Cloud Computing for Simulation in Manufacturing and Engineering: Introducing the CloudSME Simulation Platform

Simon J. E. Taylor
Department of Computer Science, Brunel University, Uxbridge, UB8 3PH, U.K.
simon.taylor@brunel.ac.uk

Tamas Kiss, Gabor Terstyanszky
Centre for Parallel Computing, University of Westminster, 115 New Cavendish St, London, W1W 6UW, UK.
t.kiss@westminster.ac.uk, G.Z.Terstyanszky@westminster.ac.uk

Peter Kacsuk
MTA SZTAKI, 1132 Budapest, Victor Hugo u. 18-22, Hungary.
kacsuk.peter@sztaki.mta.hu

Nicola Fantini
ScaleTools Schweiz AG, Huobstrasse 10, 8808 Pfäffikon, Switzerland.
Nicola.Fantini@scaletools.com

Keywords: Cloud computing, MSaaS, discrete-event simulation, computational fluid dynamics, 3D modelling, aircraft logistics

Abstract

There are many benefits of implementing simulation software on a Cloud. These include the use of scalable, on-demand access to resources that can be used to speed up simulation and the development of composable Modeling & Simulation services (Modeling & Simulation as a Service (MSaaS)). However, developing Cloud computing solutions for simulation in industry is difficult without appropriate expertise due to the complex technologies involved. This paper introduces the CloudSME Simulation Platform that is based on the proven technologies of gUSE and CloudBroker and is being used to create Cloud computing versions of industrial simulation software including Simul8's discrete-event simulation environment, Ascomp's TransAT computational fluid dynamics application, Ingecon's 3D Scan Insole Designer tool and 2MORO's Bfly software for aircraft maintenance logistics. The paper discusses the Platform and outlines how it is used to implement Cloud-based simulation software.

1 INTRODUCTION

Cloud computing is attractive as it seeks to deliver new levels of service and ICT resource integration [Velte, et al. 2009]. The *National Institute of Standards and Technology* (NIST) defines cloud computing as "a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction" [Grance and Mell 2011].

In a recent survey of what practitioners need from web-based simulation, over half indicated that fast simulation response time, as shorter time as possible to obtain results from simulation runs, was the most desirable feature [Onggo, et al. 2014]. cloud computing can potentially deliver this speedup. The ubiquitous use of Cloud computing for Modeling & Simulation (M&S) appeared frequently in a set of Grand Challenges in M&S discussions [Taylor, et al. 2012, Taylor, et al. 2013a, Taylor, et al. 2013b, Taylor et al. 2013c]. Tolk, for example, discusses the concept of using cloud computing to deliver composable M&S services, i.e. Modeling and Simulation as a Service (MSaaS) [Taylor, et al. 2013a].

To investigate how Cloud computing for Simulation can be used for M&S in an industrial setting, the Cloud Computing for Simulation in Manufacturing and Engineering project (CloudSME - www.cloudsme.eu) has developed the CloudSME Simulation Platform. This is based on the gUse [Kacsuk, et al. 2012] and CloudBroker (www.cloudbroker.com) technologies that were integrated and used successfully to implement Cloud versions of several major scientific applications in the FP7 SCI-BUS project (www.sci-bus.eu). CloudSME is developing MSaaS by porting an initial set of industrial M&S tools including Simul8's discrete-event simulation environment (www.simul8.com), Ascomp's TransAT computational fluid dynamics application (http://www.ascomp.ch/transat/), Ingecon's 3D Scan Insole Designer tool and 2MORO's Bfly software for aircraft maintenance logistics. This list will be expanded in 2014 and the development of the Cloud versions of these tools will be reported in a later paper.

This paper introduces and describes the elements of the CloudSME Simulation Platform. The paper is structured as follows. Section 2 briefly discusses Cloud computing and M&S. Section 3 introduces the architecture and major functionality of the CloudSME Simulation Platform. Section 4 discusses how simulation software can deployed using this architecture and Section 5 concludes the paper.

2 CLOUD COMPUTING AND M&S

Cloud computing has the potential to deliver flexible on-demand computing resources transparently. After defining *Cloud computing*, Grance and Mell [*ibid*] go on to identify a generic cloud model that consists of five essential characteristics:

- on-demand self-service (automatic deployment of computing capabilities),
- broad network access using multiple platforms (such mobile devices, laptops and desktop computers),
- resource pooling,
- rapid elasticity (computing capabilities can be scaled up and down to match the fluctuation in demand) and
- measured service (such as pay-as-you-use model).

Three service models have evolved and the follow terms are in common use. These are:

- Software as a Service (SaaS),
- Platform as a Service (PaaS) and
- Infrastructure as a Service (IaaS).

