, we will consider three requirements among the more critical security objectives in the context of the cloud:

- *Knowhow preservation*
- *Trust verification*
- *Data confidentiality*

Of course, these objectives are rather high level and can be refined in more fine grained objectives. Nevertheless, they are good representatives of our problematic.

3.1.3 Security risks specific to cloud computing

According to different studies [15], [14], [21], [27], cloud computing has new kinds of risks that companies never had to deal with before.

While no taxonomy has emerged as "the" reference, several works seems us good references. Among them, we can cite the ENISA report [15] which lists 35 risks related to cloud computing. They are organized in 4 main categories:

Policy and organizational risks

- Example 1, *Lock-in*: when switching costs are too high, the customer is unable to use another solution and becomes dependent on a specific provider.
- Example 2, *Loss of governance*: changes in the terms and conditions of a service may lead to a loss of compliance to the security requirements.

Technical risks

- Example 3, *Isolation failure*: failures of mechanisms separating cloud computing environments can lead to loss of sensitive data, reputation damage or service interruption.
- Example 4, *Ineffective deletion* of data, which may be available to malicious parties beyond the lifetime specified in the security policy.

Legal risks

- Example 5, *Subpoena*: when governments can confiscate physical hardware with shared tenancy, there is a higher risk of data disclosure for cloud customers.
- Example 6, *Change of jurisdiction*: an unpredictable legal framework increases the exposure to law enforcement measures which can be in some case unacceptable.

Risks not directly specific to the cloud

- Example 7, *Network breaks*: clouds are accessed through an internet connection, so this risk remains high and must still be considered.
- Example 8, *Natural disasters*: redundancy of data centers and multiple network paths should considerably lower this risk compared to traditional infrastructures.

This classification is representative of the different works as the different proposals do not differ drastically from this one, but rather tackle risks at a different level of abstraction. More importantly, they globally agree that a generic list of relevant risk can be established at a given granularity level and that there will not be different kinds of risks to consider depending on the provider. However, if providers present the same nature of risk, they do not respond to the risk in the same way and with the same accuracy; different providers

can naturally have different ratings for a given risk.

Also, depending on the security objectives, one risk will be more or less important to consider. This also will impact the choice of providers by consumers. For example, the *Subpoena*, *Change of jurisdiction* and *Isolation failure* risks impact the *Confidentiality of data* objective, while the *Distributed denial of service attack* risk seems not to.

Thus, it is on the responsibility of the consumer, with the help of a methodology and tools to choose the providers who better answer its needs. In this objective, consumers base their decision on the cloud offering (see next section).

3.1.4 Cloud offering

From the customer point of view, a cloud provider is a black box with some characteristics, and how cloud risks are treated technically is not really important. But how a provider responds to a cloud risks, and at which level of quality, are important in customer decision, especially to compare and select the best one.

In this objective, two types of contribution can help:

- First, cloud application interfaces and especially the part dedicated to risk treatment.
- Secondly, cloud metrics for measuring cloud quality and comparing cloud between them.

Cloud API.

A good way to publish interfaces is UDDI-like directories, as they are yet widely used for the selection of the functional part of services assigned to tasks. Such directories [3], [4] or [5] exist, allowing cloud providers to be sought and compared by service type, costs or reviews.

Unfortunately, while security seems an important selection criterion, it is only superficially considered in these repositories. To fill this gap, some ongoing studies can be used directly or indirectly by cloud customers. In this area, the Cloud Security Alliance (CSA), the Security, Trust and Assurance Registry (STAR) [12] is "a free, publicly accessible, registry that documents the security controls provided by various cloud computing offerings". In combination with the Cloud Control Matrix (CCM), which is a list of controls to fulfill a set of security requirements [13], the CSA has defined the *Consensus Assessments Initiative (CAI) Questionnaire* [11]: by answering a list of questions with "yes" or "no", a cloud provider can perform a self-assessment of its security controls. By downloading this questionnaire, a customer can get valuable information for selecting the most appropriate provider.

In the same way, a customer can use the ENISA Information Assurance Framework (IAF) [15], which proposes a similar questionnaire for providers about their infrastructures, controls or procedures. A CISCO study which considers the treatments to respond to the top ten cloud risks of the OWASP list [27] can be also used. But both systems are lesser mature than CSA STAR.

