SeaClouds Project

D1.7.1
1st SeaClouds Scientific Workshop

Project Acronym	SeaClouds
Project Title	Seamless adaptive multi-cloud management of service-based applications
Call identifier	FP7-ICT-2012-10
Grant agreement no.	Collaborative Project
Start Date	1st October 2013
Ending Date	31st March 2016

Work Package	WP1, WP Management and Dissemination
Deliverable Code	D1.7.1
Deliverable Title	1st SeaClouds Scientific Workshop
Nature	Report
Dissemination Level	Public
Due Date:	M12
Submission Date:	10th October 2014
Version:	1.0
Status	Final
Author(s):	Antonio Brogi, PengWei Wang (UPI); Javier Cubo (UMA) Andrea Turli (Cloudsoft); Ivan Febles (ATOS)
Reviewer(s)	James Ahtes (ATOS); Ernesto Pimentel (UMA)

Dissemination Level

Project co-funded by the European Commission within the Seventh Framework Programme		
PU	Public	X
PP	Restricted to other programme participants (including the Commission)	
RE	Restricted to a group specified by the consortium (including the Commission)	
CO	Confidential, only for members of the consortium (including the Commission)	

Table of Contents

Executive Summary

A scientific workshop was promoted and organized by the SeaClouds project, and accepted as part of the ESOCC program, held on September 2, 2014 in Manchester, United Kingdom. The objective was to provide a forum to discuss challenges, solutions and perspectives of ongoing research and standards development activities aimed at enabling an efficient and adaptive management of service-based applications across multiple clouds.

Key takeaways from the workshop include:

- The diversity of cloud solutions and vendor lock-in strategies has increased the challenges for interoperability and portability, as well as the demand for standards to solve them.

- Vendor lock-in due to lack of standards is not just a customer concern, but can also be a risk to vendors that employ the strategy: pushing customers away to an increasing amount of flexible solutions, and rising significant upgrade costs to their platforms when the underlying infrastructure is changed.

- As more product vendors are taking on service provider roles, the benefits of agility, and the interoperability and standards that allow it, become increasingly relevant.

- Multi-cloud scenarios can bring clear benefits, such as best-fit cloud solutions on the level of application modules and their orchestration, increasing optimization and agility for adopters.

- Hybrid solutions are in rising demand, and many large enterprises require cloud solutions that are compatible with their existing cloud investment. Integrating off-site clouds with local already-deployed solutions drives the need for standards.

- Interoperability in the telco world requires focus on both northbound and southbound APIs, to not only mitigate vendor lock-in for customers, but allow services to be deployed on different cloud providers when required.

- Standards development is hampered by timing: too early and they can stagger innovation, while late introduction greatly diminishes their effectiveness. But in a fast market that is increasingly more focused on iterative cycles, standards must not try to "boil the ocean", but solve practical, more concrete targets.

- Standards in the PaaS segment are a larger challenge to conquer; the diversity of applications and their requirements is evident. Light standards can cover minimal and common targets, letting platform providers focus on their innovative capabilities.

1. Introduction

1.1 Objectives

The objective of the workshop was to provide a forum to discuss problems, solutions and perspectives of ongoing research activities, aimed at enabling an efficient and **adaptive management** of service-based applications **across multiple clouds**.

This forum, held within the European Conference on Service-Oriented and Cloud Computing (ESOCC) framework, was a good opportunity to discover other interesting initiatives in the cloud panorama. It also intended to generate **knowledge-transfer** by fostering discussion about important cloud-related topics like:

- Efficient monitoring and adaptive deployment of multi-cloud applications.
- Controlled migration of application modules across multiple clouds.
- Emerging standards supporting multi-cloud application management
- Case studies and best practices in multi-cloud applications

The workshop also aimed at contributing to the SeaClouds´ **dissemination strategy** by not only raising the awareness about the project, but also establishing relationships with other ongoing research projects, which may be useful for future **collaborations**. The workshop research contributions will be published by Springer in the Communications in Computing and Information Science (CCIS) series.