SaaS translates a software application so that it runs on a cloud, it is accessible as a Cloud-based service that allows consumers and other Cloud-based services to access and run and/or interoperate with the application on demand. PaaS provides the facilities (the platform) to create and manage these application-based services. IaaS provides the overall components to build, manage and deliver a complete cloud-based system consisting of network access, computing, storage and appropriate middleware. As discussed above, the SaaS service model when used to support M&S is sometimes referred to as MSaaS.

Four service provision models are also possible:

- *Private clouds* are in-house or hosted for internal use of a particular organization only. They are multi-tenant across organization sites, departments, groups and/or users, and typically mainly focus on self-service and accountability inside the organization (e.g. using in house clusters or institutional desktop Grids deployed as Clouds).
- *Community clouds* are not only for a single organization, but for a larger, but still separate community across different organ izations and/or individuals. They follow similar principles, but more emphasize the sharing aspects of cloud (e.g. the above deployed within a common secure domain or a volunteer computing Cloud).
- *Public clouds* are offered by certain organizations or cloud providers to a larger community or basically to everybody. They are multi-tenant across organizations and/or individuals, and focus usually on the on-demand and pay-per-use advantages of cloud computing (e.g. the Clouds offered by Amazon, IBM, CloudSigma, etc.)
- *Hybrid clouds* mix public, community and/or private clouds. They typically focus on providing scalability and failover features.

What does this mean for MSaaS? If, for example, a simulation user needs to get results faster, then in an ideal world (and subject to cost!) the user should be able to select an option within the software s/he is using and *get* those results faster. Obviously this will be dependent on the characteristics of the simulation software and the model being simulated. However, the ability to get the extra computing resources on-demand in a completely transparent way is a very attractive aspect of Cloud computing. Additionally, the deployment of simulation software in a modern online environment so that it is easily interoperability with other simulation services and, for example, ERP systems is also attractive.

In terms of MSaaS research there has been much interest. However, as reviewed by Zhao, et al. [2012] and Sakellari and Loukas [2013], most of this work has been to use M&S to support cloud research rather than using cloud computing *for* M&S. There are a relatively small number of examples where MSaaS is specifically studied. For example, Rossetti and Chen [2012] investigate a cloud computing architecture for supply chain network simulation, Hong and Luo [2011] discuss the use of cloud computing to support large scale ranking and selection, and Lu et al. [2012] investigate architectures for cloud-based simulation. In industry, for example, Simul8 (www.simul8.com) has a Cloud-based version of their software (YouSimul8) that allows users to deploy, access and visualize simulations on the Web.

Overall, cloud computing offers great opportunities for MSaaS in terms of speeding up experimentation and testing, allowing more experimentation and testing during a simulation project, better accuracy of results, wider use of optimization techniques and interoperability with other simulations and systems. However, deploying effective Cloud computing infrastructures in which M&S applications can be easily developed requires quite sophisticated support and technology. The next section introduces the CloudSME Simulation Platform that has been developed to support the rapid development of industrial Cloud computing products for M&S in manufacturing and engineering.

3 THE CLOUDSME SIMULATION PLATFORM

The CloudSME Simulation Platform has been developed to support the rapid development of a range of Cloud-based implementation options for simulation software. The core architecture of the CloudSME Simulation Platform (CSSP) is shown in Figure 1. As can be seen the CSSP is composed of two major components: WS-PGRADE/gUSE and CloudBroker.