In any cases, these resources are mainly developed for helping providers to define good risk management policies, and thus are not really customer-oriented. However, customers can use them to organize their selection process by estimating cloud providers ability to lower risks, maybe with the help of third party audits or certifications.

Cloud metrics.

While such metrics seem valuable elements for cloud interfaces definition and comparison, it is clear that such metrics are currently few and incomplete. However, some initial frameworks compare cloud providers with respect to quality objectives and especially security objectives, generally contributing specific metrics.

For example, the ongoing Common Assurance Maturity Model (CAMM) has for objective a scorecard on which the provider has a grade (between 1 and 6 for example [8]) for each of the following control areas: Governance (GR), Human Resources (HR), Physical (PHY), IT Services (IT), Incident Management (IM) and Business Continuity (BM).

In the same way, the Eurocloud Star Audit proposes to certify cloud providers with regards to different security requirements. Like a ranking system, a provider can be certified from 1 up to 5 stars. To be certified on a certain level, all requirements of the levels below must be fulfilled. Examples of requirements are:

- Data deletion at the end (1 star)
- Interfaces, API, exports formats (2 stars)
- Possibility to select place of jurisdiction (3 stars)
- Providing VPN access (4 stars)
- Spread out data centers (5 stars)

These metrics are very generic, but as explained before, a cloud customer does not need to know exactly which controls or procedures are in place to protect its customer data, and a coarse granularity as considered may be enough.

3.2 Security requirements categorization and engineering

As introduced in section 2.3, we have organized security requirements, induced by security objectives and cloud risks, in three sets corresponding to the three traditional dimensions of BP modeling, i.e. requirements at the BP logical level, requirements at the organizational level, and finally requirements at the informational level.

3.2.1 Requirements at the BP logical level

Traditionally, the BP logic expresses the synchronization of the tasks to execute for achieving the process objective, following its control flow. The control flow uses task states and pertinent data to navigate in the BP model. The BP logic is generally expressed in a BP modeling language (BPMN is probably the more largely used at the design level), BPEL[3] is a good reference in the context of SaaS (Software as a Service) for modeling Web Service orchestration and choreographies. Figure 4A and 4B are examples of process model depicted in the BPMN language. Of course, it is not in our objective to define a new BP modeling language, but to contribute modeling principles to integrate security objectives in process models expressed in a traditional modeling language.

We consider these main principles for translating security objectives into security requirements at the BP logical level:

- The **Split knowledge into several BP fragments** principle is especially useful for managing confidentiality at the BP level, typically for *knowhow preservation* which is a highest level of requirement, if not the highest, for BP in the cloud context. The idea is to split the knowledge in several pieces in the objective of assigning resulting task/fragments to different clouds. In such a way, each cloud has only a partial view of the

[3]Business Process Execution Language: https://www.oasis-open.org/committees/wsbpel

process. Only the root BP fragment which assumes the integration of contributing BP fragments has a global view of the process: it can be maintained on premises or assigned to a highly trusted cloud. This is achieved by splitting tasks/BP fragments in several tasks/BP fragments and adding the corresponding control flow.

- The **Separate logic and data** principle is especially useful for managing confidentiality at the logical level, typically for hiding some relationships between data and tasks which represent an important part of the *knowhow to be preserved*. This principle is in some way yet supported by the separation of logical and informational levels in the BP modeling process, but the cloud context can request to enhance this property by splitting a task/BP fragment for introducing new tasks specific for data management.
- The **Group knowledge pieces into one BP fragment** principle is especially useful for enhancing *data integrity* by putting in the same place the more valuable artifacts so that it is easier to watch them. This is achieved by grouping several tasks/BP fragments in one task/BP fragment and adding the corresponding control flow.
- The **Replicate fragments** principle is especially useful to *verify* that clouds operate as promised (by comparing results and performance), and *that they deserve the trust put in them*. This is achieved by replicating task(s)/fragment(s) and adding task(s) to synchronize replicas. Replication can also be used to support availability but this is not central to our purpose.
- The **Add security management tasks** principle is to add nonfunctional tasks dedicated to security objectives in the cloud context. We do not think that these tasks are specific to the cloud and a taxonomy as this is defined in [18] can be reused. However, new needs, not anticipated in the initial BP model, can emerge due to the cloud context. For example, anonymization tasks, as introduced in our motivating example, can be defined to enhance *data confidentiality*, logs management ones to support the verification of conformance of cloud executions, and others to compare the performance of replicas.