1.2 Program

The workshop "**Seamless Adaptive Multi-Cloud Management of Service-based Applications**"[1] was promoted by the ongoing European research project EC-FP7-ICT-610531 SeaClouds, and accepted as an ESOCC (European Conference on Service-Oriented and Cloud Computing)[2] workshop, which was held on **2-4 September 2014 in Manchester, United Kingdom.** The Program Chairs of the workshop (Antonio Brogi, UPI, and Ernesto Pimentel, UMA) formed an international Program Committee, which included fourteen internationally recognized experts from nine different countries. Five contributions were submitted in response to the call for papers. The originality and relevance of those contributions were evaluated with a peer-review process

[1] SeaClouds' 1st Scientific Workshop, "Seamless Adaptive Multi-Cloud Management of Service-based Applications": http://seaclouds.lcc.uma.es/
[2] ESOCC 2014: http://esocc2014.cs.manchester.ac.uk/

carried over by the Program Committee, which unanimously decided to accept only one of those contributions as a regular paper, and three other contributions were accepted as presentations of work in progress (with a shorter time for presentation and fewer pages in the proceedings).

In addition to the presentations of the contributed papers, the program included a shard opening keynote on the OASIS TOSCA initiative[3], an invited talk on the OASIS CAMP initiative[4], a round table on multi-cloud interoperability, and a session devoted to presentations of the development and initial results of 8 ongoing EU research projects. About 40 people from both academia and industry (e.g., IBM, Microsoft, ATOS and Cloudsoft) attended the workshop.

Agenda

The following is the final program of the SeaClouds workshop, which included a shared opening keynote, four presentations of the contributed papers, a round table on multi-cloud interoperability, an invited talk, and a session devoted to presentations of ongoing EU projects.

Time	Plan
09:30 – 10:30	**Opening Keynote** *Simon Moser (IBM)* From TOSCA landscapes to the Foundry – A walkthrough
11:00 – 11:30	*P. Jamshidi, C. Pahl (Dublin City University)* Orthogonal Variability Modeling to Support Multi-Cloud Application Configuration
11:30 – 12:30	*B. Surajbali (CAS Software AG), A. Juan-Verdejo (Stuttgart University)* A Marketplace Broker for Platform-as-a-Service Portability *H.S. Alqahtani, G. Kouadri-Mostefaoui (University of Bedfordshire)* Towards a classification of multiple-cloud computing concepts and terms *J. Carrasco, J. Cubo, E. Pimentel (University of Malaga)* Towards a flexible deployment of multi-cloud applications based on TOSCA and CAMP
12:30 – 12:45	**Invited presentations of ongoing research projects** SeaClouds - *F. D'Andria (ATOS)*
14:00 – 15:00	**Round table on multi-cloud interoperability** *K. Djemame (Univ. of Leeds), E. Gelenbe (Imperial College), A. Heneveld (Cloudsoft), S. Moser (IBM), C. Pezuela (ATOS), Ricardo Jimenez (Tech. Univ. of Madrid), Moderator: F. D'Andria (ATOS)*

[3] TOSCA: https://www.oasis-open.org/committees/tosca/
[4] CAMP: https://www.oasis-open.org/committees/camp/

15:00 – 15:30	**Invited Talk** *Alex Heneveld (Cloudsoft)* Going to CAMP via Apache Brooklyn
16:00 – 18:00	**Invited presentations of ongoing research projects** CoherentPaaS - *Ricardo Jiménez (Tech. Univ. of Madrid)* PANACEA - *E. Gelenbe, G. Gorbil (Imperial College)* PaaSage - *A. Rossini (SINTEF)* MODAClouds - *G. Casale (Imperial College)* ARTIST - *C. Pezuela (ATOS)* ECO_2Cloud - *U. Wajid (Univ. of Manchester)* ASCETiC - *K. Djemame (Univ. of Leeds)*

Program Committee

The Program Committee of the workshop (please see the table below) included fourteen internationally recognized experts from nine different countries (France, Germany, Ireland, Italy, Norway, Portugal, Romania, Spain, U.K.).