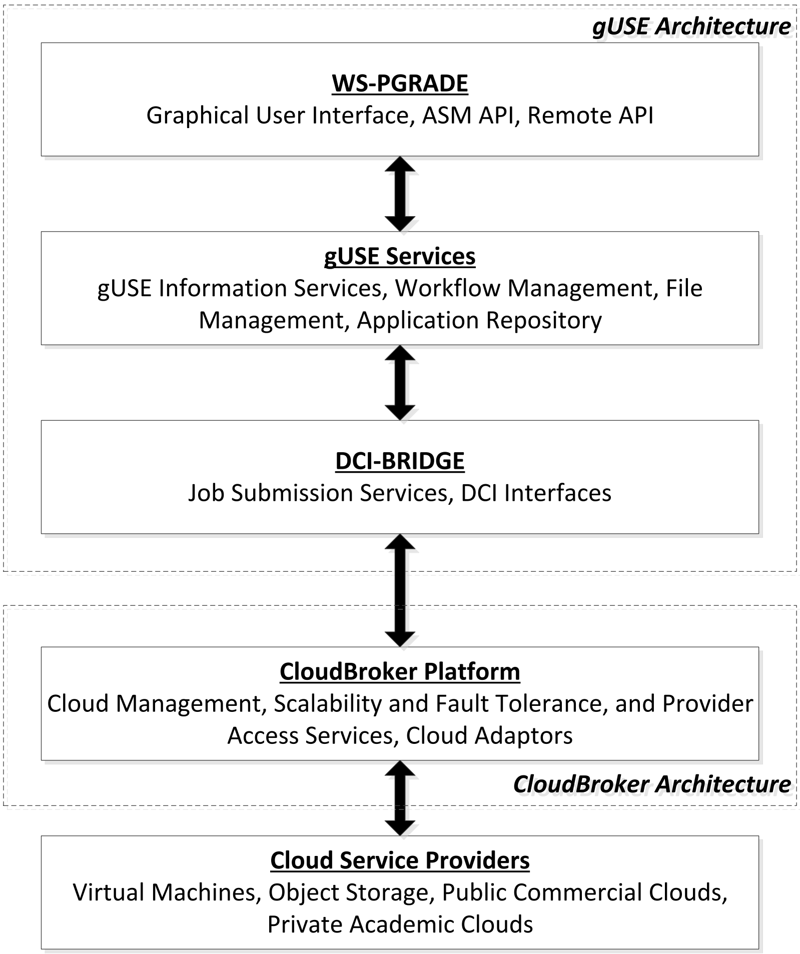

Figure 1. The CloudSME Platform (www.cloudsme.eu)

3.1 Grid User Support Environment (gUSE)

Grid User Support Environment (gUSE) was originally developed to support the needs of application development for Grid computing on different Distributed Computing Infrastructures (DCIs) such as clusters and institutional desktop Girds. This has been used to implement a wide range of scientific computing software over a wide range of DCIs. It allows developers to easily specify the workflow (the sequences of tasks to be executed on the machines of a DCI) required by their application and parallel execution (the tasks that can be executed in parallel on a DCI's machines). It is therefore ideal for Cloud computing as gUSE treats the computing resources of a Cloud (virtual machines) as just another DCI.

As shown in more detail in Figure 2, the major elements of gUSE are WS-PGRADE, gUSE Services and the DCI-BRIDGE.

Overall, to develop an application to run on a Cloud or a Grid a developer will use the graphical environment of WS-PGRADE to describe the sequence of tasks to run the application by describing the *workflow* of the application. A workflow is a directed acyclic graph of nodes and arcs where a node represents a specific task and an arc (typically) a file transfer between tasks. Arcs are connected to specific ports on each node to distinguish actions specific to those edges. Nodes can be configured to perform virtually any computing task via a visual interface. Some specialist nodes have been created to support High Performance Computing (HPC) tasks. In many HPC tasks an application will create a set of jobs that need to be executed in parallel using the machines of a DCI. These jobs need to be sent out to the DCI and the results collected when they return. As shown in Figure 3, creating a workflow for this using WS-PGRADE is very simple. A user creates three nodes and links them together by ports. The first node is specified as a *Generator* that composes the work to be done by an application into a set of jobs. The processing of these tasks is described in the Parallel Task node. For example, this node could simply run the application on each job received by the node. Setting the input port of this node to *parametric* means that a new instance of the parallel task will be created on a machine of the chosen DCI to run the task when a job is received from the Generator. This means that this node will execute the task in parallel subject to the availability of machines on the selected DCI (the DCI-BRIDGE will queue jobs that cannot be executed immediately). Specifying the last node as a *Collector* means that when a task is finished the results are returned to the

