Privacy-Aware Scheduling for Inter-Organizational Processes

Christoph Hochreiner

Distributed Systems Group, Vienna University of Technology, Austria
c.hochreiner@infosys.tuwien.ac.at

Abstract. Due to the increasing specialization of companies in a globalized world, inter-organizational process enactments have become increasingly relevant in recent years. Nevertheless there are hardly any scheduling approaches that meet the requirements of these inter-organizational processes, especially in terms of privacy aspects. In this paper we present a privacy-aware scheduling approach for hybrid clouds, which represents a vital starting point to design a holistic execution environment for inter-organizational process enactments.

Keywords: Cloud Computing, Business Process Management, Hybrid Clouds

1 Introduction

In the last couple of years Business Process Management (BPM) has become a well-adopted approach for companies to provide value-added services to customers [8]. Business processes are composed of software- as well as human-based services and their design ranges from simple sequences to complex structures involving loops, splits or choices [11]. The process enactment is conducted by a Business Process Management System (BPMS) [12] which is considered as a generic software system that manages operational business processes [2]. The management of operational business processes covers the assignment of the different process steps to the designated services which are required to realize a process enactment and schedules their instantiation. Apart from the process scheduling, a BPMS may also manage the provisioning of computational resources to instantiate the software-based services. Since BPMS are used to execute business processes, the BPMS as well as the services are often deployed on fixed resources within the company's premises, where the companies may combine their computational resources to implement a resource pool, i.e. a private cloud [7]. The most important reason for this internal hosting solution are security and privacy restrictions, since services may deal with sensitive information, e.g., health data or execute algorithms that are considered as trade secrets [10].

This paper proposes a privacy-aware scheduling approach for hybrid cloud environments. To obtain an optimal and privacy-aware scheduling respectively resource provisioning approach, we extend the Service Instance Placement Problem (SIPP) [4] which applies Mixed Integer Linear Programming (MILP).

The remainder of this paper is structured as follows: In Sect. 2 we state the motivation for our work and discuss some preliminaries in Sect. 3. We further present our privacy-aware scheduling approach in Sect. 4 and Sect. 5 concludes the paper with an outlook on our future work.

2 Motivation

Business process enactments are usually triggered by process requests. These process requests are issued by external events, e.g., customer interactions, which lead to alternating amounts of business process requests respectively changing resource requirements. In peak-times, when external events issue an extraordinary amount of process requests, a BPMS may run into an underprovisioning scenario, since there are not enough resources to enact the process requests according to their Service Level Agreements (SLAs) [9]. This leads to a lower Quality of Service (QoS), e.g., longer response times and SLA violations may also trigger penalty cost that increase the overall cost for process enactment. Besides the peak-times, a system with fixed resources is also likely to run into overprovisioning scenarios, since the computational resources will not be used adequately. This leads to economically inefficient cost structures for the companies.

Public clouds, e.g., Amazon EC2, offer a promising solution to the resource usage challenges for varying process requests. A cloud-aware BPMS is able to obtain the required resources *on demand* in an utility like fashion. This enables the BPMS to obtain *resource elasticity* by scaling the computational resources up and down, based on the changing requirements. *Measured services* further allow an exact billing of the computational resources based on the actual resource usage [7]. This elastic resource provisioning strategy avoids underprovisioning scenarios, since the public cloud provides enough resources to cover the peak-requirements. A cloud environment also avoids overprovisioning scenarios, because not required resources can be released as soon as they are not needed any more.

Besides the resource allocation there are also other challenges for BPMS, like privacy issues for service instantiations of inter-organizational processes, i.e., service choreographies. Inter-organizational processes are structured similarly to business processes. The major difference is that their process steps are assigned to software services which are provided by different companies instead of only one. Therefore software services for inter-organizational processes can be executed on a community cloud [7]. Nevertheless this common execution environment is not acceptable for some software services due to privacy restrictions. The most promising approach to tackle these issues is the creation of a hybrid cloud which consists of a community cloud and dedicated private clouds for each company [3]. Although resource scheduling for hybrid clouds already raised some attention in terms of scheduling [1] as well as privacy aware deployments [13], there are surprisingly little efforts towards privacy-aware scheduling approaches for BPMSs [6].

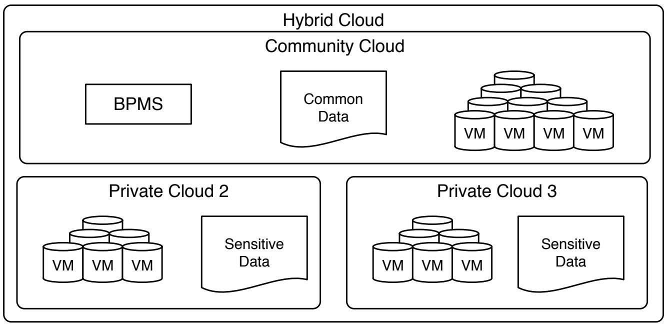

Fig. 1: Cloud Landscape

3 Preliminaries

In our previous work we presented the Service Instance Placement Problem (SIPP) [4], which provides a cost-optimized scheduling and resource provisioning plan for multiple parallel process enactments. SIPP represents a multi-objective scheduling strategy which is presented in Sect. 4.1. Up to now the SIPP only considers a single cloud for process enactments. Therefore it does not consider any security nor privacy related aspects which are relevant for process enactments in hybrid cloud environments. Before we describe the privacy related concepts in detail in Sect. 4.2, we define some preliminaries.