3.2.2 Requirements at the BP organizational level

In traditional BP settings, the BP organizational level defines for each task, the role (capacity) requested to execute the task. It also defines task assignment rules for constraining resource allocation (like separation/binding of duties), and it assigns tasks to organizational units (swimlanes).

This remains in the cloud context, but the cloud itself can be constrained by organizational rules for achieving security objectives. Especially, these new rules provide security requirements for the cloud selection in the BP configuration process. They can be directly connected to requirements at the logical level: for example it seems a good practice to assign two BP fragments, required to be separated for preserving knowledge, to two different clouds.

We list here a representative, but not exhaustive, set of such rules:

- A **Separation of knowledge** rule imposes two process fragments to execute in two different clouds. As its name indicates, its objective is to fragment knowledge and it is especially useful to support *Separation of knowledge* and *Separation of logic and data* decisions taken at the logical level.
- A **Co-location of knowledge** rule imposes two process fragments to execute in the same cloud. As its name indicates, its objective is to group knowledge and it is especially useful to support a decision to group knowledge taken at the logical level.
- An **Impose retention of knowledge at premises** rule imposes a BP task/fragment to execute at premises, typically because it is a highly critical task/fragment.

3.2.3 Requirements at the BP informational level

In traditional BP settings, the BP informational level introduces technical choices for assigning an implementation to each task. Either a task is explicitly linked here to a particular implementation (tool), or its linkage is deferred to execution, but in such a case technical constraints can restrict this linkage.

This remains in the cloud context, but technical choices must integrate security requirements with cloud properties.

At the informational level, a customer can constrain cloud selection in three main ways. She/he can:

- **Ban a given cloud** for executing a task/BP fragment because she/he does not trust it.
- **Impose a cloud** for executing a task/BP fragment, because she/he trusts it; she/he has good experience with this cloud, or simply, it has a very good reputation.
- **Impose a level of security** a cloud must have for executing a BP task/fragment. For example, in relation with metrics defined above, impose a minimum level of security a cloud must provide, either globally (for example, not less than 3 stars ranking in the Eurocloud Star Audit system for the cloud implementing BP fragment x), or regarding a specific risk (for example, grade 3 for Governance in the Common Assurance Maturity Model).

See sections 5.1, 5.2 and 5.3 for the application of these principles to our motivating example.

4. TRUSTED DEPLOYMENT OF A BP IN THE CLOUD

The trusted deployment of a BP is operated in two steps:

1. First, the cloud risk-aware BP logic is designed and a set of security requirements to restrict assignment of tasks/BP fragments to clouds is generated.
2. Then, based on the cloud risk-aware BP logic and constraints (among them security cloud constraints are an important part), the BP is configured, i.e. tasks/ BP fragments are assigned to clouds.

4.1 Cloud risk-aware BP logic

Using BP security requirements at the logical and organizational levels, the initial BP model is re-designed to manage task/BP fragments splitting, grouping, replication, and to include new added security management tasks, as requested by security requirements at the logical and organizational level.

Figure 4B depicts the *Shipping company* cloud risk-aware BP model designed accordingly to requirements defined in section 5 below with respect to guidelines in section 3.2

