

# Performance and Cost Evaluation for the Migration of a Scientific Workflow Infrastructure to the Cloud

Santiago Gómez Sáez, Vasilios Andrikopoulos, Michael Hahn,  
Dimka Karastoyanova, Frank Leymann, Marigianna Skouradaki, Karolina Vukojevic-Haupt  
*Institute of Architecture of Application Systems, University of Stuttgart, Stuttgart, Germany*  
{gomez-saez, andrikopoulos, hahn, karastoyanova, leymann, skouradaki, vukojevic}@iaas.uni-stuttgart.de

Keywords:

Workflow Simulation; eScience; IaaS; Performance Evaluation; Cost Evaluation; Cloud Migration

Abstract:

The success of the Cloud computing paradigm, together with the increase of Cloud providers and optimized Infrastructure-as-a-Service (IaaS) offerings have contributed to a raise in the number of research and industry communities that are strong supporters of migrating and running their applications in the Cloud. Focusing on eScience simulation-based applications, scientific workflows have been widely adopted in the last years, and the scientific workflow management systems have become strong candidates for being migrated to the Cloud. In this research work we aim at empirically evaluating multiple Cloud providers and their corresponding optimized and non-optimized IaaS offerings with respect to their offered performance, and its impact on the incurred monetary costs when migrating and executing a workflow-based simulation environment. The experiments show significant performance improvements and reduced monetary costs when executing the simulation environment in off-premise Clouds.

## 1 INTRODUCTION

In the last years the workflow technology has been widely adopted in several domains, e.g. business or eScience, which often have different domain-specific requirements in terms of supported functionalities and expected behavior of the underlying infrastructure. Focusing on eScience applications, simulation workflows are a well-known research area, as they provide scientists with the means to model, provision, and execute automated and flexible long running simulation-based experiments [Sonntag and Karastoyanova, 2010]. Such simulation-based experiments typically comprise large amounts of data processing and transfer and consume multiple distributed simulation services for long periods of time. Due to the access and resource consumption nature of such simulation environments, previous works have targeted the migration and adaptations of such environments to be deployed, provisioned, and executed in Cloud infrastructures [Juve et al., 2009, de Oliveira et al., 2011, Vukojevic-Haupt et al., 2013, Zhao et al., 2014].

The Cloud computing paradigm has led in the last years to an increase in the number of applications which are partially or completely running in different *Everything-as-a-Service* Cloud offerings. The increase of available and optimized Cloud services has introduced further efficient alternatives for hosting application components with special resources consumption patterns, e.g. computationally or memory intensive ones. However, such a wide landscape of possibilities has become a challenge for deciding among the different Cloud providers and their corresponding offerings. Previous works targeted such a challenge by assisting application developers in the tasks related to selecting, configuring, and adapting the distribution of their application among multiple services [de Oliveira et al., 2011, Sáez et al., 2014a]. There are multiple decision points that can influence the distribution of an application, e.g. cost, performance, security concerns, etc. The focus of this research work is to provide an overview, evaluate, and analyze the trade-off between the performance and cost when migrating a simulation environment to different Cloud

providers and their corresponding Infrastructure-as-a-Service (IaaS) offerings. The contributions of this work can therefore be summarized as follows:

- the selection of a set of viable and optimized IaaS offerings for migrating a previously developed simulation environment,
- an empirical evaluation focusing on the performance and the incurred monetary costs, and,
- an analysis of the performance and cost trade-off when scaling the simulation environment workload.

The remaining of this paper is structured as follows: Section 2 motivates this work and depicts the problems that aim to be achieved. The simulation environment used for evaluation purposes in this work is introduced in Section 3. Section 4 presents the experiments on evaluating the performance and incurred costs when migrating the simulation environment to different IaaS offerings, and discusses our findings. Finally, Section 5 summarizes related work and Section 6 concludes with some future work.

## 2 MOTIVATION & PROBLEM STATEMENT

Simulation workflows, a well-known topic in the field of eScience, describe the automated and flexible execution of simulation-based experiments. Common characteristics of such simulation workflows are that they are long-running as well as being executed in an irregular manner. However, during their execution a wide amount of resources are typically provisioned, consumed, and released. Considering these characteristics, previous works focused on migrating and executing simulation environments in the Cloud, as Cloud infrastructures significantly reduce infrastructure costs while coping with an irregular but heavy demand of resources for running such experiments [Vukojevic-Haupt et al., 2013].

