Scalable Multi-cloud Platform to Support Industry and Scientific Applications

Kiss, T.

Scalable Multi-cloud Platform to Support Industry and Scientific Applications

Tamas Kiss
University of Westminster, Centre for Parallel Computing, London, UK
e-mail: kisst@wmin.ac.uk

Abstract **- Cloud computing offers resources on-demand and without large capital investments. As such, it is attractive to many industry and scientific application areas that require large computation and storage facilities. Although Infrastructure as a Service (IaaS) clouds provide elasticity and on demand resource access, the challenges represented by multi-cloud capabilities and application level scalability are still largely unsolved. The CloudSME Simulation Platform (CSSP) extended with the Microservices-based Cloud Application-level Dynamic Orchestrator (MiCADO) addresses such issues. CSSP is a generic multi-cloud access platform for the development and execution of large scale industry and scientific simulations on heterogeneous cloud resources. MiCADO provides application level scalability to optimise execution time and costs. This paper outlines how these technologies have been developed in various European research projects, and showcases several application case-studies from manufacturing, engineering and life-sciences where these tools have been successfully utilised to execute large-scale simulations in an optimised way on heterogeneous cloud infrastructures.**

Keywords – cloud computing, application level scalability, container technologies, multi-cloud access.

I. Introduction

Cloud computing has the potential to offer computational resources on demand and dynamically scale these resources to fit the requirements of applications. Additionally, purchasing resources from public cloud providers eliminates the need to invest into expensive hardware, and to maintain these resources in-house that otherwise result in significant costs and requires expertise. Due to these features, cloud computing has proven to be rather attractive to both industry and the research community.

On the other hand, utilising cloud computing resources raises various challenges. Porting existing applications to cloud infrastructures and developing new applications that are able to fully utilise the dynamic and elastic nature of resources is not straightforward. The challenge is further increased if the application aims to be independent from the underlying cloud infrastructure. This feature may be required to avoid vendor lock-in and enable the migration of applications between various cloud providers. Additionally, more complex applications realising application workflows may also intend to utilise multiple heterogeneous cloud providers in order to further optimise execution time and costs. For example, some parts of a workflow may run on an internal private cloud and some other components may utilise a public provider to access large amount of resources.

Application developers are often facing the above described challenges when implementing industry and academic cloud applications. In order to address these issues, we developed a set of technologies within a string of European projects that supports application developers when creating flexible and scalable multi-cloud applications.

These technologies include the CloudSME Simulation Platform (CSSP) [1] that incorporates the multi-cloud CloudBroker Platform [2], the workflow-oriented WS-GRADE science gateway framework [3], and the CloudSME AppCenter. The CSSP was developed within the CloudSME (Cloud-based Simulation platform for Manufacturing and Engineering) project [4]. It represents a Platform as a Service (PaaS) solution that fastens up the development and deployment of complex multi-cloud applications (provided as SaaS – Software as a Service for the end-user), and also supports workflow execution and commercial accounting and billing. The CSSP was prototyped with eleven commercial applications from the manufacturing and engineering sector and was also used in several academic application scenarios. Currently, it is further developed in the CloudiFacturing [5] project where further 21 industry applications will apply the technology in near production quality demonstrators.

While the CSSP enables fast development and deployment of cloud applications, it does not optimise these applications either at deployment or at run-time. In order to address this issue, we are currently developing a generic pluggable framework, called MiCADO (Microservices-based Cloud Application-level Dynamic Orchestrator) [6] within the COLA (Cloud Orchestration at the Level of Applications) [7] project. MiCADO enables application developers not only to optimally deploy their applications in container-based cloud platforms, but also provides run-time monitoring and optimisation of these applications based on user defined quality of service parameters (e.g. deadline, cost, performance or security). MiCADO is fully integrated with the CSSP and provides cloud orchestration and optimisation in multi-cloud environments.

The rest of this paper provides a short overview of the above described technologies and also highlights some typical application examples and scenarios.

Identify applicable sponsor/s here. If no sponsors, delete this text box.

