PROJECT FINAL REPORT

Grant Agreement number: 317715

Project acronym: PAASAGE

Project title: PAASAGE MODEL BASED CLOUD PLATFORM UPPERWARE

Funding Scheme: FP7 ICT Call 8

Period covered: from 01/10/2012 to 30/09/2016

Name of the scientific representative of the project's co-ordinator, Title and Organisation:

Philippe Rohou

Tel: +33 4 9715 5306

Fax: +33 4 9238 7822

E-mail: philippe.rohou@ercim.eu

Project website address: http://www.paasage.eu/

Table of Contents

1. Final publishable summary report

1.1. Executive summary

The purpose of PaaSage (http://www.paasage.eu/) is to enable application developers - and users executing applications - to see a virtualisation of available CLOUD platforms so that the application is deployed optimally across multiCLOUDs, including scalable elasticity, optimal resource usage, minimised cost and respecting security, privacy and performance constraints. The PaaSage architecture – while remaining consistent in the essentials – has changed slightly during the project to accommodate some changes caused by development decisions. In particular, even more work has been done on the user interface aspects in Y3 and Y4. The CAMEL language has been developed further and – most importantly – has been re-validated by the use cases being described in CAMEL and the software components of PaaSage executing using CAMEL. Of note is the joint group working on improvements to the TOSCA standard utilising CAMEL input – this demonstrates the intellectual content of PaaSage. The major effort has been in the development of the components of the PaaSage platform and their integration. Particular emphasis has been on the development of Cloudiator to replace Cloudify for cross-platform deployment, the development of multiple reasoners/solvers and the development of a security environment. While the system has been demonstrated running at end Y2, successive improvements to the middleware and executionware have necessitated re-integration and deployment of the platform several times in Y3 and Y4 and the maintenance of the open source on the OW2 facility. In parallel 9 industrial workshops have been organised in Y3 and Y4 (4 individual, 5 jointly with other events) with appropriate training materials available on the project website. 85 publications have been produced and 67 presentations made so disseminating widely information on PaaSage. Team members have collaborated with other major CLOUD and related projects. The website received a major redesign for Y4 and various press releases have been produced and disseminated. The final result is a production-ready PaaSage system capable of dealing with applications characterised in CAMEL, and associated user constraints (such as cost and elapsed time). PaaSage plans possible deployments, selects the optimal one and executes it – with the ability to re-deploy or even re-plan if end-user constraints are in danger of being violated. PaaSage is being exploited by the use case partners in their real-world applications so providing direct business benefit in reduced cost, faster application deployment and flexibility in utilisation of CLOUD platform offerings. The software is offered open source, using the OW2 facility (so providing increased assurance) and exploitation is expected through widespread take-up and utilisation in the community. The commercial benefit is reduced cost, flexible deployment of applications across multiCLOUDs and improved confidence in application developers and end-users in utilising CLOUD platforms. The social benefit comes from the applications themselves ranging from better and faster re-scheduling of flights (LSY) through to better management of the collection and distribution of human milk in a distributed computing environment (EVRY) with better and faster simulations using high performance computing for the automobile industry (ASCS) improved (cost, speed) financial investment management (IBSCY), and better management of engineer deployment (BEWAN) as other examples. A particular example for scientific benefit is the ‘data farming’ and ‘hyperflow’ approaches for utilising computing resources in applications from AGH. It is clear that PaaSage is applicable widely across commercial and research domains, is made

available open source for maximum take-up and is capable of realising commercial advantage for application developers and end users.

1.2. Summary description of project context and objectives

The main overall objective is:

To 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.

And in detail:

- To participate to the design and standardisation of an open, powerful, and expressive modelling language together with the MODAClouds project for Cloud independent modelling of enterprise systems with the preferences and constraints, focusing on architectural styles and characteristics of the Cloud computing paradigm.
- To provide an intelligent Integrated Development Environment (IDE) supporting the modelling language and supporting the developer in the task of optimising the application.
- To provide mappers and generators allowing a modelled Cloud application to be deployed in a distributed environment, interacting with multiple Cloud providers if needed.
- To define metadata relevant for Cloud services, and provide mechanisms to acquire the metadata and critical performance indicators from running applications and to reuse the historical metadata available on the services in the application design and deployment.

The context of PaaSage is the business-driven need to be able to deploy an application across one or several CLOUD platforms automatically. Although some competing middleware has been developed during the lifetime of PaaSage none supports the range of capabilities of PaaSage. Thus PaaSage is unique and technologically advanced, and is being used by project partners for their business objectives as indicated in the subsequent sections of this report.

The technological achievements are documented in following sections. The architectural ideas have proven to be robust in practice and the software developed is leading-edge. The decision to make it open source provides the means for sustainability and improvement by the community – without precluding commercial exploitation of the software with value-added features. The development of the CAMEL language is a great success causing the leading standard in the area – TOSCA – to agree to joint work to improve TOSCA utilising the CAMEL design ideas and implementation. The use of a MDDB with knowledge gained from history of utilisation of PaaSage and external experts' input is novel and makes PaaSage more 'intelligent' in its choices of deployment strategy. The novelty of

PaaSage in all these aspects is reflected in the large number of peer-reviewed scholarly papers and presentations; the interest generated is recorded in website accesses and social media messaging.

A series of industrial workshops has showcased PaaSage with both explanations about the architecture and demos of working commercial implementations. Associated training has allowed end-users and application developers to utilise the system.

Finally, PaaSage has been a management challenge. With 18 partners including commercial companies ranging from large (Lufthansa) to SMEs (BeWAN), ICT companies such as the SME Flexiant, research laboratories such as INRIA, FORTH and STFC and universities such as Oslo and Ulm the consortium is large and diverse. The advantage is an array of diverse leading-edge talent; the disadvantage is difficulty of coordination. This challenge has been met by management coordination through the ERCIM office (Philippe Rohou) and scientific/technical coordination through a consultant (Keith Jeffery) working for ERCIM. In addition, Geir Horn, University of Oslo, was appointed and resourced to coordinate the software development and integration.

1.3. Benefits and technical findings

Benefits

The benefits of PaaSage come in several user communities:

1. For users: the PaaSage platform provides a way for a user to launch a pe-existing application, together with parameters for non-functional requirements for the particular run of the application (like elapsed time, cost) without considering the platform(s) upon which the deployment will occur. Furthermore, PaaSage ensures the running application is scaled appropriately to meet the Serie Level Agreements and/or conditions for the particular run of the application;
2. For application developers: The PaaSage platform provides the basis for application developers to characterise the application using the CAMEL language in such a way that it can be deployed optimally. This means that (a) it is not necessary to write new applications in a particular way; (b) it is not necessary to write new applications or modify existing applications for a particular target platform; (c) legacy applications can be 'wrapped' with CAMEL to enable appropriate deployment;
3. For IT managers: The PaaSage platform allows the manager to avoid committing or being bound to any one particular platform thus providing independence and dynamic choice based on platform characteristics such as performance, price, security policy etc.
4. For the CLOUD Community: the PaaSage platform provides a model of middleware for CLOUD provider independence and an evolving knowledge base of experience from deployments and added expertise collected from experts in the community;
5. For the software engineering community: the PaaSage project provides the middleware open source through OW2 so that software engineers can re-use the software for their own projects
6. For the CLOUD research and development community: PaaSage provides a language for characterising applications, latforms and deployments (CAMEL), a model for middleware, an approach that is novel and leading-edge and software to be re-used to use as a basis for further ad more advanced research and development.

Technical Findings

PaaSage is a project of 48 months with 19 partners tackling leading-edge issues in CLOUD Computing. At the start of the project it was agreed not to attempt to deal with (a) detailed aspects of security, privacy, trust in the deployment of the application (although the platform has appropriate security for the PaaSage processes of profiling, reasoning, deployment, adapting); (b) management of the locality, distribution, partitioning, migration and caching of datasets (which is the subject of the MELODIC project starting December 2016) and the new call INFRA-12-2017.

A major technical finding is that it is possible to constrict a language (CAMEL) – using different domain-specific languages – suitable for the purpose of application characterisation and deployment across CLOUD platforms. Furthermore, we discovered that the work on CAMEL is beneficial to improving the OASIS TOSCA language for CLOUDs.

A second major technical finding is the complexity of the software – both upperware and executionware – required to achieve the objectives of PaaSage. The complexity required not only excellent software design and development but also strong coordination – for which purpose we appointed a 'technical dictator' working on behalf of the overall scientific coordinator. The software design changed several times during the project not least because of changes in CAMEL needed to express the requirements of the use cases. With hindsight a more formal description of the use cases in year 1 of the project and development of CAMEL thereafter would have been better, but the reviewers of the original proposal decided that the major CAMEL development would be in the parallel MODAClouds project and so the CAMEL schedule was determined – to a large extent – outside of PaaSage.

The third major technical finding is the need for some standard set of parameters provided by the CLOUD platforms for monitoring to allow middleware to optimise (a) in upperware time preparing for deployment; (b) at execution time for dynamic adaptation and/or redeployment.