Nowadays there exists a vast amount of configurable Cloud offerings among multiple Cloud providers. However, such a wide landscape has become a challenge for deciding among (i) the different Cloud providers and (ii) the multiple Cloud offering configurations offered by such providers. We focus in this work on IaaS solutions, as there exists a lack of Platform-as-a-service (PaaS) offerings that enable the deployment and execution of scientific workflows in the Cloud. IaaS offerings describe the amount and type of allocated resources,

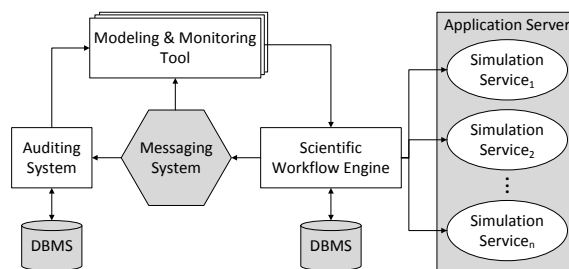


Figure 1: System Overview of the SimTech Scientific Workflow Management System (SWfMS)

e.g. CPUs, memory, or storage, and define different VM instance types within different categories. For example, the Amazon EC2<sup>1</sup> service does not only offer VM instances of different size, but also provides different VM categories which are optimized for different use cases, e.g. computation intensive, memory intensive, or I/O intensive. Similar offerings are available also by other providers, e.g. Windows Azure<sup>2</sup> or Rackspace<sup>3</sup>. The offered performance and incurred cost significantly vary among the different Cloud services, and depend on the simulation environment resource usage requirements and workload. In this work, we aim to analyze the performance and cost trade-off when migrating to different Cloud offerings a simulation environment developed and used as case study, as discussed in the following section.

## 3 THE OPAL SIMULATION ENVIRONMENT

A Scientific Workflow Management System (SimTech SWfMS) is being developed by the Cluster of Excellence in Simulation Technology (SimTech<sup>4</sup>), enabling scientists to model and execute their simulation experiments using workflows [Sonntag and Karastoyanova, 2010, Sonntag et al., 2012]. The SimTech SWfMS is based on conventional workflow technology which offers several non-functional requirements like robustness, scalability, reusability, and sophisticated fault and exception handling [Görlach et al., 2011]. The system has been adapted and extended to the

<sup>1</sup>Amazon EC2: <http://aws.amazon.com/ec2/instance-types/>

<sup>2</sup>Windows Azure: <http://azure.microsoft.com/en-us/>

<sup>3</sup>Rackspace: <http://www.rackspace.com/>

<sup>4</sup>SimTech: <http://www.iaas.uni-stuttgart.de/forschung/projects/simtech/>

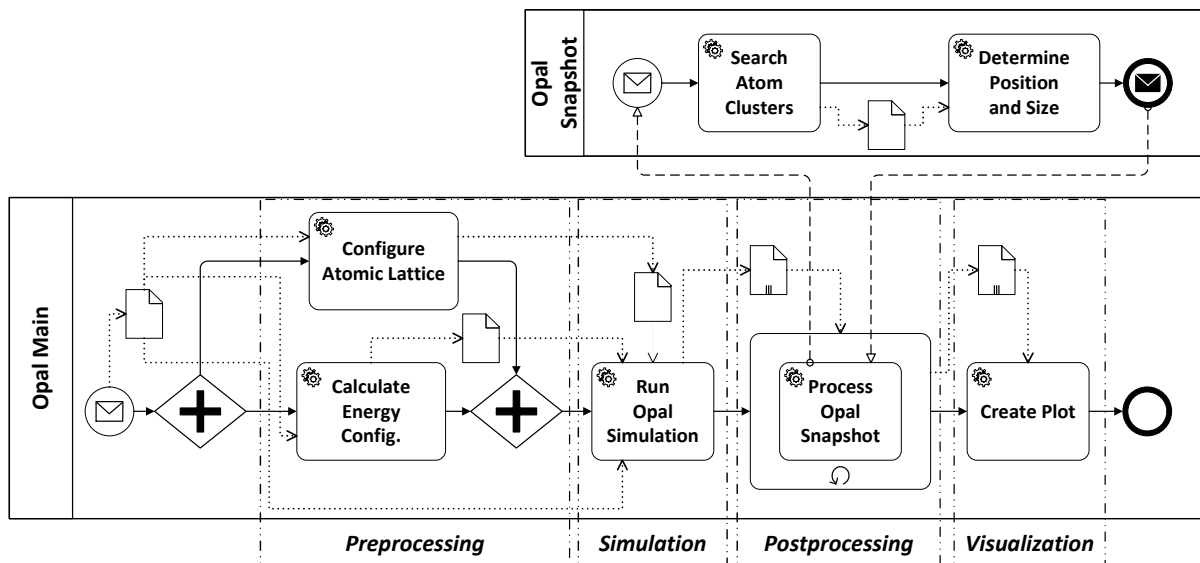


Figure 2: Simplified Simulation Workflows Constituting the OPAL Simulation Environment [Sonntag and Karastoyanova, 2013]

special needs of the scientists in the eScience domain [Sonntag et al., 2012]. During the execution of a workflow instance the system supports the modification of the corresponding workflow model, which is then propagated to the running instances. This allows running simulation experiments in a trial-and-error manner.