Program Chairs	- Antonio Brogi, University of Pisa, Italy - Ernesto Pimentel, University of Malaga, Spain
Program Committee	- Marcos Almeida, Softeam, France - Jorge Cardoso, University of Coimbra, Portugal - Martin Chapman, Oracle, Ireland - Javier Cubo, University of Malaga, Spain - Tommaso Cucinotta, Bell Laboratories, Alcatel-Lucent, Ireland - Francesco D'Andria, ATOS, Spain - Elisabetta Di Nitto, Politecnico di Milano, Italy - Christoph Fehling, University of Stuttgart, Germany - Nicolas Ferry, SINTEF, Norway - Alex Heneveld, Cloudsoft, U.K. - Simon Moser, IBM, Germany - Dana Petcu, West University of Timisoara, Romania

	- Achim Streit, Karlsruhe Institute of Technology, Germany - PengWei Wang, University of Pisa, Italy
Publicity Chair	- Michela Fazzolari, University of Pisa, Italy
Webmaster	- Adrián Pérez, University of Málaga, Spain

2. Outcomes

2.1 Multi-Cloud Solutions: Challenges and Benefits

Cloud computing has proven a major commercial success in recent years, with the appearance of many vendors to meet increasing demand. Cloud applications benefit from services provided in the cloud, such as database storage, virtual machine cloning, memory ballooning, etc. However, the development of cloud-based applications depends highly on the initial chosen cloud provider. Indeed, vendor lock-in is commonly identified as one of the most difficult issues with cloud platforms, limiting how customers can best leverage the cloud and burdening them with high switching costs. It may become practically unfeasible to change cloud arrangements once business processes and data have been assigned to cloud service providers [1]. Furthermore, there is the additional complexity for integrating multiple heterogeneous clouds and distributing services over several providers.

In the different talks and discussions during the workshop, both the benefits and challenges of adopting multi-cloud solutions were analysed. Different vendors are currently commercialising tools for the provisioning, management and automation of applications in leading public and private clouds, like Dell, BMC, Abiquo, etc. A promising perspective, opened by the availability of different cloud providers, is the possibility of distributing cloud applications over multiple heterogeneous clouds. In this sense, several aspects were exposed related to multi-cloud solutions.

Overcoming vendor lock-in and addressing the **interoperability** among heterogeneous cloud providers is one of the main ***challenges*** when a cloud application is deployed over multiple clouds. The workshop discussed various **standardisation** efforts that are currently ongoing, such as OASIS Cloud Application Management for Platforms (CAMP, [2]), DMTF Cloud Infrastructure Management Interface (CIMI, [3]), Virtualization Management (VMAN, [4]), or OASIS Topology and Orchestration Specification for Cloud Applications (TOSCA, [5]), to mention a few. Also, the large diversity of cloud providers requires the knowledge and **expertise in different platforms**, which is not an easy task. Finally, cloud adoption will be hampered if there will be no suitable way of managing data and applications across multiple clouds, and the need of **orchestration** is more evident when complex applications move to cloud environments.

Unfortunately, current cloud technologies only permit to deploy, monitor, and manage a whole application, and no support is available for synchronised deployment of application modules, unified monitoring, or for migrating the modules of an application and reconfiguring the rest of the application accordingly. Hence, services can only be deployed, managed and monitored on multiple clouds as stand-alone applications, and not as part of a composite application. Thus, in a scenario where a complex application is distributed on different cloud service providers, a solution is needed in order to manage and orchestrate the distribution of modules in a sound and adaptive way.

Regarding the ***benefits*** in multi-cloud solutions, such scenarios support both application developers (flexibility, agility, optimization, etc.) and cloud providers (differentiation, specialization, etc.), and can also expand the market, **reducing the entry barrier** for new players that offer, integrate and orchestrate such services across vendors. In addition, based on the concept of cloud-based services orchestration, automated arrangement, coordination, deployment and management of multiple services as a single, aggregated, complex application can be performed in an efficient and adaptive way, **without the need of modifying the code** of the services. This allows organisations to embrace cloud solutions and avoid risks of unreliability and vendor lock-in. Another topic discussed as a benefit is related to the possibility to deploy the application in a distributed manner, looking for the **most appropriate provider for a specific module** of the application, according to unique properties or requirements.

2.2 How Standards can Facilitate Interoperability

A workshop roundtable focused on how developing standards can alleviate multi-cloud challenges, and several outcomes of the discussion are summarized below.