Related to the above, the fourth major technical finding is the importance of the MDDB and associated generation of and extraction of knowledge. The MDDB provides to PaaSage the 'view of the world' and is the repository of data, information and knowledge of deployments of applications across CLOUD platform, assisted by expertise donated by the community. This is an important resource for the future. However, to achieve recoding of all the entities/attributes required the MDDB had to integrate several different data models (related to the DSLs in CAMEL and others) and this was not always easy or smooth.

Where results can be found

- PaaSage institutional website: http://www.paasage.eu/
  (Including all PaaSage public deliverables)
- PaaSage open source forge: https://projects.ow2.org/bin/view/paasage/
- PaaSage future developments: Check the EU project MELODIC, starting 01/12/2016.

2. Description of the main S&T results/foregrounds

2.1. CAMEL

MDE is a branch of software engineering that aims at improving the productivity, quality, and cost-effectiveness of software development by shifting the paradigm from code-centric to model-centric. MDE promotes the use of models and model transformations as the primary assets in software development, where they are used to specify, simulate, generate, and manage software systems. This approach is particularly relevant when it comes to the modelling and execution of cross-cloud applications (i.e., applications deployed across multiple private, public, or hybrid cloud infrastructures). This solution allows exploiting the peculiarities of each cloud service and hence optimising performance, availability, and cost of the applications.

Models can be specified using general-purpose languages like the Unified Modeling Language (UML). However, to fully realise the potential of MDE, models are frequently specified using domain-specific languages (DSLs), which are tailored to a specific domain of concern.

In order to support the modelling and execution of cross-cloud applications, PaaSage developed the Cloud Application Modelling and Execution Language (CAMEL) DSL, which allows to specify multiple aspects of cross-cloud applications, such as provisioning, deployment, service level, monitoring, scalability, providers, organisations, users, roles, security, and execution.

CAMEL supports models@run-time, which provides an abstract representation of the underlying running system, whereby a modification to the model is enacted on-demand in the system, and a change in the system is automatically reflected in the model.

The Topology and Orchestration Specification for Cloud Applications (TOSCA) is a specification developed by the OASIS consortium, which provides a language for specifying the components comprising the topology of cloud applications along with the processes for their orchestration. TOSCA supports the specification of types and templates, but not instances, in deployment models. In contrast, CAMEL supports the specification of types, templates, and instances. Therefore, in its current form, TOSCA can only be used at design-time, while CAMEL can be used at both design-time and run-time.

As part of the joint standardisation effort of PaaSage, MODAClouds, and ARCADIA, SINTEF presented the models@run-time approach to the TOSCA technical committee (TC) and proposed to form an ad hoc group to investigate how TOSCA could be extended to support this approach. The TC welcomed this proposal and approved the formation of the Instance Model Ad Hoc group in October 2015. The group is currently co-led by Alessandro Rossini from SINTEF and Derek Palma from Vnomic. The work performed in this group will guarantee that the contribution of CAMEL will be partly integrated into the standard.

2.2. Prototype system proving architectural concept

The PaaSage platform is essentially an IDE (Integrated development environment) driven my models (in CAMEL), utilising a knowledge base to improve deployment leading to managed execution in a multiCLOUD environment with elastic scalability. This is illustrated simply:

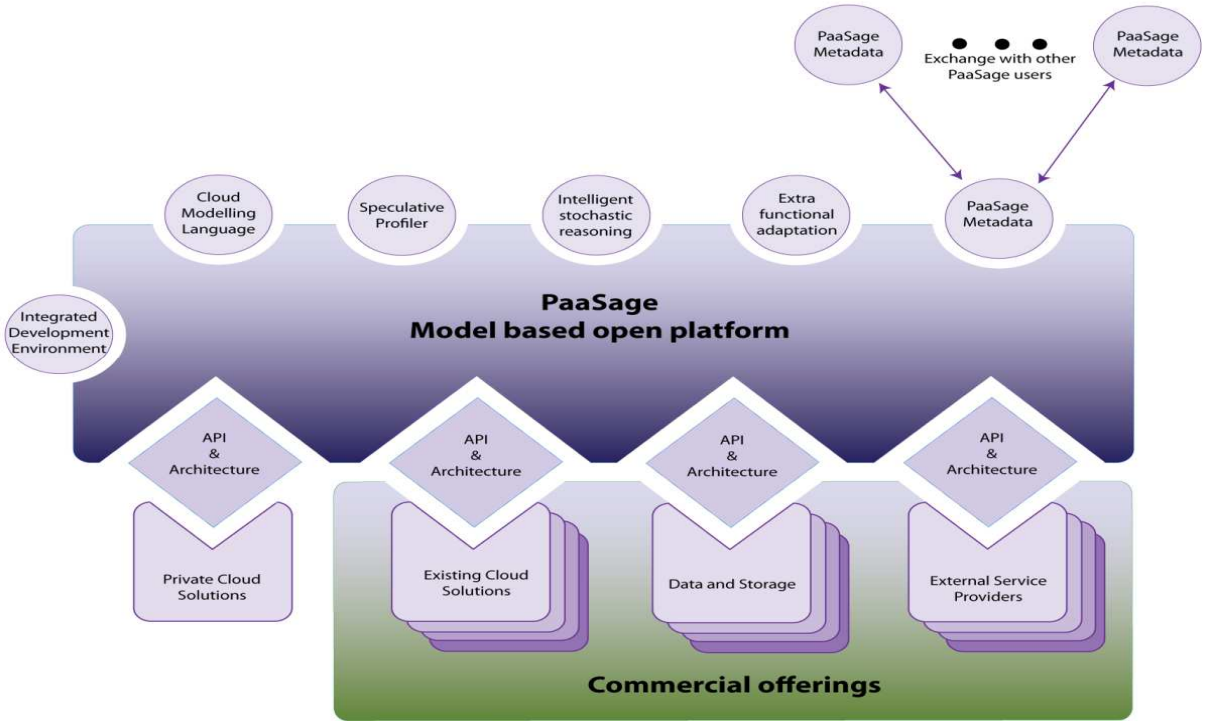

Software teams within and across the various partners developed the components in the upperware, metadata database (knowledgebase with social network) and the executionware with deployment and monitoring. As explained elsewhere, the software was developed as open source for maintain ability and community improvement leading to increased exploitation.

Needless to say, the software development and integration was difficult. This is because PaaSage is absolutely leading edge and nothing like it had been attempted before. The success is due to the skill and dedication of the software teams.

The prototype system - which will be developed further both in other projects and in the plans of various partners – achieved two major areas of success: (1) the work on assessment and validation in WP1 demonstrated its value and (2) it was used by all 7 demonstrators and worked effectively for their use cases.

The software components of the upperware (profiler, reasoner, adapter) and of the executionware are described in detail in the relevant deliverables D3.1.2 and D5.1.2 respectively.

2.3. OW2 open source

PaaSage took a novel approach in utilising open source technology and becoming a member of OW2 to promote the software and interact with the software development community (in contrast to the MDDB which interacts with the Cloud application DevOps community).

The pace of technological innovation, the constant search for differentiation against competitors, and the quest of higher investment returns through increasing adoption of its own technical solutions are some of the reasons for Cloud providers to lock their costumers into their commercial offering. Therefore, Cloud providers tend to multiply specific APIs and deployment models and have no drive to propose and adhere to standards.

In contrast, the ability to overcome constraints from Cloud providers and to freely choose Cloud offering is strongly sought by Cloud users. The impulse to get out of Cloud lock-in must come from a community effort like the PaaSage Research project where every community member contributes to reach a result that would have been too costly to achieve independently. The software outcome of such effort is highly desired by commercial partners of the community because, in addition to regaining control and freedom over software deployment and use in the Cloud, it allows to provide new services in Cloud offering selection to their customers. In the case of PaaSage, users will also benefit from the community expertise thanks to metadata gathered by the PaaSage Platform on services and providers. In the end, a shared an open source community result will also be beneficial to Cloud providers because it will drive them toward delivering better and more competitive services.

As the community effort is driven by Cloud users, it is crucial that they have a shared interest in continuous evolution of software like the PaaSage platform allowing them to restore freedom over Clouds. This therefore leads to the choice of an open source license to deliver the PaaSage Platform as well as open source methodologies to implement it. Regarding the choice of open source license, the PaaSage consortium has evaluated them according to three principles. Firstly, the license of the PaaSage Platform must be compatible with the licenses of software components brought by project partners or external components integrated into the PaaSage Platform. Secondly the license should avoid capture by an organisation for its sole benefit and protect investment done in PaaSage project. It implies that terms of the licenses force to release further improvement to PaaSage Platform as open source. Thirdly, the chosen license should not restrict commercial use of PaaSage Platform and should allow its integration into commercial closed source solutions. An additional requirement for the licence was that it should be well known in the developer communities to facilitate easy adoption of PaaSage. Studied licences were the GNU Lesser General Public License Version 3.0 (LGPL 3.0), the Eclipse Public License (EPL) the Mozilla Public License Version 2.0 (MPL2) and the Common Development and Distribution License (CDDL). MPL2 was adopted because it is compatible with all the used licenses and dual-licensing these components for PaaSage under MPL2 do therefore not create problems.