The main components of the SimTech SWfMS shown in Fig. 1 are a modeling and monitoring tool, a workflow engine, a messaging system, several databases, an auditing system, and an application server running simulation services. The workflow engine provides an execution environment for the workflows. The messaging system serves as communication layer between the modeling- and monitoring tool, the workflow engine, and the auditing system. The auditing system stores data related to the workflow execution for analytical and provenance purposes.

The SimTech SWfMS has been successfully applied in different scenarios in the eScience domain; one example is the automation of a Kinetic Monte-Carlo (KMC) simulation of solid bodies by orchestrating several Web services being implemented by modules of the OPAL application [Sonntag et al., 2011a]. The OPAL Simulation Environment is constituted by a set of services which are controlled and orchestrated through a main OPAL workflow (the *Opal Main* process depicted in Figure 2). The simulation services are implemented as Web services and divided into two main categories: (i) resource management, e.g. distributing the workload among the different servers, and (ii)

wrapped simulation packages depicted in [Binkele and Schmauder, 2003, Molnar et al., 2010]. The main workflow can be divided in four phases as shown in Fig. 2: preprocessing, simulation, postprocessing, and visualization. During the preprocessing phase all data needed for the simulation is prepared. In the simulation phase the workflow starts the Opal simulation by invoking the corresponding Web service. In regular intervals, the Opal simulation creates intermediate results (snapshots). For each of these snapshots the main workflow initiates the postprocessing which is realized as a separate workflow (*Opal Snapshot* process in Figure 2). When the simulation is finished and all intermediate results are postprocessed, the results of the simulation are visualized.

## 4 EXPERIMENTS

### 4.1 Methodology

As shown in Fig. 2, the OPAL Simulation Environment is comprised of multiple services and workflows that compose the simulation and resource management services. The environment can be concurrently used by multiple users, as the simulation data isolation is guaranteed through the creation of independent instances (workflows, services, and temporal storage units) for each user's simulation request. The experiments must therefore consider and emulate the usage of the envi-

ronment by multiple users concurrently.

The migration of the simulation environment to the Cloud opens a wide set of viable possibilities for selecting and configuring different Cloud services for the different components of the OPAL environment. However, in this first set of experiments we restrict the distribution of the simulation environment components by hosting the complete simulation application stack in one VM, which is made accessible to multiple users. Future investigations plan to distribute such environment using different Cloud offerings, e.g. Database-as-a-Service (DBaaS) for hosting the auditing databases. We therefore focus this work on *driving a performance and cost analysis when executing the OPAL Simulation Environment in on- and off-premise infrastructures, and using different IaaS offerings and optimized configurations*.

Table 1 shows the different VM categories, based on their characteristics and offered prices by three major Cloud providers: Amazon AWS, Windows Azure, and Rackspace. In addition to the off-premise VM instances types, multiple on-premise VM instances types were created in our virtualized environment, configured in a similar manner to the ones evaluated in the off-premise scenarios, and included in such categories. The on-premise VM instances configurations are based on the closest equivalent to the off-premise VM configurations within each instance category. The encountered providers and offerings showed two levels of VM categories, i.e. based on the optimization for custom use cases (*Micro*, *General Use*, *Compute Optimized*, and *Memory optimized*), and based on a quantitative assignment of virtualized resources. This fact must be taken into consideration in our evaluation due to the variation in the performance, and its impact on the final incurred costs for running simulations in different Cloud offerings. The pricing model for the on-premise scenarios was adopted from [Walker, 2009] as discussed in the following section, while for the off-premise scenarios the publicly available information from the providers was used [Andrikopoulos et al., 2013], taking into account on-demand pricing models only.

## 4.2 Setup

The scientific workflow simulation environment is constituted by two main systems: the SimTech SWfMS [Sonntag and Karastoyanova, 2010, Sonntag et al., 2012], and a set of Web services bundling resource management and the KMC sim-

ulation tasks depicted in [Binkele and Schmauder, 2003, Molnar et al., 2010]. The former comprises the following middleware stack:

- an Apache Orchestration Director Engine (ODE) 1.3.5 (Axis2 distribution) deployed on
- an Apache Tomcat 7.0.54 server with Axis2 support.
- The scientific workflow engine (Apache ODE) utilizes a MySQL server 5.5 for workflow administration, management, and reliability purposes, and
- provides monitoring and auditing information through an Apache ActiveMQ 5.3.2 messaging server.

