

## Guest Editor's Introduction: Special Issue on Cloud Computing Orchestration

Miguel Caballer · Germán Moltó ·  
Ignacio Blanquer

Published online: 17 January 2018  
© Springer Science+Business Media B.V., part of Springer Nature 2018

Cloud computing has provided users with access to computing, storage and network resources with unprecedented flexibility both from public Cloud providers and on-premises Cloud infrastructures. Cloud computing features several characteristics that explain its growth and popularity in the last decade. Maybe the most salient ones are flexibility, pay-per-use, and elasticity. The usage of Cloud technologies enables the software requirements to be defined by the user rather than by the infrastructure provider. This is an important point to ease the migration of scientific applications into the cloud. However, the resources have to be properly orchestrated (i.e., provisioned, configured and delivered) for users to leverage the benefits of these elastic infrastructures within their application domains. Complex application topologies, involving multiple Virtual Machines and multiple components with diverse software and configuration dependencies strongly benefit from the advances in orchestration.

Furthermore, in the last years the mainstream adoption of Linux containers, propelled by the popularity of Docker, enabled users to maintain customized execution environments, in the shape of lightweight container images instead of bulky Virtual Machine Images. Creating distributed applications based on containers required the ability to orchestrate a fleet of Docker container at scale, thus fostering the appearance of Container Management Platforms (CMPs) such as Kubernetes, Apache Mesos or Docker Swarm.

Therefore, orchestration tools are becoming fundamental to automate the deployment and management of application architectures on multiple Clouds. Hybrid scenarios that arise from simultaneously harnessing public and on-premises Clouds are challenges that remain to be solved for the efficient orchestration of applications on multi-Clouds. Also, the deployment of fleets of containers to achieve immutable application architectures on Clouds requires orchestration techniques.

This special issue focuses on novel solutions and innovative approaches that contribute to the field of Cloud Computing orchestration. This includes contributions related to automating the deployment of applications on Clouds, contributions related to the orchestration of scientific workflows and contributions to optimize container allocation for microservices in Cloud infrastructures.

A brief summary of the papers included in this special issue is provided for your convenience. In “Orchestrating complex application architectures in

---

M. Caballer · G. Moltó · I. Blanquer  
Instituto de Instrumentación para Imagen Molecular (I3M),  
Centro Mixto CSIC, Universidad Politécnica de Valencia,  
Valencia, Spain  
e-mail: micafer1@upv.es

G. Moltó  
e-mail: gmolto@dsic.upv.es

I. Blanquer  
e-mail: iblanque@dsic.upv.es

heterogeneous clouds”, Miguel Caballer and co-authors describe the orchestration approach for heterogeneous clouds that has been implemented in the INDIGO-DataCloud project. This orchestration model leverages the OASIS TOSCA standard as the modeling language and uses virtual machines and Docker containers in an homogeneous and transparent way providing consistent application deployment for the users on both on-premises and public Clouds. The paper entitled “QoS-Aware Orchestration of Network Intensive Software Utilities within Software Defined Data Centres” by Uroš Pačinsk and co-authors presents a new architecture for geographic orchestration of network intensive software components, in order to improve QoS in time-critical components, deciding where to run such components considering several factors modeled with qualitative machine learning techniques.

The deployment of complex applications in the Cloud is addressed in the paper “Occopus: a Multi-Cloud Orchestrator to Deploy and Manage complex scientific infrastructures”, by József Kovács and co-author, which presents an open source cloud IaaS orchestrator named Occopus. The applications are described using a domain-specific language agnostic to the infrastructure and the architecture and implementation details of the software are addressed and compared to other cloud orchestrator frameworks.

The requirement for high available applications in the Cloud is covered by the paper “Orchestrating the Deployment of High Availability Services on Multi-zone and Multi-cloud Scenarios” by R. Moreno-Vozmediano and co-authors, which introduces a new approach to achieve High Availability (HA) in multi-zone clouds through a matchmaking algorithm that considers affinity and anti-affinity properties of the resource. Rules are then created and analyzed in order to obtain a heuristic for the affinity placement.

In “The Flowbster cloud-oriented workflow system to process large scientific data sets”, Peter Kacsuk and co-authors detail a new workflow system, called Flowbster, aimed to create data pipelines which efficiently process very large data sets on cloud computing environments. It uses a choreography approach where all the nodes of communicate each other, without the need of a central service. Moreover, through the use of the Occopus cloud orchestrator, it supports several cloud backends. Finally, Flowbster includes a graphical user interface that eases the composition of workflows while hiding the complexity beneath the system.

The need to orchestrate fleets of containers for microservices architectures is addressed in the paper “Genetic Algorithm for Multi-Objective Optimization of Container Allocation in Cloud Architecture”, by Carlos Guerrero and co-authors. They propose a genetic algorithm for multi-objective optimization of container allocation in cloud architecture. The authors consider the optimizations focusing on a tight use of the resources and a reduction of the network overhead and system failure rate and present a model for cloud cluster, containers, microservices, and four optimization objectives.

Due to the importance of the container platforms in “Transparent Orchestration of Task-based Parallel Applications in Containers Platforms” Daniele Lezzi and co-authors have adapted the COMPSs runtime to support different container platforms such as Mesos, Docker or Singularity.

In summary, after a thorough review procedure, we aimed to compile in this special issue valuable contributions to the state of the art of Cloud computing orchestration, in order to illustrate the current progress and pending issues regarding this topic. Finally, the guest editors would like to thank the Editor in Chief of the Journal of Grid Computing, Prof. Peter Kacsuk, for allowing us to organize this special issue and for his continuous support during the process.