Stating that source code is available under open source license is far from enough to ensure success and adoption by an ever

-growing community. An open source community must also be formed and nurtured around the initial development team constituted by the PaaSage project consortium. To this end, the PaaSage project sought early involvement with open source communities with the goal to attach to an existing

community. The project also followed the “release early and often” principle of software engineering in order to make the project known as soon as possible outside of the project consortium.

To trigger adoption, open source software must also be available and packaged to be easily installed and used. Such goal relies on tools and services, notably source code repositories for the latter. Two kinds of repositories can be identified: the ones that just store code, like Github for example, with no particular constraint or aim other than commercial activity and the one that belongs to communities like Apache hosting modules, Eclipse hosting and OW2, aiming facilitating the development, deployment and management of open source projects and related development and management tools. Both categories of repositories provide the tools to develop build and package software according to open source good practices. Community repositories also have the advantage of active, if not vibrant, existing communities which are potentially interested by new open source projects hosted in their repositories. This led the PaaSage project to move towards a community repository and to choose the one with goals closer to those of the project. The mission of OW2 is to promote the development of open-source middleware, generic business applications, and Cloud Computing platforms. Not only it matches with PaaSage's objectives but also OW2 is the *only* European software community and hosting platform, thereby fitting PaaSage’s commitments to promote European Research and Economy.

Starting more than one year ahead of the end of the project, OW2 became the repository of the PaaSage Platform. It implied to move gradually the Platform development integration and packaging from the infrastructure of partners in the consortium to the OW2 infrastructure in order to ensure that tools required for evolution of the PaaSage Platform are ready and used before the end of the project and available also to developers outside of the project consortium. To allow a seamless transition to open source community, after the end of the project, CETIC became a member of OW2.

2.4. MDDB and knowledge base

This key success is described in some detail because of is complexity and its core positioning between (a) the Cloud devops community and PaaSage providing knowledge and expertise; (b) the monitoring of execution and parameter storage for use in the reasoner of the upperware to improve deployment; (c) the user interface and the PaaSage platform (allowing close connection with the social network); (d) the user and the system i.e. the security environment. The MDDB also has associated convertors to collect and maintain metadata using different standards.

Contributions in the context of WP4 have led to the development and delivery of a set of foundational components for the PaaSage project, according to the scope and main objectives of the main WP tasks, T4.1 and T4.2, respectively. In the sequel, we outline the main achievements in the context of the Metadata Database (MDDB), Model Importer, Knowledge Base (KB), Security Extension, Identity Provider (IdP), Policy Enforcement and Policy Decision Points (PEP/PDP), Administration API, Social Network, DSL-to-DSL transformation, and their integration with the overall platform, indicating contributing partner organisations where appropriate.

Metadata Database (MDDB): The MDDB follows a client-server architecture that relies on two main components developed in the context of WP4: the *Connected Data Objects (CDO*[1]*) Server* and *CDO Client*, used for the population and management of the MDDB.

The *CDO Server* is the server-side component responsible for the actual management and accessing of the underlying database, which is organised in the form of a model repository and realised via the Eclipse CDO technology. The management of such a model repository includes management actions performed on models, like the storage (persistence), updating, deletion and retrieval of models. The model repository can be realised via different types of CDO stores, while the selection of a certain type is handled via appropriate configuration of the CDO Server. In PaaSage, we currently support two main store types exhibiting different characteristics: HibernateStore[2] has the advantage of the use of a higher-level query language which allows the issuing of SQL-like queries on the meta-model level (i.e., by referring to classes and relationships of the respective meta-model), but does not realise all potential features of a CDO store. DBStore[3] realises all these features but offers only a SQL query capability, which translates to the need for the query issuer to know the actual mapping of models into the underlying database schema. Both types of stores can work on top of different database management systems (DBMS), while the selection of a certain DBMS is performed via appropriate CDO Server configuration. Finally, it must be highlighted that the CDO Server incorporates the CDO technology's security feature. This means that CDO sessions between respective server and client components can be secured such that the access to models is amenable to the policies specified. This CDO feature employs a RBAC access control model where roles can be mapped on policies, which can regulate access on certain repository paths or for certain meta-models or meta-model classes. Whether a CDO Server runs in a secure mode is another available point of configuration for this component. Policies should be appropriately set before the access to the underlying model repository is initiated. This component has been developed by FORTH.

The *CDO Client* is the client-side component in MDDB interactions. It is able to connect with a CDO Server via CDO Sessions, secured via a simple authentication method (based on username-password key pair mapping to the authentication information of the user for which the component acts). Each session, depending on its type, enables either a read or also a write view over the respective model repository. Both view session types are amenable to the policies specified for the authenticated user. The CDO Client component offers the following model management functionality:

- Storage of models in the MDDB or the file system
- Retrieval of models from the MDDB. Either the sole model can be retrieved or also other models pertaining to the cross-references included in that model. Moreover, cross-referencing can be handled in a recursive manner until all the models retrieved in the end do not have unaddressed cross-references.
- Loading of models from the file system. The models can be in XMI or textual form (only for CAMEL models), thus mapping to both the abstract and concrete syntax of the actual language, respectively.
- Transformation from XMI to textual form for models and the opposite

[1] eclipse.org/cdo
[2] https://wiki.eclipse.org/CDO/Hibernate_Store
[3] https://wiki.eclipse.org/CDO/DB_Store

- Complete export of the models stored in the MDDB into the file system in the form of a zip file for back-up purposes
- Import of MDDB backups to restore the content of an MDDB instance
- Deletion of models/objects
- Querying of models based on the query language supported by the underlying MDDB store
- Creation of read or write views over MDDB
- Closing of views
- Closing of CDO sessions

This component has been developed by FORTH.

Model Importer: It is an add-on to the MDDB that (a) enables generating and importing basic models which can be exploited in order to build new CAMEL models; (b) creates the appropriate structure in the CDO model repository which conforms to the model storage guidelines exemplified in [D4.1.2], thus enabling the enforcement of the security add-on solution (see analysis below). The generation of basic models involves the creation of three main models: (1) a security model which includes the definition of a set of security controls and respective domains; (2) a metric model which incorporates the definition of a set of metrics and properties plus other non-functional terms (e.g., metric formulas); (3) a location model which has been produced via processing the content of the FAO ontology[4] and which includes a hierarchical model of continents, sub-continents and countries. Use case CAMEL models can also be imported via this component. Their processing is also capable of splitting organisation from other kinds of CAMEL models which was needed initially due to the way the security add-on solution functions. This component can be considered as initialisation code to be run before the actual MDDB/CDO repository becomes available to the rest of the PaaSage components (e.g., in order to store new organisation models or other kind models where the latter can be produced during the execution of the PaaSage deployment workflow). This component has been developed by FORTH.

Knowledge Base (KB): The KB is an add-on to the MDDB that can be used for producing added-value knowledge over the MDDB via the enforcement of particular rules. Two main realisations of the KB were developed. The first realisation relies on REST-based interactions. In this case, FORTH has developed a REST-client that can communicate with the REST-service that realises the KB functionality. The second realisation comprises a standalone software program able to realise itself both the KB functionality and wrap it in the form of a KB client interface. For both realisations, specific rules have been developed to infer the knowledge of which applications or components match to each other and which are the best deployments for an application or its components. These realisations enable the complete management of the KBs and their respective sessions, where a session can be considered as a long-lasting interaction session between a KB and a respective client, which leads to the firing of rules, the production of the respective facts and thus the possibility to retrieve such facts or query them. Apart from wrapping an existing KB (Drools) into these two realisations, another novelty of the KB comes with the capability to operate directly on the underlying MDDB exploiting the full potential of the information that has been stored there.

[4] www.fao.org/countryprofiles/geoinfo/en/

Moreover, via appropriate experimental evaluations, an optimised domain model has been produced enabling the optimisation of the query performance over the KB. The added-value knowledge of KB is currently exploited by the CP Solver component (FORTH) in WP3, which enables it to save solving time by fixing parts (component to VM offering mappings) of the deployment optimisation problem to be solved. The KB component has been developed by FORTH.

Security extension: While the CDO security feature sufficiently addresses access control over the models stored in the MDDB, it comes with the following issues: (a) substantial modelling effort needs to be performed in order to model the respective access control policies; (b) it does not obviously cater for securing other facilities of the platform, like the respective APIs provided. To address these issues, a security solution (that relies and builds on the CDO security feature) was developed to secure access to all kinds of resources in the PaaSage platform, not only the model ones (MDDB). This solution is extensively described in [D4.1.2][5] and published in two conference and one journal publication. To summarise, some of the main features of this solution are: (a) it enables both programmatic and web-based authentication; (b) it supports both model and service-based authorisation; (c) it enables organisations (i.e., clients of the PaaSage platform) to control the way their resources can be accessed via the supply of respective access control rules over respective roles (either internal or external to that organisation) in the form of a CAMEL organisation model part; (d) it supports default mappings between particular roles and access control policies, thus reducing the modelling effort of an organisation; (e) it supports the mapping between CAMEL access control policies to CDO security model ones while also increasing the security level as organisations and their users do not have direct access to the underlying CDO security model; (f) it supports standards and in particular SAML[6].