The resource management and KMC simulation services are deployed as Axis2 services in an Apache Tomcat 7.0.54 server. The underlying on- and off-premise infrastructure configurations selected for the experiments are shown in Table 1. The on-premise infrastructure aggregates an IBM System x3755 M3 server<sup>5</sup> with an AMD Opteron Processor 6134 exposing 16 CPU of speed 2.30 GHz and 65GB RAM. In all scenarios the previously depicted middleware components are deployed on an Ubuntu server 14.04 LTS with 60% of the total OS memory dedicated to the SWfMS.

For all evaluation scenarios a system's load of 10 concurrent users sequentially sending 10 random and uniformly distributed simulation requests/user was created using Apache JMeter 2.9 as the load driver. Such a load aims at emulating a shared utilization of the simulation infrastructure. Due to the asynchronous nature of the OPAL simulation workflow, a custom plugin in JMeter was realized towards receiving and correlating the asynchronous simulation responses. The perceived by the user latency for each simulation was measured in milliseconds (ms). Towards minimizing the network latency, in all scenarios the load driver was deployed in the same region as the simulation environment.

The incurred monetary costs for hosting the simulation environment on-premise are calculated considering firstly the purchase, maintenance, and depreciation of the server cluster, and secondly by calculating the price of each CPU time. [Walker, 2009] proposes pricing models for analyzing the cost of purchasing vs. leasing CPU time on-premise and off-premise, respectively. The

<sup>5</sup>IBM System x3755 M3:[http://www-03.ibm.com/systems/xbc/cog/x3755m3\\_7164/x3755m3\\_7164aag.html](http://www-03.ibm.com/systems/xbc/cog/x3755m3_7164/x3755m3_7164aag.html)

Table 1: IaaS Ubuntu Linux On-demand Instances Categories per Provider (in January 2015).

Instance Category	Cloud Provider	Instance Type	vCPU	Memory (GB)	Region	Price (U\$/h)
Micro	on-premise	micro	1	1	EU (Germany)	0.13
	AWS EC2	t2.micro	1	1	EU (Ireland)	0.014
	Windows Azure	A1	1	1.75	EU (Ireland)	0.06
	Rackspace	General 1	1	1	USA	0.06
General Purpose	on-premise	large	2	4	EU (Germany)	0.26
	AWS EC2	m3.large	2	7.5	EU (Ireland)	0.154
	Windows Azure	A2	2	3.5	EU (Ireland)	0.12
	Rackspace	General 2	2	2	USA	0.074
Compute Optimized	on-premise	compute3.large	4	4	EU (Germany)	0.52
	AWS EC2	c3.large	2	3.75	EU (Ireland)	0.120
	Windows Azure	D2	2	7	EU (Ireland)	0.23
Memory Optimized	Rackspace	Compute 1-3.75	2	3.75	USA	0.1332
	on-premise	memory4.large	2	15	EU (Germany)	0.26
	AWS EC2	r3.large	2	15.25	EU (Ireland)	0.195
Memory Optimized	Windows Azure	D3	4	14	EU (Ireland)	0.46
	Rackspace	Memory 1-15	2	15	USA	0.2522

real cost of a CPU/hour when purchasing a server cluster, can be derived using the following equations:

$$\frac{(1 - 1/\sqrt{2}) \times \sum_{T=0}^{Y-1} \frac{C_T}{(1+k)^T}}{(1 - (1/\sqrt{2})^Y) \times TC} \quad (1)$$

where  $C_T$  is the acquisition ( $C_0$ ) and maintenance ( $C_{1..N}$ ) costs over the  $Y$  years of the server cluster,  $k$  is the cost of the invested capital, and

$$TC = TCPU \times H \times \mu \quad (2)$$

where  $TCPU$  depicts the total number of CPU cores in the server cluster,  $H$  is the expected number of operational hours, and  $\mu$  describes the expected utilization. The utilized on-premise infrastructure total cost breaks down into an initial cost ( $C_0$ ) of approximately 8500\$ in July 2012 and an annual maintenance cost ( $C_{1..N}$ ) of 7500\$, including personnel costs, power and cooling consumption, etc. The utilization rate of such cluster is of approximately 80%, and offers a reliability of 99%. Moreover, the server cluster runs six days per week, as one day is dedicated for maintenance operations. Such a configuration provides 960K CPU hours annually. As discussed in [Walker, 2009], we also assumed in this work a cost of 5% on the invested capital. The cost for the off-premise scenarios was gathered from the different Cloud provider's Web sites.

Table 1 depicts the hourly cost for the CPUs consumed in the different on-premise VM configurations. In order to get a better sense of the scope of the accrued costs, the total cost calculation

performed as part of the experiments consisted of predicting the necessary time to run 1K concurrent experiments. Such estimation was then used to calculate the incurred costs of hosting the simulation environment in the previously evaluated on- and off-premise scenarios. The monetary cost calculation was performed by linearly extrapolating the obtained results for the 100 requests to a total of 1K requests. The scientific library Numpy of Python 2.7.5 was used for performing the prediction of 1K simulation requests. The results of this calculation, as well as the observed performance measurements are discussed in the following.