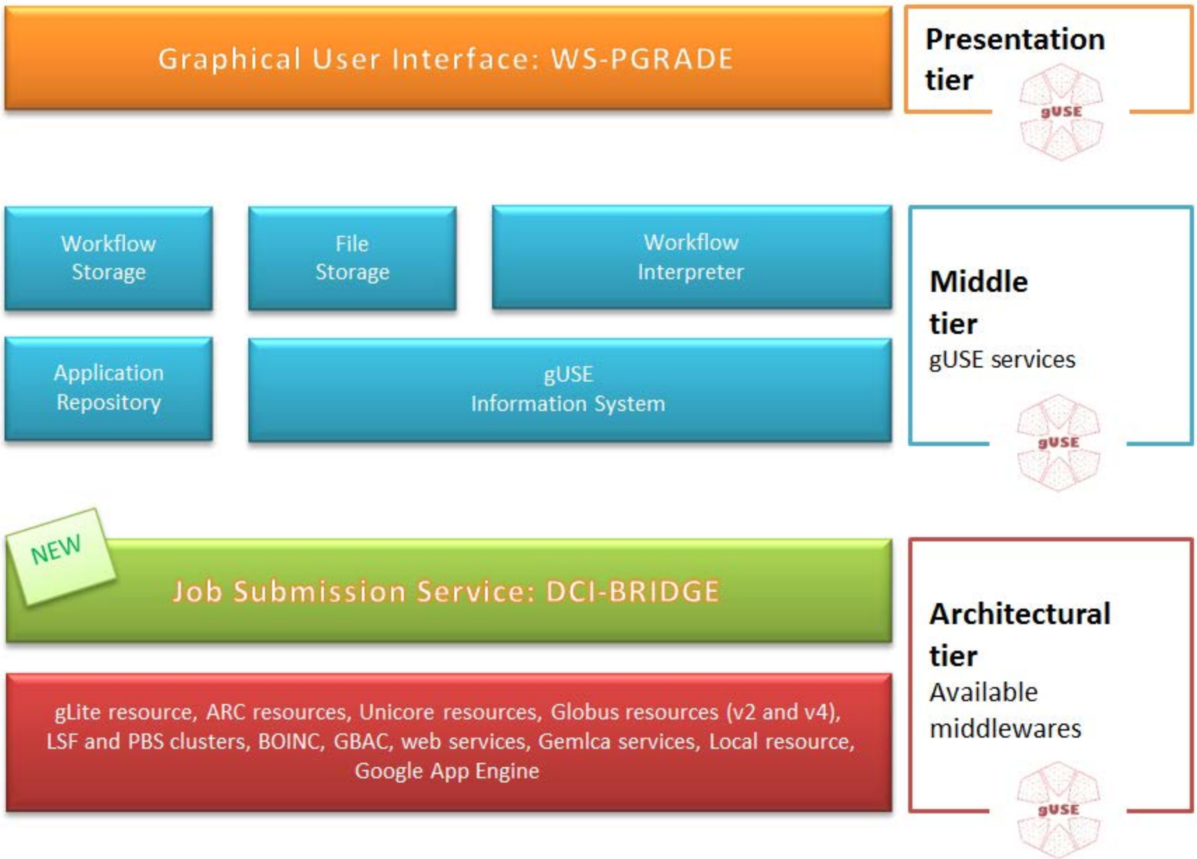

Figure 2. The Grid User Support Environment

collector and given to the user. A user can also specify which DCI the application will run on. This includes the different Cloud technologies made available by CloudBroker. Therefore, if a user wishes to deploy the above workflow on an Amazon Cloud, the user simply specifies this as the CloudBroker/Amazon DCI. For HPC applications the Cloud resources can be scaled to support the computing needs based on availability and costs. This means that a Cloud computing application can be developed extremely rapidly due to the advanced functionality of the CSSP architecture. See the WS-PGRADE Portal Cookbook (http://sourceforge.net/projects/guse/files/WS-PGRADE-Cookbook.pdf/download) for more details on this and more complex workflow development.

Once an applications developer has created a workflow it is saved in the gUSE Services environment. This a complete environment for workflow support and execution. It provides secure access and authorisation to DCI resources by using appropriate security certification. In the CSSP architecture this transparently uses the security processes of the CloudBroker platform. It enables the management, storage and execution of workflows and has flexible deployment options (single hosting/distributed hosting to optimize resource usage or increase performance.) It also offers data services (user database, proxy credential database, application code database, workflow database, application history and change log database, etc.) and control services (e.g. job submission, workflow enactment, etc.). The gUSE Information System supports workflow discovery and naming services.

To execute a workflow, gUSE uses the DCI-BRIDGE to access the resources of various DCIs (in the CSSP these are the virtual machines of the Cloud provider selected in the workflow). The DCI-BRIDGE is a web service-based application that provides standard access to DCIs via DCI plug-ins. Workflows are submitted and managed using the *Basic Execution Service* (BES) interface. This hides the access protocol and the technical details of the DCIs and uses a the standardised job description language JSDL. BES uses a job registry, an input queue, an upload manager to manage the execution of

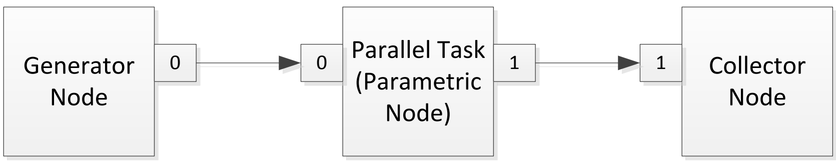

Figure 3. A Simple Workflow

jobs generated by a workflow.

For more information on these services see gUSE in a Nutshell (http://sourceforge.net/projects/guse/files/gUSE_in_a_Nutshell.pdf/download).

3.2 The CloudBroker Platform

The CloudBroker Platform is a web-based application store for the deployment and execution of compute-intensive scientific and technical software in the cloud. It is provided as a Cloud service that allows for on-demand, scalable and pay-per-use access via the Internet. CloudBroker can be deployed as a public version under https://platform.cloudbroker.com, as hosted or in-house setup, as well as licensed software. This also means that the platform can be run at different physical places and under different legislation if desired. A general overview of the platform functionality is shown in Figure 4.