The security-oriented extension to MDDB comprised the introduction of new components, namely the IdP, the PEP/PDP points and the Administration API, as well as the update of existing ones, namely the CDO Server and Client. In the following, we shortly analyse both types of extensions/modifications that have been performed.

CDO Server/Client Extensions: To support this solution, the following updates to the CDO Server and Client were needed: While the CDO security feature was enough for enabling the access to model resources, the CDO Server implementation offered by Eclipse was faulty in some cases (class and package/meta-model access control policies). In addition, as SAML-based authentication has been realised, the CDO-based authentication had to be enhanced in order to exploit SAML tokens instead of simple logic-password key pairs. As such, we modified the CDO Server to address these two issues. Moreover, the CDO client had to be modified to enable the establishment of CDO sessions with the CDO server via a SAML token and not a username-password key pair. Both modifications were performed by FORTH.

Identity Provider (IdP): This component has been developed by AGH in order to support the SAML-based authentication of users both in a programmatic and web-based manner. While the IdP has

[5] Kirkham, T., Kritikos, K., Kryza, B., Magoutis, K., Massonet, P., Papoulas, C., Korozi, M., Leonidis, A., Ntoa, S., Sheridan, C., Innes, A., Imrie, D. A.: D4.1.2 – Product Database and Social Network System (September 2015).
[6] https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security

relied on the simpleSAMLphp[7] framework, that framework was extended to be able to connect to the MDDB. This was due to the fact that security-oriented information for organisations is stored in the form of organisation models in MDDB. As such, any authentication check should have been performed against this MDDB. To enable secure access only to the respective information required for the authentication, a special user and role was developed in the CDO Server/repository and mapped to this component. Upon successful authentication, the IdP returns a specific SAML token which can then be exploited for authorisation purposes.

Policy Enforcement and Policy Decision Points (PEP/PDP): This complex component is responsible for the authorisation over service-based resources in the platform. It has been realised by STFC and has relied on the OpenAZ[8] library. The main logic here is that the PEP component checks the respective SAML token and request and retrieves the respective relevant policies. Then, these policies are checked against the user token by the PDP component which decides whether the request can be granted or not. The PEP/PDP implementation have relied on a separate XACML[9] database. This means that access control policies specified in CAMEL are transformed to XACML policies before they are stored in this database. Again, for the retrieval of the respective policy information from MDDB and its subsequent transformation, a special user and role have been created in the CDO security model.

Administration API: Apart from the capability to specify policies via CAMEL models, FORTH has developed an alternative capability that enables programmatic management of the policy-oriented information for organisations. The rationale here is that such a management is performed by administrators that may prefer to use a specific programmatic interface rather than a specific language. The administration API relies on the CDO Client for the communication with the MDDB and the respective modification of the policies for both the CAMEL models and the CDO security model. As such, it also guarantees the complete mapping between CAMEL organisational models and the CDO security model. More details about this mapping can be found in [Kritikos et al. 2016[10]] and [D4.1.2]. As this API should be accessed by an administrator of a client organisation, the CAMEL organisation model for that organisation should have been already provided preferrably in its limited form comprising only one admin user for which also the respective user credentials have been specified. Otherwise, it would not be possible to have access to this API as the respective user would not be reflected in the underlying CDO security model. As such, the minimal CAMEL organisation model is actually handled/enforced by the super-administrator of the PaaSage platform instance upon registration of that organisation in this platform instance, again through use of the Administration API. The original CDO security model includes the description of this user, its respective role and corresponding policies mapped to it.

Social Network: This component is responsible for offering social network facilities, such as messaging as well as threads & discussions along with platform-specific capabilities in the form of model storage and editing, initiation of PaaSage deployment workflows and browsing of execution

[7] https://simplesamlphp.org/samlidp
[8] http://www.openliberty.org/wiki/index.php/OpenAz_Main_Page
[9] docs.oasis-open.org/xacml/3.0/xacml-3.0-core-spec-os-en.html
[10] K. Kritikos, T. Kirkham, B. Kryza, P. Massonet. Towards a security-enhanced PaaS platform for multi-cloud applications. FGCS, 2016 (accepted)

histories. Its ultimate goal is the promotion of the use of the PaaSage platform as well as the exchange of knowledge. This component has been realised based on the Elgg[11] framework while it exploits the CDO Client facilities for the accessing and storage of the CAMEL models. It has been well integrated with the PaaSage platform API in order to be able to initiate PaaSage deployment workflows. The core SN functionality has been developed by FORTH and widely disseminated via open access to a publicly accessible site[12] and a journal publication[13]. STFC (with the assistance of FORTH) has developed the web-based CAMEL editor which can be used for the editing of models of the SN users. This editor has been integrated in the SN. STFC has also developed an elgg extension for SAML also providing OAuth support for the SN.

DSL-to-DSL Transformation: In order to be able to support standards as well as enable users to bring their existing models to the platform possibly specified in different languages than CAMEL and/or enable exploiting CAMEL as an entry point for external platforms using these standard languages, WP4 produced transformation code to support the mapping between CAMEL and standard language models. Three main languages were targeted: two are widely used in cloud computing, i.e., (1) TOSCA and (2) WS-Agreement[14], while the third concerned support for the CERIF[15] standard to cover organisation modelling. TOSCA covers the deployment models of cloud applications (i.e., mapping to the deployment meta-model in CAMEL) while WS-Agreement concerns the capturing of service level agreements and respective templates (mapping mainly to requirement and metric meta-models in CAMEL). CERIF is a standard supporting the modelling of research organisations, including the roles and users involved in them. Initially, we aimed for a bidirectional transformation. However, due to resource restrictions as well as the focus on more core functionality in WP4, it was decided to support one direction of transformation. BEWAN undertook the mapping between TOSCA and CAMEL, FORTH the mapping between CAMEL and WS-Agreement, and AGH the mapping from CERIF to CAMEL. The TOSCA-to-CAMEL transformation by BEWAN was realised via the ATL mapping framework. Only some of the TOSCA elements (*RelationshipTemplates*, *NodeTemplates*, *RequirementsTemplates* and *CapabilityTemplates*) are currently supported mainly related to the topology of an application. The CAMEL-to-WS-Agreement transformation was developed by FORTH via using a custom mapping approach based on JAVA and the EMF modelling facilities of Eclipse. Many CAMEL model kinds are taken into account, including requirement, metric, deployment and organisation models. WS-Agreement is fully supported (all elements are actually covered) and there was no need to modify its respective schema. The CERIF-to-CAMEL transformation by AGH relied on a Java-based custom mapping approach that takes as input a CERIF xml model (organisation description) as well as a CERIF vocabulary/taxonomy and creates a partial organisation model that covers the CAMEL organisation meta-model classes of *Organisation*, *User*, *Role*, *RoleAssignment*, *PaaSageCredentials*, and *CloudProvider*. An interesting feature of this transformation code is that it is accompanied by the capability to store the respective CAMEL organisation model produced in the MDDB CDO repository.

[11] https://elgg.org
[12] http://socialnetwork.paasage.eu
[13] K. Magoutis, C. Papoulas, A. Papaioannou, F. Karniavoura, D. Akestoridis, N. Parotsidis, M. Korozi, A. Leonidis, S. Ntoa, C. Stephanidis, Design and implementation of a social networking platform for cloud deployment specialists", SpringerOpen Journal of Internet Services and Applications (JISA), vol 6 no 19, pp. 1-27, 2015.
[14] https://www.ogf.org/documents/GFD.107.pdf
[15] www.eurocris.org/cerif/main-features-cerif

Metamodels. Work on certain CAMEL meta-models was undertaken in the context of WP4. This was driven by two main factors: (a) the security solution add-on and its respective modelling requirements; (b) the platform (Social Network & WP3) requirement to support the recording of execution histories. To this end, three main meta-models have been created. The *organisation meta-model* was developed by FORTH and AGH. It captures certain type of information about organisations, such as: (a) generic details about an organisation, (b) specific details concerning the types of cloud services offered in case it is a cloud provider, (c) the roles, users and assignment of users to roles for that organisation and (d) the respective access control policies over models and services involved in the PaaSage platform. Users are also mapped to certain cloud credentials via which the platform can perform deployments on their behalf. They are also mapped to a unique username-password pair that can be used for their authentication in the platform, thus being able to support single sign-on. The *security meta-model* was developed by FORTH and CETIC in order to cover mainly the filtering of the cloud provider space via high- and low-level security requirements in the form of security controls and SLOs, respectively. As such, it is related to a specific minimum but sufficient security-related concepts which are properly integrated with the rest of the concepts in CAMEL. For instance, security metrics are special cases of CAMEL metric meta-model metrics, while security SLOs are special types of CAMEL requirement meta-model SLOs which only include conditions over security metrics or properties. Potentially, security conditions can also be exploited in the definition of scalability rules (see CAMEL scalability meta-model). An interesting extension to CAMEL in order to support adaptation rules which can more meaningfully include security conditions is described in [Kritikos et al. 2016]. The *execution meta-model* has been developed by FORTH with the main goal to cover the specification of deployment episodes. These episodes cover the respective execution history of an application under a specific deployment. Such execution history includes the specification of the measurements that were produced for the application, the SLO assessments that have been performed under these measurements and the respective scaling actions that have been executed in an attempt to still satisfy the SLO requirements posed. It also includes other details like the start and end time of the deployment and its overall cost. All this information can be exploited to infer added-value knowledge via the KB. It is also utilised by the Social Network in order to enable users to browse the execution history of their deployments.