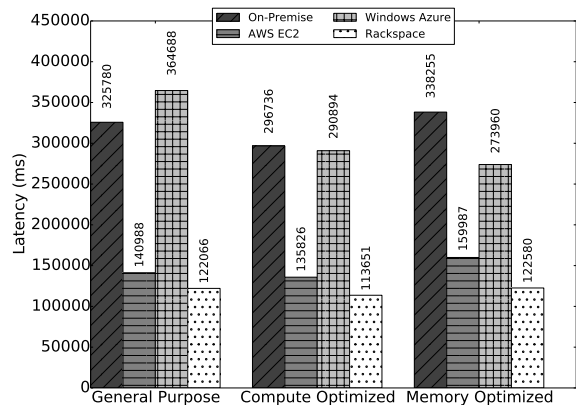



Figure 3: Average Simulation Latency per Provider & VM Category.



## 4.3 Evaluation Results

### 4.3.1 Performance Evaluation

Figure 3 shows the average observed latency for the different VM categories depicted in Table 1 for the different Cloud providers. The latency perceived in the scenarios comprising the selection of *Micro* instances have been excluded from the comparison due to the impossibility to finalize the execution of the experiments. More specifically, the on-premise micro-instance was capable of stably running approximately 80 requests (see Figure 4a), while in the off-premise scenarios the

load saturated the system with 10 requests approximately in the AWS EC2 and Windows Azure scenarios (see Figures 4b and 4c, respectively). For the scenario utilizing Rackspace, the VM micro instance was saturated immediately after sending the first set of 10 concurrent simulation requests.

With respect to the remaining instance categories (*General Purpose*, *Compute Optimized*, and *Memory Optimized*), the following performance variation behaviors can be observed:

1. the on-premise scenario shows in average a latency of 320K ms. over all categories, a 40% higher average than the perceived latency in the off-premise scenarios.

2. However, the performance is not constantly improved when migrating the simulation environment off-premise. For example, the *General Purpose* Windows Azure VM instance shows a degraded performance of 11%, while the Windows Azure *Compute Optimize* VM instance shows only a slightly performance improvement of 2%, when compared with the on-premise scenario.
3. The performance when migrating the simulation environment to the Cloud improves by approximately 56% and 62% for the AWS EC2 and Rackspace *General Purpose* VM instances, respectively,
4. 54%, 2%, and 61% for the AWS EC2, Windows Azure, and Rackspace *Compute Optimized* VM instances, respectively, and
5. 52%, 19%, and 63% for the AWS EC2, Windows Azure, and Rackspace *Memory Optimized* VM instances, respectively.

When comparing the average performance improvement among the different optimized VM instances, the *Compute Optimized* and *Memory Optimized* instances enhance the performance by 12% and 6%, respectively.

Figure 4 shows the perceived requests' latency individually. It can be observed when executing the simulation environment in the Rackspace infrastructure that the performance highly varies when increasing the number of requests (see Figure 4d). Such performance variation decreases in the on-premise, AWS EC2, and Windows Azure infrastructures (see Figures 4a, 4b, and 4c, respectively). In all scenarios, the network latency does not have an impact in the performance due to the nature of our experimental setup described in the previous section.

When comparing the performance improvement among the different VM instances categories, the Windows Azure infrastructure shows the greater when selecting a *Compute Optimized* or *Memory Optimized* VM instance over a *General Purpose* VM instance (see Figure 4c).

### 4.3.2 Cost Comparison

Figure 5 presents an overview of the expected costs for running 1K experiments among 10 users. The following pricing variations can be observed:

1. The incurred costs of hosting the simulation environment on-premise is 25\$ in average.
2. When migrating the simulation infrastructure off-premise, the cost descends in average 80%,

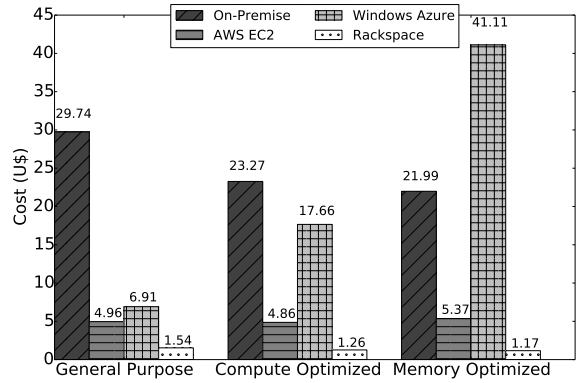


Figure 5: Cost Comparison extrapolated to 1K Simulation Requests (in January 2015 Prices)

12%, and 94% when utilizing the AWS EC2, Windows Azure, and Rackspace IaaS services.