II. THE CLOUDSME SIMULATION PLATFORM

The initial motivation when creating the CloudSME Simulation Platform was coming from the manufacturing domain. Large manufacturing companies widely utilise computer simulation to assist their design and manufacturing processes. Simulation can be used to analyse a wide range of physical and chemical processes, manufacturing systems, logistics, transportation networks and supply chains. On the other hand the take up of simulation by small and medium sized enterprises (SMEs) remains low. This is due to high hardware and software costs (simulation typically requires high performance computing capabilities and specialised software packages), and the lack of skills, technical expertise and in general low awareness of such technologies within SMEs. The aim of the CloudSME project was to utilise cloud computing and build a generic platform that supports independent software vendors (ISVs) and consultant companies when building simulation solutions for SMEs. The platform also enables the targeted end-user companies to utilise these solutions as services in the cloud and pay for it on a subscription or on a pay-as-you-go basis. Although the platform was originally built for the manufacturing sector, it was also successfully utilised in a number of academic and research oriented application scenarios.

Fig. 1 demonstrates a high level architecture of CSSP.

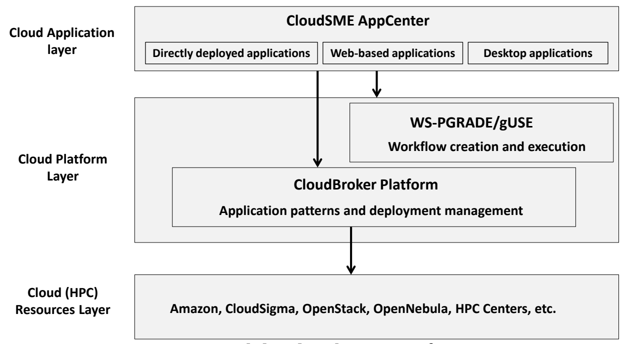

Figure 1. High level architecture of CSSP

The CSSP consists of three main layers. The **Simulation Applications Layer** allows software vendors deploying and presenting simulation products to end-users as SaaS in a wide range of scenarios and deployment models. The **Cloud Platform Layer** provides access to multiple heterogeneous cloud resources and supports the creation of complex application workflows. It is a PaaS to create and execute cloud-based simulations. Finally, the **Cloud Resources Layer** represents the IaaS clouds (and also high performance computing (HPC) resources connected to the platform.

When implementing the CSSP we concentrated on reusing, extending and integrating existing building blocks in order to speed up the development process and provide a robust and reliable solution. However, especially for the top layer that includes commercial components, some new modules also needed to be developed within the project.

The Simulation Applications Layer consists of the CloudSME AppCenter as a generic entry point to browse and execute various simulation applications. The AppCenter is a new component that was developed during the CloudSME project. The CloudSME AppCenter is a web-based one-stop-shop solution with the goal to provide software products and services by software vendors and service providers to end users that can utilise them via a single interface. For this, the AppCenter stores information about software products in an accessible way, provides usage scenarios for the software, and offers billing functionality that includes price setting, payment integration and tracking of users' spending.

The AppCenter supports three main deployment models: Directly Deployed Applications, Desktop Applications and Web-based Applications.

Directly Deployed Applications are published in the AppCenter using its native interface. End-users can access these applications directly from the AppCenter and execute them on the vendor selected cloud or HPC resources. Users will be billed based on software provider defined pricing policies, e.g. resource consumption-based pricing, fixed price per run, periodic subscription charge, etc.

Besides directly deploying and utilizing an application via the native user interface of the AppCenter, developers can also apply its API (Application Programming Interface) to embed accounting and billing functionality into existing web-based or desktop applications. Additionally, using the APIs of the Cloud Platform Layer they can also map these applications to various cloud/HPC resources or even combine then into complex workflows.

Desktop Applications are modified versions of the vendors' original software packages that include AppCenter and Cloud Platform Layer API calls, and redirect and charge for computation intensive simulations on the cloud. These applications can be downloaded from the AppCenter and executed on the users' local machine (but sending computation intensive jobs to remote computation facilities). Web-based applications are simply linked to the AppCenter and use the suitable API calls behind their web interface.

Fig. 2 shows a screenshot of the CloudSME AppCenter providing examples for each three categories. 3D Scan Insole Designer is a Directly Deployed Application that can be run from the AppCenter. TransAT is a Desktop Application that can be downloaded. Finally, CalculiX and OpenFOAM are Web-based Applications linked to the AppCenter.