CloudBroker is cross-domain and supports all kinds of non-interactive, batch-oriented applications, both serial and parallel ones, with a focus on software running under Linux. Windows is also supported. Users can access CloudBroker directly via a web-based interface or via the CloudBroker Web Service API. These give access to a set of modules that manage processes, applications, users, finance (accounting, billing and payment), and runtime issues (process monitoring, queuing, resources, storage and images). A scalability and fault handler layer supervises scalability requirements and failure issues. Cloud Provider Access Management oversees the connection to each Cloud technology. These connections are provided by plug-in modules for various different public and private Clouds. CloudBroker incorporates two main types of security features:

- communication layer security using SSL transport layer encryption both between client and platform and between platform and cloud infrastructures, and
- authentication and authorization security using login and password for each user and different levels of organization and user roles

The platform employs industry standard application and server technology and is typically operated in industry standard secure data centres. The different cloud infrastructures it utilizes usually also provide authentication mechanisms, isolated virtual machines and security-certified cloud technologies and data centres. Security is also integrated with gUSE.

CloudBroker uses IaaS from resource providers and provides PaaS to software vendors and SaaS to end users. It offers a marketplace where users can register, deploy, charge for and use their own Cloud resources and application software or use resources and software provided by others (e.g., CloudBroker itself). Surcharges for platform usage are derived as percentage of the resource and software prices. From this follows a freemium model, that is, if resource providers and software vendors set zero prices for their resources and software, the corresponding platform usage is also for free.

CloudBroker can be accessed in several ways. Its main two operation modes to manage and use

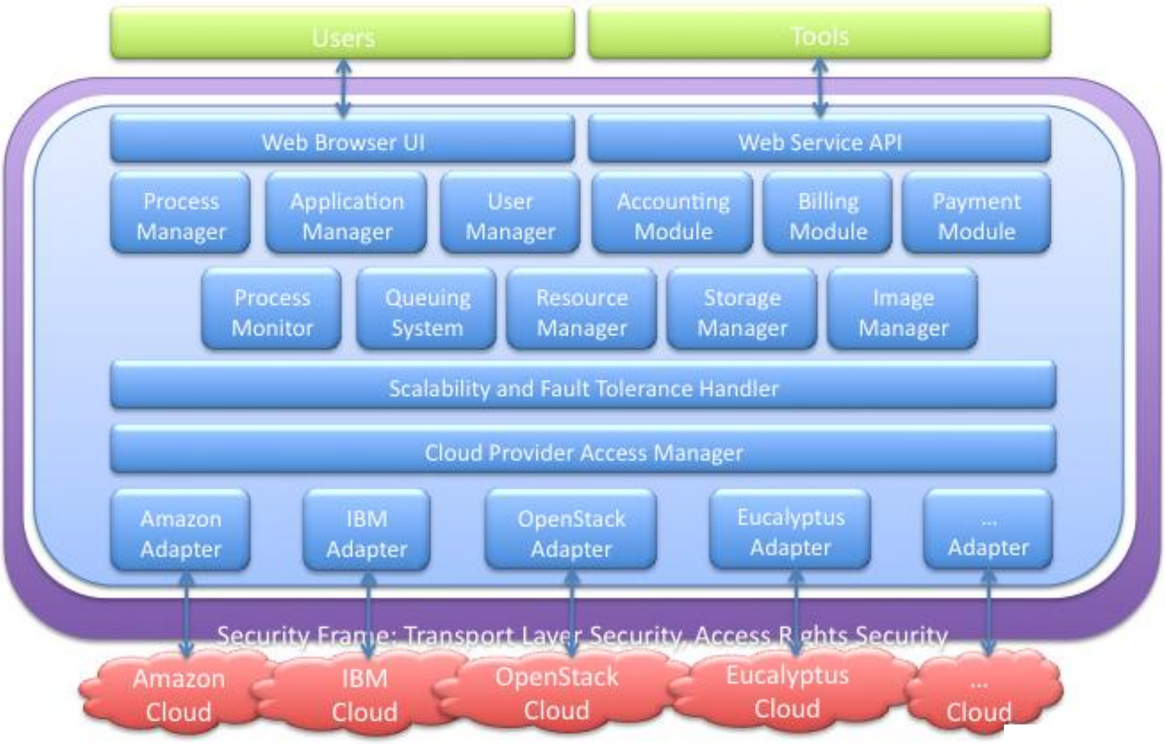

Figure 4. The CloudBroker Platform (www.cloudbroker.com)