3. When comparing the incurred costs among the different VM categories, the *Memory Optimized* categories are in average 61% and 47% more expensive when compared to the *Compute Optimized* and *General Purpose* VM categories, respectively.
4. Among the different off-premise providers, Windows Azure is in average 900% more expensive for running the simulation environment.

## 4.4 Discussion

The experiments driven as part of this work have contributed to derive and report a bi-dimensional analysis focusing on the selection among multiple IaaS offerings to deploy and run the OPAL Simulation Environment. With respect to performance, it can be concluded that:

1. The migration of the simulation environment to off-premise Cloud services has an impact on the system's performance, which is beneficial or detrimental depending on the VM provider and category.
2. The selection of *Micro* VM instances did not offer an adequate availability to the simulation environment in the off-premise scenarios. Such a negative impact was produced by the non-automatic allocation of swap space for the system's virtual memory.
3. When individually observing the performance within each VM category, the majority of the selected off-premise IaaS services improved the performance of the simulation environment. However, the *General Purpose* Windows Azure VM instances showed a degradation of the

performance when compared to the other IaaS services in the same category.

4. The perceived by the user latency was in average reduced when utilizing *Compute Optimized* VM instances. Such an improvement is in line with the compute intensity requirements of the simulation environment.

The cost analysis derived the following conclusions:

1. There exists a significant monetary cost reduction when migrating the simulation environment to off-premise IaaS Cloud services.
2. Despite of the improved performance observed when running the simulation environment in the *Compute Optimized* and *Memory Optimized* VM instances, scaling the experiments to 1K simulation requests incurred in an average increase of 9% and 61% with respect to the *General Purpose* VM instances cost, respectively.
3. The incurred monetary costs due to the usage of Windows Azure services tend to increase when using optimized VM instances, i.e. *Compute Optimized* and *Memory Optimized*. Such behavior is reversed for the remaining off-premise and on-premise scenarios.
4. Due to the low costs demanded for the usage of Rackspace IaaS services (nearly 40% less in average), the final price for running 1K simulations is considerably lower than the other off-premise providers and hosting the environment on-premise.

The previous observations showed that the IaaS services provided by Rackspace are the most suitable for migrating our OPAL Simulation Environment. However, additional requirements may conflict with the migration decision of further simulation environments, e.g. related to data privacy and transfer between EU and USA regions, as Rackspace offers a limited set of optimized VMs in their European region.

## 5 RELATED WORKS

We consider our work related to the following major research areas: performance evaluation of workflow engines, workflow execution in the Cloud, and migration and execution of scientific workflows in the Cloud.

When it comes to evaluating the performance of common or scientific workflow engines, a stan-

dardized benchmark is not yet available. A first step towards this direction is discussed in [Skouradaki et al., 2015], but propose approach is premature and could not be used as the basis for this work. Beyond this work, performance evaluations are usually custom to specific project needs. Specifically for BPEL engines not much work is currently available. For example [Röck et al., 2014] summarize nine approaches that evaluate the performance of BPEL engines. In most of the cases, workflow engines are benchmarked with load tests with a workload consisting of 1-4 workflows. Throughput and latency are the metrics most frequently used.

There are only few Cloud providers supporting the deployment and execution of workflows in a Platform-as-a-Service (PaaS) solution. The WSO2 Stratos Business Process Server [Pathirage et al., 2011] and Business Processes on the Cloud is offered by IBM Business Process Manager<sup>6</sup> offer the necessary tools and abstraction levels for developing, deploying and monitoring workflows in the Cloud. However, such services are optimized for business tasks, rather than for supporting simulation operations.

Scientific Workflow Management Systems are exploiting business workflows concepts and technologies for supporting scientists towards the use of scientific applications [Sonntag et al., 2011b, Sonntag and Karastoyanova, 2010]. Zhao et al. [Zhao et al., 2014] develop a service framework for integrating Scientific Workflow Management Systems in the Cloud to leverage from the scalability and on-demand resource allocation capabilities. The evaluation of their approach mostly focuses on examining the efficiency of their proposed PaaS based framework.