The Cloud Platform Layer consists of the science gateway framework WS-PGRADE/gUSE and the CloudBroker Platform.

The CloudBroker Platform supports the management and execution of software on different cloud provider resources. The platform provides access to a wide range of resources including open source and proprietary clouds, and also various HPC resources. It can support non-interactive, serial and parallel batch processing applications on both Linux and Windows operating systems. The Platform consists of a set of modules that manage processes, applications, users, finance (accounting, billing and payment), and runtime issues (process monitoring, queuing, resources, storage and images). Application "patterns" are deployed to CloudBroker in a form that allows the platform when instructed to run the application on a particular cloud and cloud instance type. Two typical patterns are direct installation (an application package and deployment script that allows the installation of the software on a cloud instance) or virtualisation (virtual machine image containing installed software that allows direct deployment to a cloud instance).

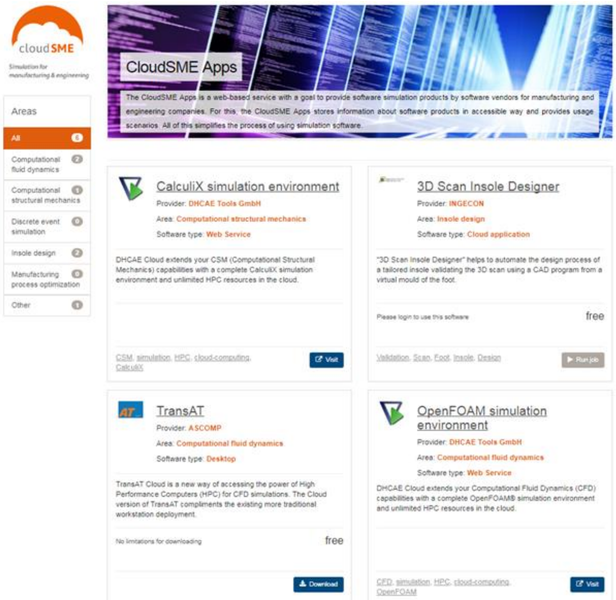

Figure 2. The CloudSME AppCenter

WS-PGRADE/gUSE supports the development and deployment of complex application workflows. It consists of three components: a workflow presentation layer, a workflow management layer and a workflow execution layer. The workflow presentation layer (WS-PGRADE) includes a set of Liferay portlets to create, start and control workflows, monitor their execution on various DCIs (Distributed Computing Infrastructures), and present results to users. The workflow management layer (gUSE) consists of services to support workflow management and execution (file storage, workflow repository, and workflow interpreter to manage the execution of workflows). The workflow execution layer manages job submission as specified in the workflow.

WS-PGRADE/gUSE and CloudBroker were integrated within the SCI-BUS [8] and CloudSME projects enabling the execution of WS-PGRADE workflows on multiple heterogeneous cloud resources supported by the CloudBroker platform.

The bottom layer of CSSP is the Cloud Resources Layer that consists of a range of clouds and HPC resources accessible via the CloudBroker Platform. These currently include Amazon and CloudSigma public clouds, various private clouds based on either OpenStack or OpenNebula, and HPC resources, for example the Cineca Galileo Cluster or the ETH Euler Cluster.

III. Applications on the CSSP

In order to prove the applicability of the CSSP, eleven near production quality demonstrators, in two waves, were developed in CloudSME. Altogether 24 companies, all of them SMEs, were involved in the project. These companies included cloud resource and platform providers responsible for the underlying resources and infrastructure, independent software vendors and software development/consultant companies providing various simulation software solutions, and manufacturing and engineering companies as end-users of the developed technology.

The strength of the CSSP is that it provides support for the fast and efficient development and deployment of a wide variety of applications. These application scenarios fall into three main categories.

Parallel multiprocessor applications use cloud instances or HPC resources to run massively parallel applications on multiple CPUs of one instance or on multiple instances of a computing cluster or HPC Center. Such CloudSME application example is described in [9] that presents TransAT (Transport Phenomena Analysis Tool), a computational fluid dynamics (CFD) application by ASCOMP GmbH [10], and analyses its performance speed-up when applying the CSSP.