2.5. 7 demonstrators

System demonstration, evaluation and showcasing of the features of the PaaSage platform by using the different use case applications developed throughout the PaaSage project is a necessary part to proof the concepts of PaaSage. Therefore, the system demonstrators can be regarded as the closing bracket of the requirements definition exercise documented in the deliverable *Final Requirements (D6.1.2)* done in M24.

Besides the demonstration of the features of PaaSage, another focus of the demonstrators is to encourage interest in take-up and further exploitation of the PaaSage framework. Since industrial integrators and providers are a key target for the PaaSage dissemination activities, this focus is particularly relevant. These companies will either deploy or offer at a large scale the PaaSage technology, making the benefits available to a wide range of industrial users.

The seven demonstrators created by the use case partners can be divided into different categories (partner names in brackets):

- Industrial demonstrations (BEWAN and LSY)
- eScience demonstrations (ASCS, AGH, USTUTT)
- Public sector demonstrations (EVRY, supported by FORTH)
- Financial sector demonstrations (IBSCY, supported by UCY)

In-depth information about each use case application and the demonstrators can be found in deliverable D7.1.1.

3. Potential Impact, Dissemination and Exploitation

3.1. Impact

During PaaSage, the consortium has elaborated system demonstration, evaluation and showcasing of the features of the PaaSage platform by using the different use case applications developed throughout the PaaSage project.

Besides the demonstration of the features of PaaSage, another focus of the demonstrators was to encourage interest in take-up and further exploitation of the PaaSage framework. Since industrial integrators and providers are a key target for the PaaSage dissemination activities, this focus is particularly relevant. These companies will either deploy or offer at a large scale the PaaSage technology, making the benefits available to a wide range of industrial users.

Through these demonstrators, PaaSage has demonstrated:

- **cross-cloud deployment**: The ability to distribute the application components over multiple, diverse cloud environments was demonstrated for one usage scenario of the LSY use case application. This scenario can be described as follow-the-sun and the idea is to distribute the application over different geographical regions worldwide, e.g., Europe, Asia and North America. This way, different users can work on the same airline schedule during their regular working hours. To provide a reasonable user experience parts of the application (e.g., the RM's) must be located in data centres near to each client. These data centres are not necessarily hosted by the same cloud provider; hence a cross-cloud deployment is necessary.
- **interoperability across heterogeneous platforms**: PaaSage supports accessing different cloud providers in a uniform way. Together with the cloud provider independent modelling capabilities offered by the Cloud Application Modelling and Execution Language (CAMEL), this prevents the application owner from a vendor lock-in. This has been demonstrated in the LSY scenario
- **knowledge-based, learning optimisation**: PaaSage allows deployment optimisation based on requirements and scalability rules. It uses data collected from social network and platform own learning capabilities. In the EVRY use case scenario, in order to always be compliant with the different milk banks laws and regulations and maintaining the service as cost-effective as possible, PaaSage will help us track and propose changes in the deployment architecture if needed. The requirements are modelled in CAMEL language. In the Computer Aided Engineering (CAE) use case scenario lead by ASCS, scalability rules are defined in the CAMEL model. For the CAE use case, when the CPU load is above a predefined threshold, new worker VMs will be deployed by PaaSage automatically to join the parallel simulation.

PaaSage has contributed to the TOSCA standard through CAMEL.

3.2. Dissemination

PaaSage has been very active in disseminating towards both the scientific community and the target industries, as well as the general audience. The main objectives of the dissemination activities were:

- To initiate the building of the Open Source Community;
- To initiate the industrial uptake of PaaSage.

Major initial decisions were:

- PaaSage and CAMEL are the brand names for marketing the key outcomes of the project;
- OW2 is the open source repository for accessing the PaaSage code.

3.2.1. Website

At the beginning of the project, a web site has been created (www.PaaSage.eu) under the form of a project presentation (project objectives, structure, consortium, organisation, documents…).

During the project year 4, the PaaSage Web site has been completely re-designed as a 'product Web site' in order to support a wide adoption of the PaaSage platform by partners and third-party users and developers.

The PaaSage Website product refers globally to four different web sites:

- The main PaaSage web site (http://www.paasage.eu/) is the main entry point for target visitors (developers and users);
- The PaaSage social network (http://socialnetwork.paasage.eu/) is the collaboration space for the PaaSage users community;
- The PaaSage open source repository (https://projects.ow2.org/bin/view/paasage/) hosted by OW2 is the interaction space for PaaSage developers and provides access to the source code of the platform.
- The CAMEL web site (http://camel-dsl.org/) providing an introduction along with the documentation of CAMEL (Cloud Application Modelling and Execution Language).

In addition to these four major web sites, several partners have built and are maintaining specific web sites (such as the ExecutionWareToolset Cloudiator, containing documentation, a tutorial section and a tutorial video at https://cloudiator.github.io).

The main PaaSage web site (http://www.paasage.eu/) appears now as follows:

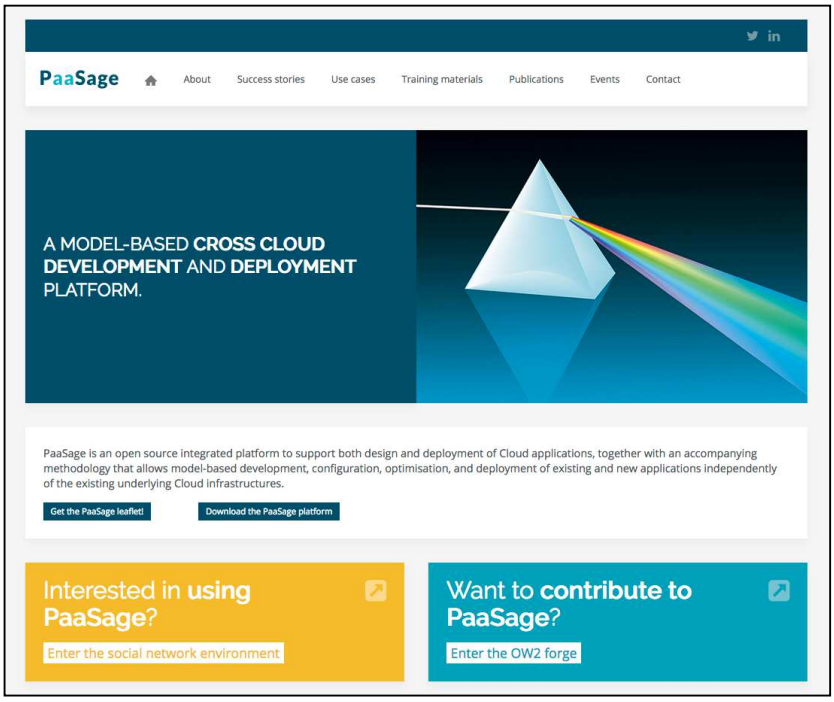

The visibility towards a wider audience was completed by a number of press releases:

- Announcing the launch of the PaaSage project (October 2012)
- Announcing the partnership between PaaSage and OW2 to publish the platform on the AppHub marketplace (June 2015)
- Announcing the release of the first version of the PaaSage platform (September 2015)
- Announcing the industrial workshop in Stuttgart (February 2016)
- Announcing the industrial workshop in Budapest (September 2016)
- Reporting about the Future of Clouds conference in Budapest (September 2016)

PaaSage has set up accounts on Twitter and LinkedIn that have been very active in discussing major events (meetings, product releases, success stories).

A PaaSage Wikipedia page (https://en.wikipedia.org/wiki/PaaSage) has been set up. The PaaSage product has been registered with the Open Cloud Directory (http://www.ocdirectory.org/) as a way to promote the prototype in a wider community.

Complementary to the consortium-wide activities, many partners (and especially the industrial members) have been very active in promoting their participation to the PaaSage project towards their ecosystems and customers' network.

3.2.2. Supporting materials

Supporting communication materials have been created during the project lifetime:

- An initial leaflet was created early on, and was redesigned into a professional-looking product-oriented leaflet, intended to support the promotion of the PaaSage platform towards candidate adopters (from industry and academia).
- A roll-up to be used in industry events and conferences.
- A poster that can be tuned for specific scientific events.

3.2.3. Scientific Talks & Papers

The academic partners of PaaSage have actively published all along the project duration. A total of 88 scientific papers have been produced, and a high number of scientific presentations have been performed. A list of published papers is available in an Appendix of this report.

At the Cloud Forward 2015 conference in Pisa, the paper "Beyond Mere Application Structure: Thoughts on the Future of Cloud Orchestration Tools", by Jörg Domaschka, Frank Griesinger, Daniel Baur and Alessandro Rossini won the Best Paper Award.