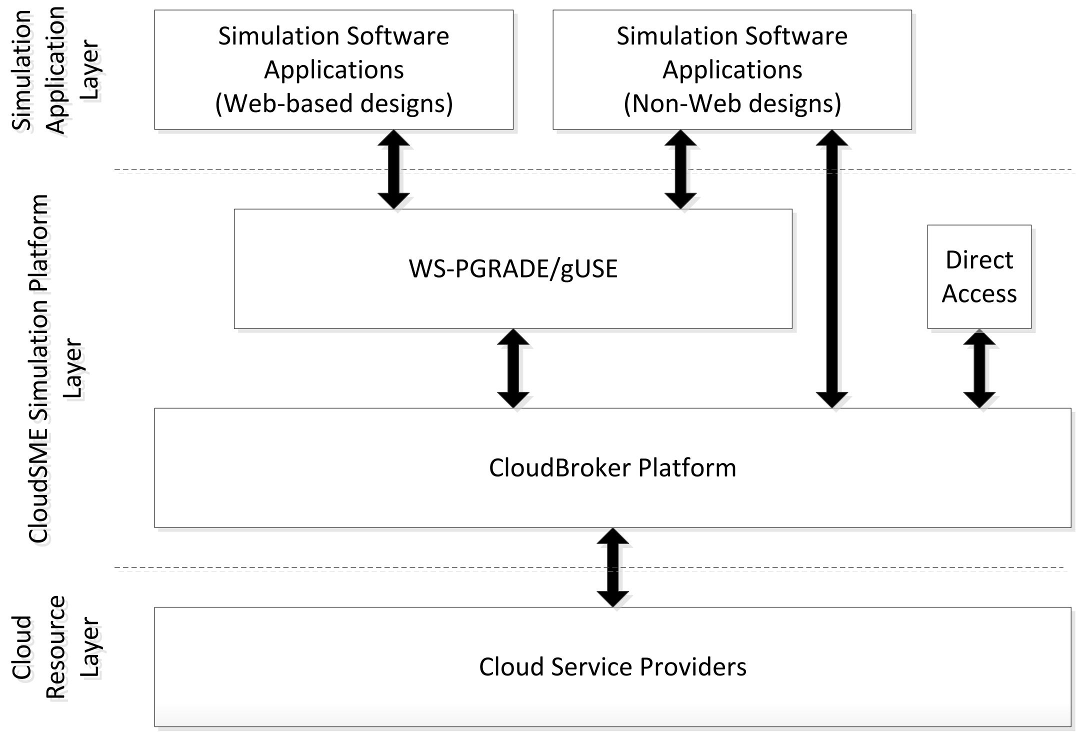

Figure 5. CloudSME Simulation Platform Design Options

software on a Cloud are either as direct front-end or as back-end middleware service. For the former, the platform can be accessed with any regular Web browser as a service via the Internet. This is fine for first-time and individual usage. However, for frequent, advanced and automatic usage, API access is provided. These include a REST web service interface, Java client library and Linux shell command line interface. Using the different APIs, CloudBroker can be utilized by front-end software as middleware to allow access to applications in the cloud. More information on this can be found on the CloudBroker website cloudbroker.com.

4 Cloud-based Simulation Software Design Options

Figure 5 shows the CloudSME Simulation Architecture. This shows the deployment of simulation software in the *Simulation Application Layer*, the CSSP deployed in the *Cloud Simulation Platform Layer* and Cloud computing resources made available in the *Cloud Resource Layer*.

The Figure shows the several available design options available to implement simulation software on a Cloud based on the functionality of the CloudSME Simulation Platform. These are:

- Web-based simulation software
- Non-Web-based simulation software (implemented, for example, as a Windows package)
- Direct access to the simulation software implemented on the CloudBroker Platform.

Web-based and non-Web-based simulation software may be then deployed using either gUSE or CloudBroker directly.

Assuming that a workflow has been already developed, using gUSE the design options include:

- using the WS-PGRADE graphical interface,
- building a specific web-based graphical interface using the Application Specific Module (ASM) API that launches the workflow via gUSE Services,
- integrating access directly from an existing application's user interface using the Remote API that again launches the workflow via gUSE Services, and
- accessing the DCI-BRIDGE BES interface directly (the application must then implement its own equivalent gUSE Services as necessary)

Using CloudBroker the design options are:

- using the CloudBroker web interface directly,
- using the CloudBroker REST web service interface API,
- using the CloudBroker Java Client Library API, and
- using a Linux shell command line interface.

These options also do not preclude the possibility of Web-based and non-Web-based simulation Cloud-based applications being deployed separately from a "unifying" Portal. In CloudSME this is planned as being a "One-Stop-Shop" type Portal.