The second category is parameter sweep applications that use multiple independent cloud instances to speed up experimentation. In this scenario large number of independent jobs has to be executed independently from each other typically followed by a shorter evaluation of results. A CloudSME application representing this scenario is also presented in [9] where Repast Simphony [11], an open source agent-based modelling and simulation tool was benchmarked on the CSSP.

The final category is scalable web applications where the challenge is the dynamically changing number of users and fluctuating load of the application that require a flexible amount of resources. A commercial application example where a custom web portal was implemented and connected to the CSSP for uploading and validating scanned images for custom foot insole design is presented in [12]. However, please note that this solution developed within CloudSME does not offer automated scalability and only scales up following operator intervention. This issue will be further addressed in Section IV.

Besides industry utilization, the usability of CCSP was also demonstrated in several academic and research application scenarios. The Raccoon2 [13] molecular docking desktop application was successfully extended with multi-cloud support as described in [14], and the automated setup and workflow-based execution of applications using Hadoop [15] was offered for users of the European Grid Infrastructure Federated Cloud (EGI FedCloud) based on the work in [16].

IV. Automated Scalability of Cloud Applications

One of the limitations of CSSP is that the resources utilized for the execution of applications are configured at the beginning of the experiment and cannot be dynamically adjusted at run-time. However, many scientific and commercial applications require access to computation, data or network resources based on dynamically changing requirements. This challenge is addressed by the MiCADO Framework that is currently being developed within the COLA European project.

The concept of MiCADO is to provide application level scalability of cloud computing resources. The aim is to create a framework, where automatically adjusted cloud service supply can be arranged, based on application demands. Such framework allows cloud application developers to build cost and performance optimization mechanisms into their applications using a high level definition language. The suggested solution is based on microservices [17] and their dynamic orchestration in a cloud computing environment.

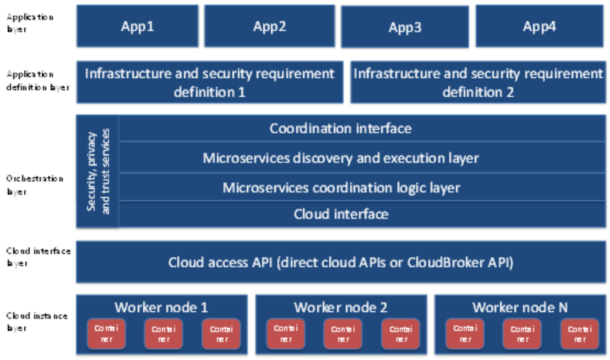

Figure 3. MiCADO Layered Generic Architecture

The generic layered architecture of MiCADO is shown in Fig. 3. The *Application layer* contains actual application code and data to make an application definition functioning in such a way that a desired functionality is reached. The *Application definition layer* forms the basis to define a functional architecture of the application using application templates. This layer describes the application topology and the deployment and scalability policies using TOSCA (Topology and Orchestration Specification for Cloud Applications), an OASIS standard [18]. The *Orchestration layer* is divided into four horizontal and one vertical sub-layer. The horizontal sub-layers include the *Coordination interface API* that provides access to orchestration control and decouples the orchestration layer from the application definition, the *Microservices discovery and execution layer* that manages the execution of microservices and keeps track of services running, the *Microservices coordination logic layer* that gathers and processes information in order to support decisions regarding scaling up or down the infrastructure, and the *Cloud interface API* that abstracts cloud access from layers above. The orchestration layer also includes a vertical sub-layer called *Security, privacy and trust services* that shields application developers from detailed security management. The *Cloud interface layer* provides means to launch and shut down cloud instances. Finally, the *Cloud instance layer* contains cloud instances provided by IaaS cloud providers.

As the cloud interface layer of the above described architecture can be implemented using the CloudBroker platform, MiCADO can seamlessly extend the CSSP when deploying and managing applications that require dynamic automated scalability.

The first prototype implementation of MiCADO, presented in [6] is based on the integration and extension of widely used open source container-based cloud technologies, such as Docker Swarm [19], Occopus [20], and Prometheus [21]. Docker Swarm is a widely used docker-based service where the Docker master node is responsible for the dynamic creation and management of Docker containers at the worker nodes of the Swarm cluster. In order to take the decision when the number of workers in the Swarm cluster should be increased or decreased, Prometheus is used as monitoring tool to observe the load of the worker nodes. The collected data is used to create alerts. These alerts are events which if triggered then they are sent out to other services of MiCADO. The alerts are generated by the Alert manager service. Prometheus provides only a framework to create the Alert manager service therefore it was our task to develop the alert conditions and code. Similarly the Alert executor service was developed by us based on the Prometheus framework concept. The Alert executor instructs the Occopus cloud orchestrator to deploy new Swarm worker nodes or remove existing ones depending on the information provided by the Alert manager service.