4.2 Security cloud constraints

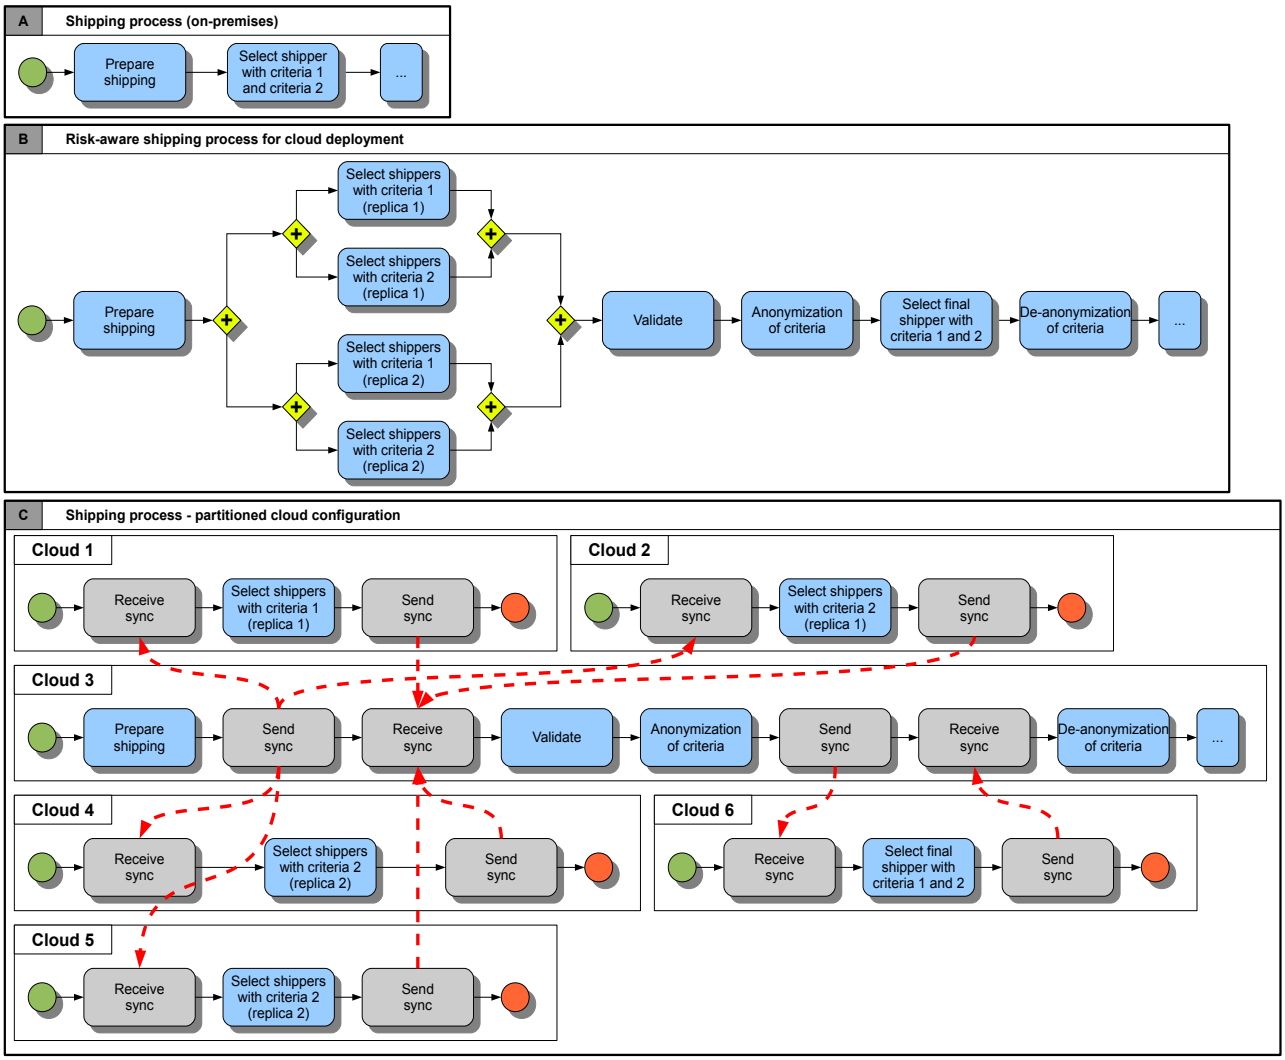

Figure 4: The 3 different phases of the shipping process: on-premises, risk-aware and final cloud configuration

Using BP security requirements at the organisational and informational levels and the cloud offerings, a set of constraints on cloud organization (assignment of tasks/BP fragments in different/same cloud(s)) and cloud properties (security level...) are generated.

A set of cloud constraints related to the *Shipping company* is given in section 5 below.

4.3 BP configuration

Its main objective is to assign tasks/ BP fragments to clouds for selecting an optimized cloud configuration answering customer requirements, involving security requirements.

Concretely, such a selection algorithm has to find the good equilibrium between different constraints, security constraints being one important input, but just one among others. Other requirements concern minimizing the cost of computing and the overhead of communication, and maximizing the overall quality of service [16].