As indicated earlier in the paper, CloudSME is developing MSaaS by porting an initial set of industrial M&S tools including Simul8's discrete-event simulation environment (www.simul8.com), Ascomp's TransAT computational fluid dynamics application (http://www.ascomp.ch/transat/), Ingecon's 3D Scan Insole Designer tool and 2MORO's Bfly software for aircraft maintenance logistics.

The flexibility of this technology means that the developers of the above Cloud-based simulation software can easily add in different forms of licensing and Cloud resource costing. Moving to Cloud means that software can be more widely accessed (more users) and be charged on a per use basis (cheaper use costs). This also means that a specific deployment is not tied into a specific Cloud resource provider and can be easily redeployed to another provider if costs increase and/or service levels worsen (this is only limited by the availability of a specific CloudBroker/Cloud Resource provider adaptor). Future developments in gUSE may also mean that Cloud-based simulation software could also use *several* DCIs in parallel thus combining, for example, institutional desktop Grid resources and external Cloud resources.

For simulation software that requires simple hosting on a Cloud, i.e. deployed as SaaS, then direct interfacing between a Web-based or non-Web-based application (such as one installed under Windows) and CloudBroker might be sufficient. However, if the simulation software requires HPC then first developing and then deploying a workflow via WS-PGRADE/gUSE, and then using either the ASM or Remote APIs of gUSE as appropriate would rapidly give the simulation software the ability to use Cloud-based HPC resources to speed up its simulation task. The current set of applications, their deployment on our Platform, and examples of how the simulation software is making use of HPC on a Cloud will be discussed in a later paper and will be reported on at the Conference.

For further information on Cloud-based simulation software development using the CloudSME Simulation Platform see the CloudSME project's web pages at www.cloudsme.eu.

5 Conclusions

This paper has presented the CloudSME Simulation Platform and has discussed its major technologies and their functionality. As demonstrated with the successful development of Cloud-based scientific applications in the SCI-BUS project, the combination of workflow specification, Cloud management and multiple design options means that this Platform will assist simulation vendors in the rapid development of Cloud-based versions of their simulation software.

Acknowledgments

This work is part funded by the CloudSME—Cloud-based Simulation platform for Manufacturing and Engineering project No. 608886 (FP7-2013-NMP-ICT-FOF). The authors kindly acknowledge the efforts of the other beneficiaries on the project.

References

Liu, X., Qiu, X., Chen, B. and Huang, K. 2012. Cloud-Based Simulation: The State-of-the-Art Computer Simulation Paradigm. In Proceedings of the 2012 ACM/IEEE/SCS 26th Workshop on Principles of Advanced and Distributed Simulation (PADS 2012). ACM Press, NY, 71-74.

Hong, J. and Luo, J. 2011. Large-Scale Ranking and Selection Using Cloud Computing. In *Proceedings of the 2011 Winter Simulation Conference.* ACM Press, NY, 4046-4056.

Rossetti, M.D. and Chen, Y. 2012. A Cloud Computing Architecture for Supply Chain Network Simulation. In *Proceedings of the 2012 Winter Simulation Conference*. ACM Press, NY, 1-12.

Sakellari, G. and Loukas, G. 2013. A Survey of Mathematical Models, Simulation Approaches and Testbeds used for Research in Cloud Computing. *Simulation Modelling Practice and Theory*. 39:92-103.

Zhao, W., Peng, Y., Xie, F. and Dia, Z. 2012. Modeling and Simulation of Cloud Computing: A Review. In *Proceedings of the 2012 IEEE Asia Pacific Cloud Computing Congress (APCloudCC)*. IEEE Press. 20-24.

Kacsuk P., Farkas Z., Kozlovszky M., Hermann G., Balasko A., Karoczkai K., and Marton I. 2012. WSPGRADE/gUSE generic DCI gateway framework for a large variety of user communities. *Journal of Grid Computing* 2012; 10(4): 601-630.

Velte, T, Velte, A and Elsenpeter, R. 2009. *Cloud Computing, a Practical Approach.* McGraw-Hill, Inc.

Onggo, B.S.S., Taylor, S.J.E. and Tulegenov, A. *(to appear).* Web-Based Simulation as a Collaborative Modelling and Community Development Tool. In *Proceedings of the 2014 Operational Research Society Simulation Workshop*. UK Operational Research Society, Birmingham, UK.

Grance T. and Mell P. 2011. The NIST Definition of Cloud Computing. [Online]. Available at: http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf [Accessed 23/12/12].

Taylor, SJE, Fishwick, PA, Fujimoto, R, Page, EH, Uhrmacher, AM, Wainer, G.. Panel on Grand Challenges for Modeling and Simulation. 2012. In: *Proceedings of the 2012 Winter Simulation Conference*. ACM Press, NY., Article No. 232.