IV. APPLICATIONS OF MICADO

The two main targeted application categories for MiCADO are cloud-based services where scalability is achieved by scaling up or down the number of containers and virtual machines based on load, performance and cost, and the execution of a large number of (typically parameter sweep style) jobs where a certain number of these jobs need to be executed by a set deadline.

The first MiCADO implementations were primarily focusing on the first application category where the scalability of cloud-based services needs to be automated. Such application example is presented in [6] where the Data Avenue data staging service [22] was successfully extended with MiCADO based automated scalability. Data Avenue enables to transfer files or directories between different data storages having various storage access protocols (e.g. HTTP, HTTPS, SFTP, GSIFTP, SRM, iRODS, S3). Data Avenue is also used in gUSE/WS-PGRADE portals either as a portlet or inside WS-PGRADE workflows to enable file transfer based communication between workflow nodes. The original version of Data Avenue proved to be a bottleneck when a large number of users wanted to transfer large amounts of data. The MiCADO-based implementation overcomes this problem as it is highly scalable and can be automatically replicated whenever the usage of the service exceeds a pre-defined threshold. Similarly, once the load is reduced below this threshold, under-loaded replicas can be removed from the cloud.

The COLA project currently implements three large scale industry and public sector demonstrators to showcase the capabilities of MiCADO. One of these examples, Magician, an application for Social Media data mining by Spanish company Inycom, is described in [23]. Aragon Regional Government in Spain has set the goal to develop communication channels with citizens to become aware of their opinion about the local government's services, and how these can be improved. The authorities also want to offer information to entrepreneurs and companies in the region that can be used to improve businesses or develop new ones. The local government already collects large amounts of data resulting from interactions with its citizens (e.g. applying for public services such as grants, aids, subsidies, and licenses). This data can be combined and extended with information publicly available in social networks. The aim is to set up a business gateway that is supported by the intelligent analysis and utilisation of all available information.

Two relevant aspects of such social media analysis tools are the ever increasing size of the amount of data available and, at the same time, the unpredictable fluctuation of computing load required due to the high level of uncertainty on how much information will be collected by the crawlers to be processed later. Additionally, databases and software processes may run on different Clouds requiring access to a heterogeneous set of resources in a dynamic way. In order to meet the above challenges, Magician is currently being prototyped within COLA to extend it with automated scalability.

V. Conclusion and Future Work

This paper presented an overview of technologies and solutions developed in a string of European research projects to address scalable multi-cloud deployment of industry and academic applications. The CloudSME Simulation platform has already been utilized for the development and cloud deployment of various applications from the manufacturing and engineering domain and also by the research community. In the currently running CloudiFacturing project the technology is expected to mature further and be used by up to 21 new industry applications.

In order to address the automated scalability of such cloud applications, we are developing the MiCADO framework that supports application developers to embed such capabilities into their applications without requiring low level custom development and cloud expertise.

Future work is twofold. We are further developing the CSSP based on the requirements of the CloudiFacturing application scenarios. On the other hand, MiCADO is being extended with policy keeper and optimiser modules in order support various application types and scalability parameters.

Acknowledgment

The work presented in this paper was supported by the following projects: CloudSME Project No. 608886, COLA Project No. 731574, and CloudiFacturing Porject No. 768892.

References

[1] S. J. E. Taylor, T. Kiss, G. Terstyanszky, P. Kacsuk and N. Fantini: Cloud Computing for Simulation in Manufacturing and Engineering: Introducing the CloudSME Simulation Platform, Proceedings of ANSS 14, Annual Simulation Symposium 2014, April 13 - 16, 2014, Tampa, FL, USA, Article No. 12, Society for Computer Simulation International.

[2] The CloudBroker Platform website – http://www.cloudbroker.com

[3] P. Kacsuk, ed., Science Gateways for Distributed Computing Infrastructures, Springer International Publishing, Cham, 2014. doi:10.1007/978-3-319-11268-8.