Vendor Lock-In Risks to Industry
Most cloud providers recognize a lack of standards in IaaS. There have been some attempts to re-use Amazon de-facto standard APIs in open source solutions like Openstack and Eucalyptus, but that was clearly not enough. For dominant cloud players, standards are typically a threat and they prefer to not fill this gap as it often reflects a market advantage. However, for the health of the ecosystem, the lack of standards is a barrier to entry.

There are two main risks in having this strategy. One can be called an “external risk”: if you lock in your customers in this way, you could be inadvertently pushing them towards open alternatives, as flexibility between vendors is a primary concern for adopters. There is an “internal” risk as well. For example, if a cloud provider has IaaS and PaaS (based on its IaaS) offerings, any time the Infrastructure layer changes there is often a huge investment needed at the platform layer in order to adapt (ripple effect).

Multi-Vendor Integration
Standards can ease the interoperability between different systems. This is the case in a hybrid cloud scenario, for example, where a large enterprise needs to partially consume both public and private clouds. As many enterprises already have investments on cloud technologies, it is quite difficult to integrate existing solutions without such standards to facilitate interoperability.

Northbound/Southbound Interoperability for Telcos
Another interesting use case to stress the need for interoperability is in the telecom sector. In order to deliver added value services offered by telco operators, it is important that the operator offers a standard northbound interface to its customers (to avoid vendor lock-in) but also consumes a standard cloud API from its southbound

interface, to have the freedom to move its services on top of different cloud providers when required (interoperability).

Consideration on Standards Development

The roundtable presented open questions on the future of standards and their development process:

- How and which standards will help the new era of applications that are already cloud native? For example, it's hard to imagine if TOSCA or CAMP will be relevant in 5-10 years, but still the technology will help commoditize most of the application management strategies with standards. As usual, the introduction of standards too early brings a risk of freezing the market, which in the case of cloud, is particularly fluid and dynamic, while a late introduction of standards would be considered useless.

- Is a PaaS standard going to be more useful than IaaS? At the level of applications it seems harder than simply standardizing the Infrastructure resources' interfaces. Applications can be so diverse and the requirements so specific that it is hard to find a good interface for all types.

Standards and open source are key enablers for agility: certainly the customers but even the dominant players should be interested. Of course, unless dominant players invest time and money on creating standards, these gains will not take place. Those are technical enablers of an agile model but the business prospective plays also an important role to facilitate the adoption of standards: as long as the revenues from a vendor lock-in strategy will sustain the business, standardization adoption will be postponed. But revenues should not be considered as the only metric to monitor. In fact, many product companies are turning into service companies, and for their own benefit, in order to have a smoother and successful transition, they should consider adopting standards to improve internal and external agility.

Regarding TOSCA and CAMP, the workshop attendees think that there is space for both, as they are starting from different needs: TOSCA looks to solve the design time while CAMP is set to manage applications. There is an overlap on the deployment plan and live model, but this outcome could eventually converge.

Another interesting advantage given by both TOSCA and CAMP standards is the idea of the meta-model first: having an XML description of the plan is beneficial because it is parsable and machine-readable, so as to enable more automation and cloud adoption.

The round table also concluded that the standardization process is not sufficiently dynamic, as the market is still moving much quicker than standards development. But this asks the question of timely adoption: shall vendors wait for the "perfect standard", or can a more proactive implementation of smaller, partial solutions be an industry goal, particularly for a market segment that relies on iterative cycles.

2.3 Related Research on Multi-Cloud

The workshop brought together several EU research projects focused on interoperability and multi-cloud scenarios, presenting contributed research papers and initial outcomes of their development.

2.3.1 Participating EU Projects

The development and initial results of 8 ongoing EU research projects were presented at the workshop. In the following, we provide a brief introduction of these projects and their presentations. More details can be found from their presentation materials (see Annex 1) and their websites.

(1) The SeaClouds[5] project– Seamless Adaptive Multi-cloud Management of Service-based Applications

SeaClouds aims to tackle the problem of deploying and managing, in an efficient and adaptive way, complex multi-services applications over technologically heterogeneous clouds environments. This allows organisations to embrace cloud solutions and, at the same time, avoid risks of unreliability and vendor lock-in. The project will perform a seamless adaptive multi-cloud management of service-based applications, by developing Cloud Service Orchestrators and a set of tools to manage complex applications, thus avoiding the problem of cloud lock-in. This will be achieved by supporting the migration, replication, and distribution of modules that compose cloud-based applications over multiple and technologically diverse clouds offerings, by using a unified management API and universal metrics for monitoring and verifying functional and non-functional properties. The results delivered by SeaClouds will be compatible with the current standards related to cloud interoperability, such as OASIS TOSCA and CAMP.