Published PaaSage scientific publications are openly available at http://www.paasage.eu/scientific-publications with links to abstracts and full papers (when publicly available).

3.3. Workshops and training

3.3.1. Scientific Workshops

PaaSage partners have been very active in contributing to scientific workshops.

During project year 1

Three collaboration workshops were organised with the MODAClouds and ARTIST projects:

- an organizational meeting at the EC Internet of Services Collaboration meeting in Brussels (10/2012);
- a workshop on Multi-Cloud Applications and Federated Clouds co-organized by PaaSage in Prague (4/2013). More information at http://www.paasage.eu/training-events/events/8-multi-cloud-2013.
- a special session focused on the activities of the three projects at the 7th Annual Advanced School on Service-Oriented Computing in Hersonissos, Greece (7/2013). More information at http://www.paasage.eu/training-events/events/27-summersoc-13.

Additionally, PaaSage presentations were organised within the summer school SummerSOC'13

During projet year 2

PaaSage participation has been provided to following events:

- the Krakow Grid Workshop 2013 (CGW'13, 04-05.11.2013, Krakow)
- the EU-Mexico Workshop: Exploring common research interests in the Future Internet and Cloud Computing (20.03. 2014, Athens)
- MultiCloud 2014 (03-05.04.2014, Barcelona), in conjunction with CLOSER 2014, The 4th International Conference on Cloud Computing and Services Science

- the European Open Cloud Collaboration Workshop (15.05.2014, Brussels, organised by the OCEAN project).
- the 8th Advanced Summer School in Service Oriented Computing (30.06-05.07.2014, Hersonissos)
- the 2nd International Workshop on Cloud Service Brokerage (CSB 2014, 02-04.09. 2014, Manchester)
- the SeaClouds workshop, held in conjunction of ESOCC 2014 (http://seaclouds.lcc.uma.es/, 02.09.2014, Manchester)
- the Open Data Center Alliance Forecast 2014 (22-24.09.2014, San Francisco)

During projet year 3

PaaSage co-organized or participated in the following scientific workshops:

- Cracow Grid Workshop (CGW), Krakow, Poland, October 27-29, 2014 (organization of a PaaSage session, see http://www.cyfronet.krakow.pl/cgw14/
- RELATE Winter School on Engineering and Provisioning of Cloud Applications – Research and Entrepreneurship, Würzburg, 17.-20. February 2015
- AppHub workshop, 9 March 2015, Brussels, within the CloudScape VII event (see http://www.apphub.eu.com/bin/view/Main/)
- CloudWatch Concertation Meeting 2015, 25 March 2015, Brussels (See http://www.cloudwatchhub.eu/turning-cloud-research-innovative-software-services)
- 9th Symposium and Summer School On Service-Oriented Computing, Heraklion, Crete, Greece, 28 June - 3 July 2015
- Co-organization of the 2nd EAI International Conference on IoT as a Service (iotaas.org), October 26–27, 2015 in Rome, Italy, (by Kostas Magoutis (FORTH) and Benny Mandler coordinator of COMPOSE) and participation of PaaSage in this conference through a paper titled “Privacy Aware On-Demand Resource Provisioning for IoT Data Processing” by Tom Kirkham, Arnab Sinha, Nikos Parlavantzas, Bartosz Kryza, Paul Fremantle , Kyriakos Kritikos, Benjamin Aziz.

During project year 4

PaaSage co-organized or participated in the following scientific workshops and tutorials carried out during the period (this list does not include all the numerous conferences attended by PaaSage members):

- UULM presented the ExecutionWare Toolset Cloudiator and its interplay with PaaSage at the EU Projects Track of ESOCC'16 (http://esocc2016.eu/, September 2016).
- AGH with HLRS presented the e-Science use case of PaaSage at the EU Projects Track of ESOCC'16 (http://esocc2016.eu/, September 2016).
- AGH organized a tutorial session on “Parameter studies on heterogeneous computing infrastructures with the Scalarm platform” at the International Conference on High Performance Computing and Simulation (HPCS’16, Innsbruck, http://hpcs2016.cisedu.info/, July 2016)
- FORTH organized a PaaSage training session at the SummerSOC'16 Summer school with the participation of several partners (http://www.summersoc.eu/, 29.06.2016)

- UiO participated and promoted PaaSage at the Swedish Cloud Control Workshop (http://cloudresearch.org/workshops/9th/, June 2016)
- INRIA co-organized the OpenStack WorkShop in Lyon (June 2016), in conjunction with the French OpenStack community, and contributed a PaaSage session (https://openstack.fr/2016/05/25/openstack-workshop-lyon-2016-3/).
- UULM presented PaaSage and the ExecutionWare Toolset Cloudiator in a booth at the International Supercomputing 16 Exhibition in Frankurt, Germany, (http://www.etp4hpc.eu/en/events/international-supercomputing-conference-isc16_306.html, June 2016).
- AGH Presented PaaSage and the e-Science use case at the Workshop on Container Strategies for Data & Software Preservation that Promote Open Science, University of Notre Dame, USA (https://daspos.crc.nd.edu/index.php/workshops/container-strategies-for-data-software-preservation-that-promote-open-science, May 2016)
- EVRY, SINTEF and UiO organised a PaaSage event at the Software 2016 conference (industrial audience) in Oslo (https://software2016.sched.org/event/4kzR, February 2016)
- CETIC demonstrated PaaSage at the AppHub SQuAT Fest in Brussels (http://www.digitalmeetsculture.net/article/apphub-squat-fest/, 26/01/2016)
- A position paper 'PaaSage – making cloud usage easy' was submitted at the EU-Brazil Cloud Connect event in Rio (Brazil, http://eubrazilcloudconnect.eu/position-paper/9-paasage-making-cloud-usage-easy, December 2015).
- UCY presented PaaSage at the 8th IEEE/ACM International Conference on Utility and Cloud Computing (UCC 2015, http://cyprusconferences.org/ucc2015/?page_id=290, December 2015).
- A PaaSage workshop entitled "A discovery workshop through PaaSage" was organised and presented by UiO at the Paris Open Source Summit, 8-19 November 2015, http://opensourcesummit.paris/preinscription-conferences.html).
- CETIC presented PaaSage and some PaaSage use cases at the Cloud Computing and Data Center BeLux forum (17 November 2015, Brussels, http://cloudcomputingforum.brussels/en/).
- CETIC presented PaaSage at the UNIWAN Software Defined eXperience, on 20 November 2015 (Belgium).
- SINTEF arranged an event on Big Data and Cloud aimed at industry, where EVRY contributed with a keynote (Oslo, 20 October 2015).
- EVRY presented the milkbank use case at the European Milk Bank Association (EMBA) annual congress (http://www.biomedia.net/embacongress2015, Lyon, France, 08-09.10.2015).
- Three PaaSage papers were given at CF2015 (Cloud Forward 2015 – a conference organised by the HOLACloud Project, see http://cf2015.holacloud.eu/), Pisa 6-8 October 2015.
- ERCIM represented PaaSage at the Cloud Standards Collaboration (CSC) workshop organised by ETSI in Brussels (01-02.10.2015, http://csc.etsi.org/).

3.3.2. Industrial Workshops

A key action in the defined PaaSage product launch strategy is the organisation of a series of Industrial Workshops, specifically targeted for supporting the building of PaaSage communities of users and of developers and for accelerating the take-up of PaaSage by industrial stakeholders.

In order to maximize the impact of these industrial events, they have been organised by PaaSage industrial partners and targeted to that specific industrial partner ecosystem in its region.

Two kinds of industrial events have been organised:

- Joining a related event, and having a specific session / booth for presenting PaaSage:
  - In Cyprus, coordinated by IBSCY, with the support of UCY (March 2015).
  - In UK, coordinated by STFC, with the support of Flexiant (March 2015)
  - In Belgium, coordinated by CETIC (2 events in November 2015)
  - In Norway, coordinated by EVRY, with the support of SINTEF and UiO (February 2016)
  - In France, coordinated by INRIA (at the OpenStock Workshop Lyon 2016, June 2016)
- PaaSage-specific events, synchronized with a PaaSage consortium meeting:
  - In Belgium, coordinated by BE.WAN (September 2015)
  - In Germany, coordinated by ASCS (April 2016)
  - In Norway, coordinated by EVRY (September 2016)
  - In Hungary, coordinated by LSY (September 2016)

For the second kind of events (that are PaaSage-specific), the event organisation as well as the related communication campaign were in the hands of a PaaSage industrial partner. The workshop duration was typically ½ day, in the afternoon. As the industrial workshops were co-located together with a PaaSage consortium meeting, PaaSage partners were able to attend and contribute to the workshop presentations without additional cost, and also to gain experience for subsequent workshop organisation.

The typical agenda of Industrial Workshops was as follows:

- Welcome and introduction
- Presentation of PaaSage & CAMEL (objectives, architecture, main features…)
- Use case presentation (focus on the industrial partner use case)
- Demo or open/interactive booth
- Panel discussion, Q/A
- Cocktail and networking.