Taylor, SJE, Khan, A, Morse, KL, Tolk, A, Yilmaz, L, Zander, J. 2013a. Grand Challenges on the Theory of Modeling and Simulation. In: *Proceedings of the 2013 Symposium on the*

Theory of Modeling and Simulation. SCS, Vista, CA., Article 34.

Taylor, SJE, Balci, O, Cai, W, Loper, M, Nicol, D and Riley, G. 2013b. Grand Challenges in Modeling and Simulation: Expanding Our Horizons. In: *Proceedings of the 2013 SIGSIM-PADS Conference*, ACM Press, NY. 403-408.

Taylor, SJE, Brailsford, S, Chick, SE, L'Ecuyer, P, Macal, CM and Nelson, B. 2013c. Modeling and Simulation Grand Challenges: An OR/MS Perspective. In: *Proceedings of the 2013 Winter Simulation Conference*. ACM Press, NY. *To appear.*

BIOGRAPHIES

SIMON J E TAYLOR is the Founder and Chair of the COTS Simulation Package Interoperability Standards Group under SISO. He is the Editor-in-Chief of the UK Operational Research Society's (ORS) Journal of Simulation, the Series Editor of Palgrave-Macmillan's OR Essentials and is convening a new initiative for Grand Challenges in M&S. He was Chair of ACM's SIGSIM (2005-2008) and is a member of the SIGSIM Steering Committee. He is a Reader in the Department of Computer Science, Brunel University, and leads the Modelling & Simulation Group. He has published over 150 articles in modelling and simulation, has attracted substantial funding and has contributed to substantial cost savings and revenue in the M&S industry. His recent work has focused on the development of standards for distributed simulation and grid- and cloud-based simulation (www.cloudsme.eu) as well as the spread of the international Grid Infrastructure into Africa (www.eI4Africa.eu). His email address is simon.taylor@brunel.ac.uk and his web address is www.brunel.ac.uk/~csstsjt.

TAMAS KISS is a Reader at the Department of Business Information Systems, Faculty of Science and Technology, University of Westminster. He holds a PhD in Distributed Computing and his research interests include distributed and parallel computing, cloud, cluster and grid computing. He has been involved in several European research projects such as CoreGrid, EDGeS, EDGI, SHIWA, and SCI-BUS, and FP7 support action projects such as DEGISCO, ER-flow and IDGF SP. He led and coordinated application support activities and work packages in these projects. Currently he is Project Director of the FP7 CloudSME project that develops a cloud-based simulation platform for manufacturing and engineering SMEs. He co-authored more than 80 scientific papers in journals, conference proceedings and as book chapters. His email address is T.Kiss@westminster.ac.uk.

PETER KACSUK is the Director of the Laboratory of the Parallel and Distributed Systems in the Computer and Automation Research Institute of the Hungarian Academy of Sciences. He received his MSc and university doctorate degrees from the Technical University of Budapest in 1976 and 1984, respectively. He received the kandidat degree (equivalent to PhD) from the Hungarian Academy in 1989. He habilitated at the University of Vienna in 1997. He recieved his professor title from the Hungarian President in 1999 and the Doctor of Academy degree (DSc) from the Hungarian Academy of Sciences in 2001. He served as full professor at the University of Miskolc and at the Eötvös Lóránd University of Science Budapest. He has been a part-time full professor at the Cavendish School of Computer Science of the University of Westminster. He has published two books, two lecture notes and more than 200 scientific papers on parallel computer architectures, parallel software engineering and Grid computing. He is co-editor-in-chief of the Journal of Grid Computing published by Springer.

GABOR TERSTYANSZKY is a Professor of Distributed Computing at the Faculty of Science and Technology. He is the Director of the Centre for Parallel Computing. Before joining Westminster in 2000. His research in the field of distributed and parallel computing addresses the interoperability of large-scale distributed computing infrastructures and workflows. Working closely with diverse scientific communities, he has contributed to breakthroughs in interoperability research that have enabled researchers in different fields to more effectively exploit computing infrastructures in their scientific experiments and to execute computationally-hungry and data-demanding applications. it takes to produce research results, and producing many more research results than before. Professor Terstyanszky has supervised more than ten PhD students. He attracted in excess of £1.4 million in European and UK research funding to the University of Westminster. He has led or contributed to numerous collaborative research projects funded by European Commission Research Framework Programmes and by UK Research Councils. He has published more than 100 papers in academic journals and in proceedings of European and international conferences. He is regularly invited to join program committees of scientific conferences, and reviews scientific papers in academic journals on a regular basis. His email address is G.Z.Terstyanszky@westminster.ac.uk.