The execution environment for inter-organizational processes consists of a a community cloud and dedicated private clouds for privacy sensitive services, as illustrated in Fig. 1. The community cloud hosts all non privacy sensitive services as well as the BPMS. The BPMS schedules process steps, provisions resources for the software-based services on the community cloud and also triggers the deployment of the privacy sensitive services in the dedicated private clouds, based on the privacy restrictions issued for the services. These restrictions are described in detail in Sect. 4.2. In terms of computational resources, we assume that the private clouds offer a limited amount of computational resources, which are only sufficient to run the privacy sensitive services whereas the community cloud offers theoretically unlimited resources.

The inter-organizational processes are composed of multiple process steps that represent the software-based services provided by the participating companies. Fig. 2 represents an exemplary inter-organizational process, which shows the collaboration among 3 companies. Step 3 and 5 are annotated as privacy sensitive and must only be executed in the dedicated private clouds, whereas all other steps can be executed in the community cloud. To execute a process step, the BPMS triggers the deployment of the software-based service on a Virtual Machine (VM) either on the community cloud or on a private cloud. This deployment results in a service instance that can be invoked by the BPMS to execute the service and therefore execute the process step.

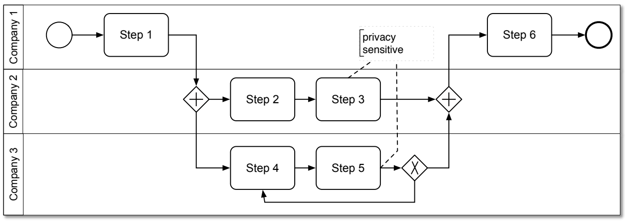

Fig. 2: Inter-organizational Process Incorporating 3 Companies

4 Privacy-Aware Scheduling

4.1 Service Instance Placement Problem

The SIPP is represented by a set of different constraints and equations, which are described in detail in [4]. In Eq. 1 the objective for the SIPP optimization model, i.e., the minimization of the overall execution cost, is shown. This objective comprises four terms, where the first term represents the overall leasing cost of the computational resources by summing up the amount of leased VMs $\gamma_{(v,t)}$ multiplied by their cost c_v. The second term shows the penalty cost, which arise, if a process is not finished within the given time. Hereby it sums up all delayed process instances $e^p_{i_p}$ and multiplies them with predefined penalty cost $c^p_{i_p}$. In order to keep the overall cost as low as possible, the optimization model penalizes idle resources (CPU ($f^C_{k_v}$) and RAM ($f^R_{k_v}$)) which are multiplied by the constants ω^C_f and ω^R_f. The last term is designed to prioritize process steps $x_{(j_{i_p},k_v,t)}$, so that steps with a closer deadline DL_{i_p} are executed first.

$$\begin{aligned}\min \sum_{v\in V} c_v \cdot \gamma_{(v,t)} + \sum_{p\in P}\sum_{i_p\in I_p} c^p_{i_p} \cdot e^p_{i_p} + \sum_{v\in V}\sum_{k_v\in K_v} (\omega^C_f \cdot f^C_{k_v} + \omega^R_f \cdot f^R_{k_v}) \\ - \sum_{p\in P}\sum_{i_p\in I_p}\sum_{j_{i_p}\in J^*_{i_p}} \frac{1}{DL_{i_p} - \tau_t} x_{(j_{i_p},k_v,t)}\end{aligned} \tag{1}$$

4.2 Privacy Extensions

Since SIPP is designed for a single cloud, we introduced additional constraints and additional SLA policies to enable the enactment of inter-organizational processes while respecting the privacy constraints of the services.

We extended the set of VM types to distinguish between different deployment locations with their different privacy policies.

$$V = \bigcup_{loc\in Loc} V_{loc} \tag{2}$$

$$K = \bigcup_{loc\in Loc} K_{loc} \tag{3}$$

This differentiation is possible by introducing the identifier loc ($loc \in Loc$), which represents the type of the cloud, e.g., community cloud or 1 of the private clouds. The new set of available VMs is then defined by the union of all VMs, which can be instantiated in the different clouds (Eq. 2). Analogously we also extended the set of all currently instantiated VMs (Eq. 3).

In terms of privacy restrictions, there are 2 specification possibilities to restrict the execution of services regarding the type of the cloud. The first approach is blacklisting: the SLA lists for every process all services, which must not be executed on specific clouds. The major downside of this approach is that the SLA needs to be updated when additional clouds are added.