In other words, security management is a part of an overall framework for optimizing the assignment of business process fragments to clouds, given a cloud offering. Concretely our cloud selection algorithm mixes security constraints with QoS constraints (maybe we should say that we consider security constraints as a special case of QoS constrains). We overview this framework in section 6. Figure 4C depicts the *Shipping company BP* deployment in clouds.

5. APPLICATION. THE SHIPPING COMPANY

In this section, we discuss the motivating example introduced in section 2. Studying the Shipping Company security objectives, i.e. *Knowhow preservation*, *Trust verification* and *Data confidentiality*, discussed above, the following decisions have been taken at the logical level.

5.1 Security requirements at the logical level

As illustrated in figure 4B, the logic has evolved as follows:

- The activity *select shipper with criteria 1 and 2* has been split into 3 activities: *select shipper with criteria 1*, *select shipper with criteria 2* and *select final shipper with criteria 1 and 2*.
- Activities *anonymization* and *de-anonymization* have been introduced for hiding how the company use criteria for optimizing shipping.
- Activities *select shipper with criteria 1* and *select shipper with criteria 2* have been duplicated.
- *and-split* gateways have been added for managing new activities.

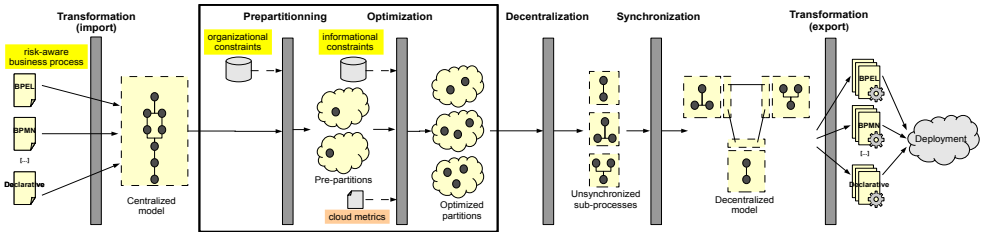

Figure 5: Overview of the implementation

In such a way, a cloud either computes partial data (only one criterion, thus preserving the company *knowhow*) or anonymous data (contributing to *data confidentiality*). Duplication of *select shipper with criteria 1* and *select shipper with criteria 2* allows to compare the quality of results of each duplicate, and as a consequence to evaluate the good functioning of the corresponding clouds (*trust verification*).

5.2 Security requirements at the organizational level

Considering the security objectives, the following decisions have been taken at the organizational level:

- Separate in different clouds all *select shipper with criteria x* tasks to promote separation of knowledge.
- Separate *select shipper with criteria 1 (replica 1)* and *select shipper with criteria 1((replica 2)*, and *select shipper with criteria 2 (replica 1)* and *select shipper with criteria 2 ((replica 2)*, to support comparison of providers.
- Co-locate in the same cloud *Prepare Shipping*, *Anonymization*, and *De-anonymization* as a way to localize critical data treatment and maintain it on-premises.
- Separate *Prepare Shipping*, *Anonymization*, and *De-anonymization* cloud from other clouds to promote separation of logic and data.

5.3 Security requirements at the informational level

This level restricts the choice of clouds regarding the services they can render and at which level of security.

In the shipping company:

- Ban clouds which do not guarantee effective data deletion.
- Impose a specific cloud for task *select shipper with criteria 1 and 2* because it is an excellent specialist of path optimization that the company trusts.
- Select a not less 3 stars ranking in the Eurocloud Star Audit system for any cloud.
- Select a not less 5 stars ranking in the Eurocloud Star Audit system for the cloud in charge of the *Select final shipper with criteria 1 and 2* task.

Figure 4C depicts the deployment of our motivating example in four clouds. All clouds are supposed to have a ranking better than 2 and cloud 6 better than 4 in the Eurocloud Star Audit system.

6. IMPLEMENTATION AND VALIDATION.

Implementing our approach needs to face two problems:

- selection of clouds respecting the security requirements
- BP decentralization in the selected cloud configuration

Our solution is mainly based on two of our previous work. In the first [16], we have developed a tool to decentralize a BP in a SaaS context by transforming orchestrations into choreographies, while minimizing communication costs and maximizing QoS. In the second [17], this tool was extended to support the objective of preparing processes for deployment on the cloud with yet a first but limited set of security constraints (co-locate, separate). It is currently being extended to introduce additional security constraints as requested above.