(2) The CoherentPaaS [6] project - Coherent and Rich PaaS with a Common Programming Model

Today companies build applications on top of multiple data store technologies. This results in: 1) The loss of data consistency due to the lack of ACID transactions; 2) Queries across data stores have to be programmed and optimized manually. CoherentPaaS aims at solving these two pains in current cloud data store management by providing a unified cloud data management framework. It provides a full ACID coherent, ultra-scalable, and efficient environment integrating NoSQL, SQL and CEP data management technologies, while allowing application developers to program using a simple and powerful programming model and query language.

[5] SeaClouds project: http://www.seaclouds-project.eu
[6] CoherentPaaS project: http://coherentpaas.eu/

(3) The PANACEA[7] project - Proactive Autonomic Management of Cloud Resources

PANACEA aims to propose innovative solutions for a proactive autonomic management of cloud resources to address the growing complexity of the cloud, based on a set of advanced machine learning techniques and virtualization. By predicting anomalies (time to failure of cloud applications and DDoS attacks) before they occur, PANACEA will realize the following autonomic properties:

- Self-healing against anomalies by recovering from multiple node and link failures, and using proactive rejuvenation of applications and servers for preventing crashes and increasing the availability, predicting the threshold violation of response time of servers,
- Self-configuring by efficiently mapping user's requirements onto distributed clouds and configuring on-the-fly in the presence of anomalies,
- Self-optimizing using proactive migration of virtual machines from one cloud resource to another, maintaining the quality of service of end-to-end flows,
- Self-protecting using proactive reconfiguration of overlay networks to protect against DDoS attacks.

(4) The PaaSage[8] project - Modelling and Execution of Multi-cloud Applications

Cloud solutions are currently still insufficient and require a high level of expertise on the part of the developer and the provider to properly exploit the capabilities offered by cloud technologies. Cloud infrastructures are not standardized and porting an existing application to a cloud platform is still a very challenging task, leading to a strong interdependence between the client application and the cloud platform. Developing once and deploying on many clouds is not a viable proposition as things stand. This is the challenge that the PaaSage project will address. PaaSage will deliver an open and integrated platform to support both design and deployment of cloud applications, together with an accompanying methodology that allows model-based development, configuration, optimisation, and deployment of existing and new applications independently of the existing underlying cloud infrastructures. With this platform, the developers of enterprise systems can access services of cloud platforms in a technology neutral approach that abstracts the technical details while guiding them to configure their applications for best performance.

(5) The MODAClouds[9] project - Model-Driven Approach for Design and Execution of Applications on Multiple Clouds

Current Cloud's offer is becoming day by day wider providing a vibrant technical environment, where SMEs can create innovative solutions and evolve their services. Cloud promises cheap and flexible services to end-users at a much larger scale than

[7] PANACEA project: http://projects.laas.fr/panacea-cloud/
[8] PaaSage project: http://www.paasage.eu/
[9] MODAClouds project: http://www.modaclouds.eu/

before. However, Cloud business models and technologies are still in their initial hype and characterized by critical early stage issues, which pose specific challenges and require advanced software engineering methods. The main goal of MODAClouds is to provide methods, a decision support system, an open source IDE and run-time environment for the high-level design, early prototyping, semi-automatic code generation, and automatic deployment of applications on multi-clouds with guaranteed QoS. Model-driven development combined with novel model-driven risk analysis and quality prediction will enable developers to specify cloud-provider independent models enriched with quality parameters, implement these, perform quality prediction, monitor applications at run-time and optimize them based on the feedback, thus filling the gap between design and run-time. Additionally, MODAClouds will provide techniques for data mapping and synchronization among multiple clouds.