3.3.3. Training Materials

During the project duration, the consortium has completed and published a large set of training materials that are to be used in conjunction with the use case reports, and intended to ease the adoption of the PaaSage platform by third-parties (system admins, business users, application developers and IT architects, software engineers or simple users).

The training materials are structured along following categories:

- PaaSage in a nutshell (presentation)
- The PaaSage social network (videos)
- Modelling with CAMEL (documentation, presentation)
- Execution deployment and usage (video)
- Downloading, installing and configuring PaaSage (presentation)

Training materials are available on the PaaSage Web site at http://www.paasage.eu/training-materials.

3.4. Exploitation plans

In PaaSage, we have identified the following non technological exploitation paths: contribution to standards, exploitation by academic and CRO partners and the open source exploitation.

- **Contribution to standards**: PaaSage contributes to the work of TOSCA, valorising and exploiting results created by the project
- **Exploitation by academic and CRO partners and open source:** The value proposition of PaaSage is "Write Once, Deploy Anywhere". In order to make PaaSage futureproof, the following sustainability plan has been put in place.
    - The PaaSage platform is open source and available on OW2 community.
    - A governance model ensures the sustainability of the PaaSage Platform with clear roles and responsibilities
    - The academic and CRO's will promote PaaSage in further project proposals.
    - When possible, the PaaSage components listed in section "6 Exploitable results" will be enhanced and will contribute to the improvement of the PaaSage Platform as a whole. Indeed, the list of project results and components composing the PaaSage platform have been listed and the partners have stated their interest in maintaining these components so as to ensure the sustainability of the whole platform.

In terms of individual exploitation, all partners have stated their commitments towards the sustainability of the project results and they have elaborated an individual exploitation plan.

In the following we list the exploitation plans from industrial and use case partners:

- **FLEX** will utilise the results using knowledge transfer as a valuable input to our Concerto commercial product via our product steering group and implementation roadmap. This will aid Flexiant as an innovative SME to become the vanguard of a European charge towards market penetration in the lucrative application and data services marketplace.
- In order to ensure the sustainability and the exploitation of the PaaSage results, **ASCS** will transfer the results through its network to the European industry using Computer Aided Engineering (CAE) methods. The main goal is to force the implementation of the flexible and adaptive features of PaaSage for new cost effective CAE cloud environments.
- **BEWAN** will use the PaaSage platform and the related technologies (CAMEL, Social Network) in order to support the deployment and in order to operate be.wan's (multi)cloud based applications. Be.wan also commits to develop a PaaSage consultancy activity to

partners & customers that operate “value added” be.wan cloud applications or their own applications.

- **EVRY** will transfer knowledge into our work with cloud-enabled modernization of public sector. EVRY will also promote the platform towards our large customer base as a part of our story on innovative cloud and application management services.
- **IBSCY** will exploit the PaaSage platform to support different business requirements (scenarios) of existing and new clients in order to enrich our cloud installations clientele related with Quorum (financial institutions). Also, IBSCY Ltd is committed to promote the PaaSage platform within our group of companies in order to provide the opportunity to model other applications that are sold in other countries of the group.
- **LSY** states the optimizing the utilized cloud infrastructures is a common issue but not a critical one concerning competition. Therefore, tackling this issue as a community and implementing an open source solution is the best approach from Lufthansa Systems point of view. PaaSage’s open source software is the most promising approach moving into the desired direction.
- At **AGH**, we commit ourselves to continue further research activities regarding the development of HyperFlow and Scalarm solutions to support complex scientific applications on clouds and other distributed infrastructures. HyperFlow and Scalarm are research products and we plan to continue not only research in this area, but also further development of these products so that they are usable outside the project. This will include maintenance, support for our users on a best effort basis and exploitation of the project results in the future research projects including EU-funded H2020 and RFCS programmes. It is worth mentioning that both HyperFlow and Scalarm are already installed in the software stack at the Academic Computer Center AGH and are used as services by the registered users from the Krakow scientific community.
- In order to ensure the sustainability and the exploitation of the PaaSage results, **USTUTT** will integrate the PaaSage platform within its planned high-performance Cloud services that will be accessible for industrial as well as academic users. Furthermore, USTUTT will distribute the benefits and findings of Cloud application modelling through industrial trainings and academic workshops.

Here are additional specific exploitation plans by non-industrial partners:

- **UULM** will continue the development of the Cloudiator component (Executionware) including the realisation of new features and the provisioning of bug fixes. We will further continue support for the tool including technical assistance for members joining the community and if needed provide the infrastructure to run continuous integration processes.
- **CETIC** will maintain the source code of the PaaSage platform that has been submitted on OW2. CETIC will also manage the issues related to this source code by dispatching the requests for action to the corresponding partners
- In order to ensure the sustainability and the exploitation of the PaaSage results, **FORTH** is committed to pursue opportunities to exploit the PaaSage social network platform and

knowledge mining on top of the PaaSage metadata database, a key component of the social network.

- **SINTEF** will maintain the Cloud Application Modelling and Execution Language (CAMEL) along with its Textual Editor, and push the results from CAMEL into the Topology and Orchestration Specification for Cloud Applications (TOSCA) standard developed by the Organization for the Advancement of Structured Information Standards (OASIS).
- At **INRIA**, we commit ourselves to continuing the development of the CP Generator Model-to-Solver, BPPH solver and Adapter components, providing software maintenance and technical support also for Solver-to-deployment component. Moreover, Inria will disseminate PaaSage results through academic publications and teaching courses in local universities.
- At **ERCIM**, We commit ourselves to disseminate the project results via our media (web site, ERCIM News) towards our member institutions as well as towards a large academic and industrial research community. We also commit ourselves to support further collaborative research initiatives that may further develop, improve and extend the PaaSage platform.
- **UIO** will continue to develop, expand and improve the open source results: The LA-based Solver component and the corresponding Learning Automata (LA) template library. Maintenance and support will be offered to the users of these results. We will also continue to push for the reuse of the PaaSage platform within other research projects, and thereby the maintenance of the platform should the project be funded
- **UCY** will integrate the PaaSage platform within the cloud computing infrastructure that was setup at UCY during the PaaSage project, and make it accessible for industrial as well as academic users in Cyprus. Moreover, we will drive knowledge transfer to local researchers (mainly on Cloud Computing and Model-Driven Development) and pursue future research collaborations based on PaaSage results, as well as the local industry on how to utilise the (open-source) PaaSage platform or its components.
- We commit ourselves to the use of PaaSage within our internal and international Research networks and infrastructures. PaaSage provides **STFC** with the ability to better capture the computation needs of end users and help find appropriate resource. This includes specialist private through to public clouds.
- In order to ensure the sustainability and the exploitation of the PaaSage results, **GWDG** will promote and transfer knowledge about PaaSage platform to its clients and affiliates, so the lifecycle management of applications deployed on GWDG cloud can be automated and self-controlled by users. GWDG will also raise awareness about PaaSage in its other projects.

This shows that PaaSage has a bright sustainability plan.

3.5. Section A – Tables

TABLE A1: LIST OF SCIENTIFIC (PEER REVIEWED) PUBLICATIONS, STARTING WITH THE MOST IMPORTANT ONES

NO.	Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers[16] (if available)	Is/Will open access[17] provided to this publication?
1	Cardinality-based feature models with constraints: a pragmatic approach	Clément Quinton, Daniel Romero, and Laurence Duchien	SPLC 2013	ISBN: 978-1-4503-1968-3	ACM New York, NY, USA ©2013		2013	162-166		http://dx.doi.org/10.1145/2491627.2491638
2	A vision for a stochastic reasoner for autonomic cloud deployment	Geir Horn	NordiCloud 2013	ISBN: 978-1-4503-2307-9	ACM New York, NY, USA ©2013		2013	46-53		http://dx.doi.org/10.1145/2513534.2513543
3	A vision for better cloud applications	Keith Jeffery, Geir Horn, and Lutz Schubert	MultiCloud 2013	ISBN: 978-1-4503-2050-4	ACM New York, NY, USA ©2013		2013	7-12		http://dl.acm.org/citation.cfm?id=2462329
4	An Architecture for Evaluating Distributed Application Deployments	Antonis Papaioannou and Kostas Magoutis	CloudCom 2013	ISBN: 978-0-7695-	IEEE Computer Society		2013	547-554		http://dl.acm.org/citation.cfm?id=2568575

[16] A permanent identifier should be a persistent link to the published version full text if open access or abstract if article is pay per view) or to the final manuscript accepted for publication (link to article in repository).