Simulation experiments are driven in the scope of different works [Binkele and Schmauder, 2003, Molnar et al., 2010]. Later research efforts focused on the migration of simulations to the Cloud. Due to the diverse benefits of Cloud environments the approaches evaluate the migration with respect to different scopes. The approaches that study the impact of migration to the performance and incurred monetary costs is considered more relevant to our work. In [de Oliveira et al., 2011] the authors examine the performance of X-Ray Crystallography workflows executed on the SciCumulus middleware deployed in Amazon EC2. Such workflows are CPU-intensive and requires the execution of high parallel techniques. Like-

<sup>6</sup><http://www-03.ibm.com/software/products/en/business-process-manager-cloud>



wise, in [Juve et al., 2009] the authors compare the performance of scientific workflows migrated from Amazon EC2 to a typical High Performance Computing system (NCSA’s Abe). In both approaches the authors conclude that migration to the Cloud can be viable but not equally efficient to High Performance Computing environments. However, Cloud environments allow the provisioning of specific resources configurations irregularly during the execution of simulation experiments [Strauch et al., 2013]. Moreover, the performance improvement observed in Cloud services provide the necessary flexibility for reserving and releasing resources on-demand while reducing the capital expenditures [Ostermann et al., 2010]. Research towards this direction is a fertile field. Juve et al. [Juve et al., 2013] execute nontrivial scientific workflow applications on grid, public, and private Cloud infrastructures to evaluate the deployments of workflows in the Cloud in terms of setup, usability, cost, resource availability, and performance. This work can be considered complementary to our approach, although we focused on investigating more on public Cloud providers and took into account the different VM optimization categories.

## 6 CONCLUSION AND FUTURE WORK

Simulation workflows have been widely used in the eScience domain due to their easiness to model, and flexible and automated runtime properties. The characteristics of such workflows together with the usage patterns of simulation environments have made these type of systems suitable to profit from the advantages brought by the Cloud computing paradigm. The existence of a vast amount of Cloud services together with the complexity introduced by the different pricing models have become a challenge to efficiently select which Cloud service to host the simulation environment. The main goal of this investigation is to report the performance and monetary cost findings when migrating the previously realized OPAL simulation environment to different IaaS solutions.

A first step in this experimental work consisted of selecting a set of potential IaaS offerings suitable for our simulation environment. The result of such selection covered four major deployment scenarios: (i) in our on-premise infrastructure, and in (ii) three off-premise infrastructures (AWS EC2, Windows Azure, and Rackspace). The selection of the

IaaS offerings consisted of evaluating the different providers and their corresponding optimized VM instances (*Micro*, *General Purpose*, *Compute Optimized*, and *Memory Optimized*). The simulation environment was migrated and its performance evaluated using an artificial workload. A second step in our analysis consisted on extrapolating the obtained results towards estimating the incurred costs for running the simulation environment on- and off-premise. The analyses showed a beneficial impact in the performance and a significant reduction of monetary costs when migrating the simulation environment to the majority of off-premise Cloud offerings.

Despite our efforts towards analyzing and finding the most efficient IaaS Cloud service to deploy and run our simulation environment, our experiments solely focused on IaaS offerings. Future works focus on analyzing further service models, i.e. Platform-as-a-Service (PaaS) or Database-as-a-Service (DBaaS), as well as evaluating the distribution of the different components constituting the simulation environment among multiple Cloud offerings. Investigating different autoscaling techniques and resources configuration possibilities is also part of future work, e.g. feeding the application distribution system proposed in [Sáez et al., 2014b] with such empirical observations.

## ACKNOWLEDGEMENTS

The research leading to these results has received funding from the FP7 EU project ALLOW Ensembles (600792), the German Research Foundation (DFG) within the Cluster of Excellence in Simulation Technology (EXC310), and the German DFG project BenchFlow (DACH Grant Nr. 200021E-145062/1).

## REFERENCES

- Andrikopoulos, V., Song, Z., and Leymann, F. (2013). Supporting the migration of applications to the cloud through a decision support system. In *Cloud Computing (CLOUD), 2013 IEEE Sixth International Conference on*, pages 565–572. IEEE.
- Binkele, P. and Schmauder, S. (2003). An atomistic Monte Carlo Simulation of Precipitation in a Binary System. *Zeitschrift für Metallkunde*, 94(8):858–863.
- de Oliveira, D., Ocaña, K. A. C. S., Ogasawara, E. S., Dias, J., Baião, F. A., and Mattoso, M. (2011).