Concretely, security management is a part of an overall framework for optimizing the assignment of business process fragments to clouds. In other words, our cloud selection algorithm mixes security constraints with other constraints for minimizing the cost of computing, the overhead of communication, and maximizing the overall quality of service.

The current tool architecture is depicted in figure 5.

In the *Transformation* phase, the tool takes in input a risk-aware BP model, as the one described in figure 4B, alternatively defined in BPMN, BPEL or JSON, and transforms it into an internal graph structure.

During *Pre-partitioning*, requirements of the organizational level, as depicted in section 3.2, are used to generate pre-partitions. A first set of partitions is built based on *co-locate*, *separate* and *impose* requirements.

The remaining tasks are then added to the pre-partitions with an algorithm which tries to *Optimize* the Quality of Service of the entire process. The requirements of the informational level are used during this phase as constraints. We have integrated a 1-dimensional provider scoring system, similar to the Eurocloud Star Audit, to rate the potential providers. Security levels are assigned to tasks, which can only be deployed on a provider if its security level is greater or equal to that of the task.

Afterward, the control and data flow of each partition is built (BP *Decentralization*) and the different fragments are *Synchronized* with respect to the original BP model. The obtained process can finally be *exported* in an output format like BPMN, BPEL or dotGraph.

We have tested our tool with different processes and generated deployment ready BPEL files. These files have been successfully deployed and executed in a hybrid cloud environment (1 private and several public clouds). The private cloud consisted in a local ApacheODE platform [1], whereas WSO2 Stratoslive [2] was chosen as public cloud.

7. RELATED WORK

We use the terminology introduced in the Information System Security Risk Management (ISSRM) reference model [24]. Based on this work, [19] presents a framework for security requirements elicitation but this remains at a general information system level.

A lot of work has already been done to model security requirements/goals in the context of business process modeling, especially with application to the BPMN notation ([23], [30], [25], [28], [26]). But most of these works do not specify

how to obtain the different security requirements and target only a minimal set of security aspects in business processes. More complete, [6] aligns the BPMN constructs with the ISSRM domain model to allow modelers to express secure assets, risks and risk treatment using BPMN, but this work does not consider the cloud context.

More close to the cloud context, some work consider the security dimension in the context of web service composition ([7], [9]), but in general do not consider the BP globally, and are rather concerned with matching security requirements with service capabilities than elucidating security requirements. [10] is good representative of this category. In this vein, we have proposed heuristics for composite web services decentralization optimized with QoS constraints [16]. In [17], we have shown that this approach can be extended to the cloud context, but with a very restricted security model and without methodological consideration.

Directly related to our context, [22] tackles the cloud security challenges by adopting multiple clouds and introduces the concepts of replication and fragmentation used in our paper, but remains at the level of principles. Another input of our work is [29] which proposes a multi-level security model to partition workflows over federated clouds, but considers simple workflows and is not concerned with design aspects.

8. CONCLUSION AND FUTURE WORK

In this paper we have proposed a comprehensive approach for a trusted deployment of a BP in clouds. It is based on the one hand on risk-aware modeling techniques and on the other hand on security constrained cloud selection. While the work is still ongoing and the current state of art can hardly support experimentations on real case studies, we must content us with such demonstrators and simulations as presented here to demonstrate the technical feasibility of the approach.

We are currently developing a taxonomy work to draw a more complete and precise relationship between cloud risks, security objectives and security requirements, including a multi-dimensional scoring system to allow a more precise evaluation and selection of cloud providers. In application, we are working on the automation of a BP deployment and execution among a pre-configured set of cloud providers.

9. REFERENCES

[1] Apache orchestration director engine. `http://ode.apache.org/`.

[2] WSO2 Business Process Server. `http://wso2.com/products/business-process-server`.

[3] `http://www.clouddir.com`, 2013.

[4] `http://www.cloudbook.net/directories/product-services/cloud-computing-directory`, 2013.

[5] `http://cloud-computing.findthebest.com/`, 2013.

[6] O. Altuhhova, R. Matulevicius, and N. Ahmed. Towards definition of secure business processes. In *CAiSE Workshops*, 2012.

[7] M. Anisetti, C. A. Ardagna, E. Damiani, and F. Saonara. A test-based security certification scheme for web services. *ACM Trans. Web*, 2013.