[4] CloudSME Project website - http://cloudsme-project.eu/

[5] CloudiFacturing Project Website - https://www.cloudifacturing.eu/

[6] T. Kiss, P. Kacsuk, J. Kovacs, B. Rakoczi, A. Hajnal, A. Farkas, G. Gesmier, G. Terstyanszky, MiCADO—Microservice-based Cloud Application-level Dynamic Orchestrator, Future Generation Computing. Systems (2017). doi:10.1016/j.future.2017.09.050

[7] COLA Project Website - https://project-cola.eu/

[8] SCI-BUS Project Website - https://www.sci-bus.eu/

[9] S.J.E. Taylor, T.Kiss, A.Anagnostou, G. Terstyanszky, P. Kacsuk, J. Costes, N. Fantini: The CloudSME Simulation Platform and it Applications, submitted to Future Generation Computing. Systems (2018)

[10] ASCOMP TransAT Suite - http://ascomp.ch/products/transat-suite/

[11] M. J. North, N. T. Collier, J. Ozik, E. R. Tatara, C. M. Macal, M. Bragen, and P. Sydelko, "Complex adaptive systems modeling with repast simphony," Complex Adaptive Systems Modeling, vol. 1, no. 1, p. 3, Mar 2013. [Online]. Available: https://doi.org/10.1186/2194-3206-1-3

[12] G. Terstyanszky, T. Kiss, S. Taylor, A. Anagnostou, M. Subira, G. Padula, E.D.M. Alonso, J.M.M. Rapun, Validating Scanned Foot Images and Designing Customized Insoles on the Cloud, in: 2016 49th Hawaii Int. Conf. Syst. Sci., IEEE, 2016: pp. 3288–3296. doi:10.1109/HICSS.2016.411.

[13] S. Forli et al., "Computational protein-ligand docking and virtual drug screening with the AutoDock suite". *Nat. Protoc.*, vol. 11, no. 5, pp. 905-919, Apr. 2016.

[14] D. Temelkovski, T. Kiss, and G. Terstyanszky, "Molecular docking with Raccoon2 on clouds: extending desktop applications with cloud computing", in *9th International Workshop on Science Gateways*, Poznan, Poland, 2017.

[15] Apache Hadoop - http://hadoop.apache.org/

[16] S. Gugnani, C. Blanco, T. Kiss, G. Terstyanszky: Extending Science Gateway Frameworks to Support Big Data Applications in the Cloud, Journal of Grid Computing, 13 June 2016, DOI 10.1007/s10723-016-9369-8, ISSN 1570-7873, Vol 14, Issue 4.

[17] Balalaie, A., Heydarnoori, A., & Jamshidi, P. (2015). Migrating to Cloud-Native Architectures Using Microservices: An Experience Report

[18] OASIS Topology and Orchestration Specification for Cloud Applications Version 1.0, Available from: http://docs.oasis-open.org/tosca/TOSCA/v1.0/TOSCA-v1.0.html

[19] Docker Documentation, [online] Available from: < https://docs.docker.com/>

[20] G. Kecskemeti, M. Gergely, A. Visegradi, Z. Nemeth, J. Kovacs, P. Kacsuk: One Click Cloud Orchestrator: Bringing Complex Applications Effortlessly to the Clouds, In: Euro-Par 2014. Lecture Notes in Computer Science (8806), Springer, Cham, pp. 38-49. ISBN 978-3-319-14312-5 10.1007/978-3-319-14313-2_4

[21] Prometheus web page - https://prometheus.io/

[22] A. Hajnal, Z. Farkas, P. Kacsuk: Data Avenue, Remote Storage Resource Management in WS-PGRADE, in proceedings of IWSG 2014, 6th International Workshop on Science Gateways, IEEE, 25 August 2014, DOI: 10.1109/IWSG.2014.7

[23] G. Pierantoni, T. Kiss, G. Gesmier, J Des Luriers, G. Terstyanszky, J. M. Martin Rapun: Flexible Deployment of Social Media Analysis Tools, to appear in Proceedings of IWSG 2018, International Workshop on Science Gateways, 13 – 15 June, 2018, Edinburg, UK.