(6) The ARTIST[10] project - Advanced Software-based Service Provisioning and Migration of Legacy Software

For owners and developers of non-cloud software applications, ARTIST aims at facilitating the modernization of non-cloud software assets and businesses to the cloud. ARTIST offers a set of methods and tools which provide an end-to-end and assisted migration service to transform non-cloud software applications not merely to run on cloud but to take full advantage of cloud features. In addition, ARTIST provides pre-analysis of migration feasibility, analysing the software to gauge complexity, and to indicate cost and predict complications based on historical data. This data helps the customer make the investment decision. Following migration, ARTIST carries out validation and certification of migrated software and its functionality. Consequently, ARTIST forms a business investment and leads to lower costs and uncertainty, improved performance, greater customer satisfaction, more innovation and greater competitiveness. The steps of ARTIST migration process related to multi-clouds:

- Selection of the best cloud provider in terms of cost, performance and availability;
- Deployment of migrated application once it has been optimized;
- Monitoring of the migrated application once it is running in the cloud provider.

(7) The ECO_2Clouds[11] project - Experimental Awareness of CO_2 in Federated Cloud Sourcing

ECO_2Clouds investigates strategies that can both ensure effective application deployment on the cloud infrastructures and reduce energy consumption and CO_2 emissions. The project will provide a challenging and innovative approach to cloud computing service through different actions:

- Developing mechanisms to collect eco-metrics data and quantify the environmental impact of operations at infrastructure and application level;

[10] ARTIST project: http://www.artist-project.eu/
[11] ECO_2Clouds project: http://eco2clouds.eu/

- Developing a comprehensive approach to multi-cloud application deployment;
- Developing evaluation mechanisms and optimization algorithms.

ECO_2Clouds will develop and test an application deployment strategy which will optimize the environmental impact. In addition it will propose an optimization model for cloud infrastructures.

(8) The ASCETiC[12] project - Adapting Service Lifecycle towards Efficient Clouds

The increased usage of ICT, together with growing energy costs and the need to reduce greenhouse gases emissions call for energy-efficient technologies that decrease the overall energy consumption of ICT. So far, efforts in eco-efficiency have mainly targeted hardware and data centre technologies. Less attention has been given to software eco-efficiency. ASCETiC is focused on providing novel methods and tools to support software developers aiming to optimise energy efficiency and minimise the carbon footprint resulting from designing, developing, deploying, and running software in clouds, which include:

- Development of models for green and efficient software design, supporting sustainability and high quality of service levels at all stages of software development and execution.
- Development and evaluation of a framework with identified energy efficiency parameters and metrics for cloud services.
- Development of methods for measuring, analysing, and evaluating energy use in software development and execution, complementing quality measures.
- Energy and quality efficiency integration into service construction, deployment, and operation leading to an Energy Efficiency Embedded Service Lifecycle.

2.3.2 Additional Research Contributions

In this workshop, four contributions related to multi-cloud were accepted after the peer-review process, one as a regular paper and other three as presentations of work in progress. The accepted research contributions will be published by Springer in the Communications in Computing and Information Science (CCIS) series, and their full presentations at the workshop can be found in Annex 1. In the following, we present a brief introduction and discussion for the four research contributions.

[12] ASCETiC project: http://www.ascetic-project.eu/

(1) Orthogonal Variability Modelling to Support Multi-Cloud Application Configuration

Cloud service providers benefit from a vast majority of customers due to variability and making profit from commonalities between the cloud services that they provide. Recently, application configuration dimensions has been increased dramatically due to multi-tenant, multi-device and multi-cloud paradigm. This challenges the configuration and customization of cloud-based software that are typically offered as a service due to the intrinsic variability. In this research, the authors present a model-driven approach based on variability models originating from the software product line community to handle such multi-dimensional variability in the cloud. They exploit orthogonal variability models to systematically manage and create tenant-specific configuration and customizations. They also demonstrate how such variability models can be utilized to take into account the already deployed application parts to enable harmonized deployments for new tenants in a multi-cloud setting. The approach considers application functional and non-functional requirements to provide a set of valid multi-cloud configurations.