[8] Brian Honan. Is it important to know how secure your cloud is ? `http://www.common-assurance.com/resources/Camm-Dublin-7th-Oct-2010.pdf`, 2012.

[9] J. Cardoso, T. Binz, U. Breitenbücher, O. Kopp, and F. Leymann. Cloud computing automation: Integrating usdl and tosca. In *CAiSE*. Springer-Verlag, 2013.

[10] B. Carminati, E. Ferrari, and P. C. K. Hung. Security conscious web service composition. In *ICWS*, 2006.

[11] Cloud Security Alliance. Cloud Security Alliance - Consensus Assessments Initiative Questionnaire. `https://cloudsecurityalliance.org/research/cai/`.

[12] Cloud Security Alliance. Cloud Security Alliance - Security, Trust & Assurance Registry. `https://cloudsecurityalliance.org/research/initiatives/star-registry/`.

[13] Cloud Security Alliance. CSA - Cloud Control Matrix. `https://cloudsecurityalliance.org/research/ccm/`.

[14] Cloud Security Alliance. CSA - Top Threats to Cloud Computing v1.0. `https://www.cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf`, 2010.

[15] European Network and Information Security Agency. Benefits, risks and recommendations for information security. `http://www.enisa.europa.eu/activities/risk-management/files/deliverables/cloud-computing-risk-assessment/at_download/fullReport`, 2009.

[16] W. Fdhila, M. Dumas, C. Godart, and L. García-Bañuelos. Heuristics for composite web service decentralization. *Software & Systems Modeling*, to appear 2013.

[17] E. Goettelmann, W. Fdhila, and C. Godart. Partitioning and cloud deployment of composite web services under security constraints. In *IEEE International Conference on Cloud Engineering (IC2E)*, 2013.

[18] N. Gonzalez, C. Miers, F. Redigolo, T. Carvalho, M. Simplicio, M. Naslund, and M. Pourzandi. A quantitative analysis of current security concerns and solutions for cloud computing. In *CLOUDCOM*, 2011.

[19] C. Haley, R. Laney, J. Moffett, and B. Nuseibeh. Security requirements engineering: A framework for representation and analysis. *IEEE Trans. Softw. Eng.*, 2008.

[20] ISO, International Organisation for Standardisation, Geneva, Switzerland. ISO/IEC 27005: Information technology - Security techniques - Information security risk management, 2008.

[21] Jeff Beckham. The Top 5 Security Risks of Cloud Computing. `http://blogs.cisco.com/smallbusiness/the-top-5-security-risks-of-cloud-computing/`, 2011.

[22] M. Jensen, J. Schwenk, J.-M. Bohli, N. Gruschka, and L. L. Iacono. Security prospects through cloud computing by adopting multiple clouds. In *CLOUD*, 2011.

[23] K. Knorr and S. Röhrig. Security requirements of e-business processes. In *Proceedings of the IFIP Conference on Towards The E-Society: E-Commerce, E-Business, E-Government*, 2001.

[24] R. Matulevičius, N. Mayer, H. Mouratidis, E. Dubois, P. Heymans, and N. Genon. Adapting secure tropos for security risk management in the early phases of information systems development. In *CAiSE*, 2008.

[25] E. Paja, P. Giorgini, S. Paul, and P. H. Meland. Security requirements engineering for secure business processes. In *BIR Workshops*, 2011.

[26] A. Rodríguez, A. Caro, C. Cappiello, and I. Caballero. A bpmn extension for including data quality requirements in business process modeling. In *BPMN*, 2012.

[27] Shankar Babu Chebrolu and Vinay Bansal and Pankaj Telang. Top 10 cloud risks that will keep you awake at night. `https://www.owasp.org/images/4/47/Cloud-Top10-Security-Risks.pdf`.

[28] S. H. Turki, F. Bellaaj, A. Charfi, and R. Bouaziz. Modeling security requirements in service based business processes. In *BPMDS/EMMSAD*, 2012.

[29] P. Watson. A multi-level security model for partitioningworkflows over federated clouds. In *CLOUDCOM*, 2011.

[30] C. Wolter, M. Menzel, and C. Meinel. Modelling security goals in business processes. In *Modellierung*. Köllen, 2008.