The alternative approach is whitelisting, i.e., the SLA for each process lists all service instantiation possibilities in the different clouds. Since this SLA pursues a defensive permission approach, there is no need to update the SLA in contrast to the blacklisting approach, when the cloud environment grows by additional private or community clouds. Eq. 4 shows an exemplary SLA for the process presented in Fig. 2.

$$SLA_P = \begin{cases} \text{communityCloud} & \text{(Service 1, Service 2, Service 4)} \\ \text{privateCloud1} & \text{(Service 3)} \\ \text{privateCloud2} & \text{(Service 5)} \end{cases} \tag{4}$$

Based on this SLA for services, a MILP-solver for the optimization problem is able to generate the instantiation possibilities according to Eq. 5. The constraints in Eq. 5 evaluate whether a specific process step j_{i_p} can be instantiated on a specific VM $k_{v_{loc}}$ by querying whether the process step is listed on the whitelist for the given cloud loc. If the SLA does not explicitly allow the instantiation of the process step on the specific VM, the constraint rules out the deployment option. Otherwise the MILP-solver decides based on other constraints whether the service is deployed on the specific VM (1) or not (0) as stated in the alternative branch of the constraint.

$$x_{(j_{i_p},k_{v_{loc}},t)} = \begin{cases} 0 & \text{, if } j_{i_p} \notin SLA_{P_{loc}}, loc \in Loc \\ \{0,1\} & \text{, else} \end{cases} \tag{5}$$

5 Outlook

In this paper we focused on the formalization of privacy constraints for the deployment and enactment of inter-organizational processes in hybrid cloud environments. Although these privacy extensions are relevant for process enactment, they only represent a small step towards a holistic process scheduling and resource allocation approach for inter-organizational processes in hybrid cloud environments. Other relevant topics are data transfer aspects among the different clouds, which become increasingly relevant for big data applications or a cost-efficient resource allocation across the cloud environment. In our future work we will evaluate the proposed privacy constraints in a hybrid cloud environment based on the Vienna Platform for Elastic Processes (ViePEP) [5]. Here we plan to evaluate our approach against other privacy ensuring methods, e.g., encryption of privacy-sensitive data in terms of performance and cost-efficiency. Further we will also investigate other areas, like data transfer aspects or pricing policies for hybrid clouds to enable the enactment of inter-organizational processes in an economically efficient manner.

Acknowledgments This paper is partially supported by TU Vienna research funds and supported by the European Union within the SIMPLI-CITY FP7-ICT project (Grant agreement no. 318201).

References

1. van den Bossche, R., Vanmechelen, K., Broeckhove, J.: Cost-optimal scheduling in hybrid iaas clouds for deadline constrained workloads. In: 2010 IEEE 3rd International Conference on Cloud Computing (CLOUD). pp. 228–235. IEEE (2010)
2. van Der Aalst, W.M., Ter Hofstede, A.H., Weske, M.: Business process management: A survey. In: Business process management, pp. 1–12. Springer (2003)
3. Goyal, P.: Enterprise usability of cloud computing environments: issues and challenges. In: 2010 19th IEEE International Workshop on Enabling Technologies: Infrastructures for Collaborative Enterprises (WETICE). pp. 54–59. IEEE (2010)
4. Hoenisch, P., Schuller, D., Hochreiner, C., Schulte, S., Dustdar, S.: Elastic process optimization - the service instance placement problem. Tech. Rep. TUV-1841-2014-01, Distributed Systems Group, Vienna University of Technology (October 2014)
5. Hoenisch, P., Schulte, S., Dustdar, S., Venugopal, S.: Self-Adaptive Resource Allocation for Elastic Process Execution. In: 6th International Conference on Cloud Computing (CLOUD 2013). pp. 220–227. IEEE (2013)
6. Huang, Z., van der Aalst, W.M.P., Lu, X., Duan, H.: Reinforcement learning based resource allocation in business process management. Data & Knowledge Engineering 70(1), 127–145 (2011)
7. Mell, P., Grance, T.: The NIST Definition of Cloud Computing. Recommendations of the National Institute of Standards and Technology (2011)
8. Rosemann, M., vom Brocke, J.: The Six Core Elements of Business Process Management. In: Handbook on Business Process Management 1, pp. 107–122. Springer (2010)
9. Schulte, S., Janiesch, C., Venugopal, S., Weber, I., Hoenisch, P.: Elastic Business Process Management: State of the Art and Open Challenges for BPM in the Cloud. Future Generation Computer Systems 46, 36–50 (2015)
10. Subashini, S., Kavitha, V.: A survey on security issues in service delivery models of cloud computing. Journal of Network and Computer Applications 34(1), 1 – 11 (2011)
11. van der Aalst, W.M.P., ter Hofstede, A.H.M., Kiepuszewski, B., Barros, A.P.: Workflow Patterns. Distributed and Parallel Databases 14(1), 5–51 (2003)
12. Weske, M.: Business Process Management: Concepts, Languages, Architectures. Springer, 2nd edn. (2012)
13. Zhang, K., Zhou, X., Chen, Y., Wang, X., Ruan, Y.: Sedic: privacy-aware data intensive computing on hybrid clouds. In: Proceedings of the 18th ACM conference on Computer and communications security. pp. 515–526. ACM (2011)