(2) A Marketplace Broker for Platform-as-a-Service Portability

Platform as a Service (PaaS) has become a strong technological solution in particular for small medium enterprises (SMEs) to achieve cost savings and rapid time to market of their software solutions. However, for SMEs how to choose the PaaS provider becomes a bottleneck due to the number of offerings each PaaS vendor offers. Another challenge often faced by enterprises is how to match their software system requirements to PaaS offerings. Furthermore, searching for the best PaaS offering is even more difficult when enterprises need to migrate their existing software solutions from one PaaS vendor to another. In such cases, the portability of the enterprises data and application components becomes cumbersome. If their software solutions are bound to a specific PaaS provider, enterprises suffer from vendor lock-in. This study addresses these portability challenges by proposing a high-level architecture to ease the portability of software solutions over PaaS vendors taking into consideration the various life-cycle stages, such as identifying and analysing PaaS offerings in the market; selecting the best PaaS offering according to organisation's requirements; and the deployment, management, and monitoring of the software solutions.

(3) Towards a Classification of Multiple-cloud Computing Concepts and Terms

In order to enhance the quality of the delivered services, dependency avoidance, and operational costs optimization, the multiple-cloud computing emerged. The complexity of multiple-clouds causes a lot of opacity in the definitions and classifications as well. This study attempts to provide an accurate definition for multiple-clouds and classification for the available types based on well-defined criteria. It also outlines the differences between the terms that have been used to define and refer to the multiple-cloud computing concept.

(4) Towards a Flexible Deployment of Multi-cloud Applications Based on TOSCA and CAMP

Cloud Computing platforms offer diverse services and capabilities with own features. Hence, the provider services could be used by end users to compose a heterogeneous context of multiple cloud platforms in order to deploy their cloud applications made up of a set of modules, according to the best capabilities of the cloud providers. However, this is an ideal scenario, since the cloud platforms are being conducted in an isolated way by presenting many interoperability and portability restrictions, which complicate the integration of diverse provider services to achieve a heterogeneous deployment of multi-cloud applications. In this ongoing work, the authors present an approach based on model transformation to deploy multi-cloud applications by reusing standardization efforts related to the management and deployment of cloud applications. Specifically, using mechanisms specified by both standards, TOSCA and CAMP, they propose a methodology to describe the topology and distribution of modules of a cloud application and to deploy the interconnected modules over heterogeneous clouds.

3. Conclusion

From the organizational point of view, this workshop has been a success not only for the number of attendees and their positive feedback, but also for the enthusiastic participation around the round table and the wide range of research projects presented in the afternoon sessions.

Regarding the content of the workshop, there were several aspects discussed by participants, focused around the multi-cloud domain and developing standards to enable the needed interoperability. **A collection of key takeaways can be found in the Executive Summary.**

The importance of the **cloud standards** and their relevance for the health of the cloud ecosystem, was analysed during the different sessions of the workshop, as well as presentations on OASIS' CAMP and TOSCA.

Multi-cloud was another topic of the workshop. The evolution experienced by multi-cloud in the last years, which has evolved from a market trend to an increasing business need, generated an interesting debate. The added complexity generated by multi-cloud is challenge, but outweighed by the benefits provided (e.g. enhanced QoS, dependency avoidance, operational costs optimization, etc.).

The success experienced by this format, will set the foundations for future workshops (both industrial and scientific) to be held in 2015, the second year of the SeaClouds project. Initially, due to the success for the first edition, the ESOCC conference will be again the target venue to collocate the event.

Annex: Additional Material

The additional materials are included in a separate PDF file attached to this deliverable, which contains all of the presentation materials at the workshop.

References

1. D. Petcu, G. Macariu, S. Panica, and C. Craciun. Portable Cloud applications: From theory to practice. Future Generation Computer Systems, 29(6):1417-1430, 2013.

2. OASIS: Cloud Application Management for Platforms v1.1 (CAMP 1.1), 2014. Available at http://docs.oasis-open.org/camp/camp-spec/v1.1/camp-spec-v1.1.pdf

3. DMTF DSP0263, Cloud Infrastructure Management Interface (CIMI) Model and REST Interface over HTTP, An Interface for Managing Cloud Infrastructure v1.0.0, 2012. Available at http://dmtf.org/sites/default/files/standards/documents/DSP0263_1.0.1.pdf

4. DMTF Virtualization Management (VMAN), 2014. Available at http://dmtf.org/standards/vman

5. OASIS: Topology and Orchestration Specication for Cloud Applications v1.0 (TOSCA 1.0), 2013. Available at http://docs.oasis-open.org/tosca/TOSCA/v1.0/TOSCAv1.0.pdf