- A Performance Evaluation of X-Ray Crystallography Scientific Workflow Using SciCumulus. In Liu, L. and Parashar, M., editors, *IEEE CLOUD*, pages 708–715. IEEE.
- Görlach, K., Sonntag, M., Karastoyanova, D., Leymann, F., and Reiter, M. (2011). *Conventional Workflow Technology for Scientific Simulation*, pages 323–352. Guide to e-Science. Springer-Verlag.
- Juve, G., Chervenak, A., Deelman, E., Bharathi, S., Mehta, G., and Vahi, K. (2013). Characterizing and Profiling Scientific Workflows. *Future Gener. Comput. Syst.*, 29(3):682–692.
- Juve, G., Deelman, E., Vahi, K., Mehta, G., Berriman, B., Berman, B., and Maechling, P. (2009). Scientific Workflow Applications on Amazon EC2. In *E-Science Workshops, 2009 5th IEEE International Conference on*, pages 59–66.
- Molnar, D., Binkele, P., Hocker, S., and Schmauder, S. (2010). Multiscale Modelling of Nano Tensile Tests for different Cu-precipitation States in  $\alpha$ -Fe. In *Proc. of the 5th Int. Conf. on Multiscale Materials Modelling*, pages 235–239.
- Ostermann, S., Iosup, A., Yigitbasi, N., Prodan, R., Fahringer, T., and Epema, D. (2010). A Performance Analysis of EC2 Cloud Computing Services for Scientific Computing. In *Cloud Computing*, pages 115–131. Springer.
- Pathirage, M., Perera, S., Kumara, I., and Weerawarana, S. (2011). A Multi-tenant Architecture for Business Process Executions. In *Proceedings of the 2011 IEEE International Conference on Web Services, ICWS '11*, pages 121–128, Washington, DC, USA. IEEE Computer Society.
- Röck, C., Harrer, S., and Wirtz, G. (2014). Performance Benchmarking of BPEL Engines: A Comparison Framework, Status Quo Evaluation and Challenges. In *26th International Conference on Software Engineering and Knowledge Engineering (SEKE)*, pages 31–34, Vancouver, Canada.
- Sáez, S. G., Andrikopoulos, V., Leymann, F., and Strauch, S. (2014a). Design Support for Performance Aware Dynamic Application (Re-)Distribution in the Cloud. *IEEE Transactions on Services Computing (to appear)*.
- Sáez, S. G., Andrikopoulos, V., Wessling, F., and Marquezan, C. C. (2014b). Cloud Adaptation & Application (Re-)Distribution: Bridging the two Perspectives. In *Proceedings EnCASE'14*, pages 1–10. IEEE Computer Society Press.
- Skouradaki, M., Roller, D. H., Frank, L., Ferme, V., and Pautasso, C. (2015). On the Road to Benchmarking BPMN 2.0 Workflow Engines. In *Proceedings of the 6th ACM/SPEC International Conference on Performance Engineering ICPE 2015*, pages 1–4. ACM.
- Sonntag, M., Hahn, M., and Karastoyanova, D. (2012). Mayflower - Explorative Modeling of Scientific Workflows with BPEL. In *Proceedings of the Demo Track of the 10th International Conference on Business Process Management (BPM 2012)*, *CEUR Workshop Proceedings, 2012*, pages 1–5. CEUR Workshop Proceedings.
- Sonntag, M., Hotta, S., Karastoyanova, D., Molnar, D., and Schmauder, S. (2011a). Using Services and Service Compositions to Enable the Distributed Execution of Legacy Simulation Applications. In Abramowicz, W., Llorente, I., Surridge, M., Zisman, A., and Vayssière, J., editors, *Towards a Service-Based Internet, Proceedings of the 4th European Conference Service-Wave 2011, Poznan, Poland, 2011*, pages 1–12. Springer-Verlag.
- Sonntag, M., Hotta, S., Karastoyanova, D., Molnar, D., and Schmauder, S. (2011b). Workflow-based Distributed Environment for Legacy Simulation Applications. In *ICSOFT (1)*, pages 91–94.
- Sonntag, M. and Karastoyanova, D. (2010). Next Generation Interactive Scientific Experimenting Based On The Workflow Technology. In Alhajj, R., Leung, V., Saif, M., and Thring, R., editors, *Proceedings of the 21st IASTED International Conference on Modelling and Simulation (MS 2010)*, 2010. ACTA Press.
- Sonntag, M. and Karastoyanova, D. (2013). Model-as-you-go: An Approach for an Advanced Infrastructure for Scientific Workflows. *Journal of Grid Computing*, 11(3):553–583.
- Strauch, S., Andrikopoulos, V., Bachmann, T., Karastoyanova, D., Passow, S., and Vukojevic-Haupt, K. (2013). Decision Support for the Migration of the Application Database Layer to the Cloud. In *Cloud Computing Technology and Science (Cloud-Com), 2013 IEEE 5th International Conference on*, volume 1, pages 639–646. IEEE.
- Vukojevic-Haupt, K., Karastoyanova, D., and Leymann, F. (2013). On-demand Provisioning of Infrastructure, Middleware and Services for Simulation Workflows. In *Service-Oriented Computing and Applications (SOCA), 2013 IEEE 6th International Conference on*, pages 91–98. IEEE.
- Walker, E. (2009). The Real Cost of a CPU Hour. *IEEE Computer*, 42:35–41.
- Zhao, Y., Li, Y., Raicu, I., Lu, S., Lin, C., Zhang, Y., Tian, W., and Xue, R. (2014). A Service Framework for Scientific Workflow Management in the Cloud. *Services Computing, IEEE Transactions on*, PP(99):1–1.