[17] Open Access is defined as free of charge access for anyone via Internet. Please answer "yes" if the open access to the publication is already established and also if the embargo period for open access is not yet over but you intend to establish open access afterwards.

	in Multi-Clouds			5095-4	Washington, DC, USA ©2013					
5	Beyond IaaS and PaaS: An Extended Cloud Taxonomy for Computation, Storage and Networking	Steffen Kächele, Christian Spann , Franz J. Hauck und Jörg Domaschka	UCC 2013	ISBN: 978-0-7695-5152-4	IEEE Computer Society Washington, DC, USA ©2013		2013	75-82		http://dl.acm.org/citation.cfm?id=2588676
6	Data Farming as a Service	Dariusz Krol, Michal Orzechowski, Renata Slota, and Jacek Kitowski	ICCS 2014				2014			
7	Handling Constraints in Cardinality-Based Feature Models: The Cloud Environment Case Study	Clément Quinton, Daniel Romero, and Laurence Duchien	ICSE 2014	ISBN: 978-1-4503-1968-3	ACM New York, NY, USA ©2013		2014			http://dl.acm.org/citation.cfm?id=2491638
8	Lifecycle Management of Service-based Applications on Multi-Clouds: A Research Roadmap	George Baryannis, Panagiotis Garefalakis, Kyriakos Kritikos, Kostas Magoutis, Antonis Papaioannou, Dimitris Plexousakis and Chrysostomos Zeginis	MultiCloud 2013	ISBN: 978-1-4503-2050-4	ACM New York, NY, USA ©2013		2013	13-20		http://dx.doi.org/10.1145/2462326.2462331
9	Managing elasticity across multiple cloud providers	Fawaz Paraiso, Philippe Merle, and Lionel Seinturier	MultiCloud 2013	ISBN: 978-1-4503-2050-4	ACM New York, NY, USA		2013	53-60		http://dx.doi.org/10.1145/2462326.2462338

					©2013					
10	Managing multi-cloud systems with CloudMF	Nicolas Ferry, Franck Chauvel, Alessandro Rossini, Brice Morin and Arnor Solberg	NordiCloud 2013	ISBN: 978-1-4503-2307-9	ACM New York, NY, USA ©2013		2013	38-45		http://dx.doi.org/10.1145/2513534.2513542
11	Managing Service Performance in the Cassandra Distributed Storage System	Maria Chalkiadaki and Kostas Magoutis	CloudCom 2013				2013			http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6753779
12	Model Based Cloud Application Development using PaaSage	Bastian Koller	inSiDE				2013			
13	Models@Runtime to support the Iterative and Continuous Design of Autonomous Reasoners	Franck Chauvel, Nicolas Ferry, Brice Morin, Alessandro Rossini and Arnor Solberg	MRT 2013				2013			
14	soCloud: A service-oriented component-based PaaS for managing portability, provisioning, elasticity, and high availability across multiple clouds	Fawaz Paraiso, Philippe Merle, and Lionel Seinturier	SCJ				2013			https://hal.inria.fr/hal-01019385/file/paper.pdf
15	Strengthening Consistency in the Cassandra Distributed Key-Value Store	Panagiotis Garefalakis, Panagiotis Papadopoulos, Ioannis Manousakis and Kostas Magoutis	DAIS 2013				2013			http://link.springer.com/chapter/10.1007/978-3-642-38541-4_17

16	Towards Cross-layer Monitoring of Multi-Cloud Service-based Applications	Chrysostomos Zeginis, Kyriakos Kritikos, Panagiotis Garefalakis, Konstantina Konsolaki, Kostas Magoutis and Dimitris Plexousakis	ESOCC'13				2013			https://pdfs.semanticscholar.org/44a8/23ad99d6de6936f29363ab28944de8f1163b.pdf
17	Towards model-driven provisioning, deployment, monitoring, and adaptation of multi-cloud systems	Nicolas Ferry, Alessandro Rossini, Franck Chauvel, Brice Morin and Arnor Solberg	CLOUD 2013	Electronic ISSN: 2159-6190	IEEE		2013			http://dx.doi.org/10.1109/CLOUD.2013.133
18	Towards multi-cloud configurations using feature models and ontologies	Clément Quinton, Nicolas Haderer, Romain Rouvoy, and Laurence Duchien	MultiCloud 2013	ISBN: 978-1-4503-2050-4	ACM New York, NY, USA ©2013		2013			http://dx.doi.org/10.1145/2462326.2462332
19	WSSL: A Fluent Calculus-Based Language for Web Service Specifications	George Baryannis and Dimitris Plexousakis	CAiSE 2013	978-3-642-38708-1	Springer Berlin Heidelberg		2013	256-271		http://dx.doi.org/10.1007/978-3-642-38709-8_17
20	Optimal Negotiation of Service Level Agreements for Cloud-based Services using Autonomous Agents	Edwin Yaqub, Ramin Yahyapour, Philipp Wieder, Constantinos Kotsokalis, Kuan Lu, Ali Imran Jehangiri	SCC 2014	ISBN: 978-1-4799-5066-9	IEEE Computer Society Washington, DC, USA ©2014		2014	59-66		http://dl.acm.org/citation.cfm?id=2680847.2681496
21	Towards Bridging the Gap Between Scalability and Elasticity	Nicolas Ferry, Gunnar Brataas, Alessandro Rossini, Franck Chauvel and Arnor	CLOSER 2014	ISBN: 978-1-4503-2307-9	ACM New York, NY, USA		2014	38-45		http://dx.doi.org/10.1145/2513534.2513542

		Solberg			©2013					
22	Assessing and monitoring cloud security: from controls frameworks to Service Level Agreements	Jesus Luna, Ruben Trapero, Alain Pannetrat, Philippe Massonet	SECRYPT 2014				2014			
23	Metaheuristics-based Planning and Optimization for SLA-aware Resource Management in PaaS Clouds	Edwin Yaqub, Ramin Yahyapour, Philipp Wieder, Ali Imran Jehangiri, Kuan Lu, Constantinos Kotsokalis	UCC 2014	ISBN: 978-1-4799-7881-6	IEEE Computer Society Washington, DC, USA ©2014		2014	288-297		http://dl.acm.org/citation.cfm?id=2760050
24	A Strategic Framework for Automating SLA Negotiation in Emerging Service Oriented Markets	Edwin Yaqub, Ramin Yahyapour, Philipp Wieder, Constantinos Kotsokalis, Kuan Lu, Ali Imran Jehangiri	FGCS				2014			
25	QoS-Aware VM Placement in Multi-Domain Service Level Agreements Scenarios	Kuan Lu, Ramin Yahyapour, Philipp Wieder, Constantinos Kotsokalis, Edwin Yaqub, Ali Imran Jehangiri	CLOUD 2013	Electronic ISSN: 2159-6190	IEEE		2013			http://dx.doi.org/10.1109/CLOUD.2013.112
26	QoS-based Resource Allocation Framework for Multi-Domain SLA Management in Clouds	Kuan Lu, Ramin Yahyapour, Philipp Wieder, Constantinos Kotsokalis, Edwin Yaqub, Ali Imran Jehangiri	IJCC				2013			
27	ACaZoo: A Distributed Key-Value Store based on	Panagiotis Garefalakis,	SRDS 2014	Electronic ISBN: 978-	IEEE		2014			DOI: 10.1109/SRDS.2014.4

	Replicated LSM-Trees	Panagiotis Papadopoulos and Kostas Magoutis		1-4799-5584-8						3
28	Scalable entity-based summarization of web search results using MapReduce	Ioannis Kitsos, Kostas Magoutis, Yannis Tzitzikas	CLOUD 2014				2014			http://dx.doi.org/10.1007/s10619-013-7133-7
29	Consistency Checking for the Evolution of Cardinality-based Feature Models	Clément Quinton, Andreas Pleuss, Daniel Le Berre, Laurence Duchien and Goetz Botterweck	SPLC 2014	ISBN: 978-1-4503-2740-4	ACM New York, NY, USA ©2014		2014	122-131		http://dx.doi.org/10.1145/2648511.2648524
30	SALOON: A Platform for Selecting and Configuring Cloud Environments	Clément Quinton, Daniel Romero, Laurence Duchien	SP&E				2014			https://hal.inria.fr/hal-01103560/document
31	Towards a Generic Language for Scalability Rules	Jörg Domaschka, Kyriakos Kritikos, Alessandro Rossini	CSB 2014	Online ISBN: 978-3-319-14886-1	Springer International Publishing		2014	206-220		DOI: 10.1007/978-3-319-14886-1_19
32	A Cloud-based Data Farming Platform for Molecular Dynamics Simulations	Dariusz Krol, Michal Orzechowski, Jacek Kitowski Christoph Niethammer, Anthony Sulistio, Amer Wafai	CloudAM 2014	Electronic ISBN: 978-1-4799-7881-6	IEEE		2014			DOI: 10.1109/UCC.2014.89
33	'CLOUDS a large Virtualisation of Small Things' Keynote; FICLOUD2014 Barcelona August 27-29 2014	Keith Jeffery	FICloud2014				2014			

34	Towards a Model-Based Execution-Ware for Deploying Multi-Cloud Applications	Daniel Baur, Stefan Wesner, Jörg Domaschka	CSB 2014	Online ISBN 978-3-319-14886-1	Springer International Publishing		2014	124-138		DOI: 10.1007/978-3-319-14886-1_13
35	Model-Based Execution of Scientific Applications on Cloud Infrastructures Scalarm Case Study	Dariusz Krol, Michal Orzechowski, Jakub Liput, Renata Slota, and Jacek Kitowski	CGW14				2014			
36	Optimising Multi-Cloud Deployments with Security Controls as Constraints	Philippe Massonet, Jesus Luna, Alain Pannetrat, Ruben Trapero	ESSOS 2015				2015			
37	soCloud: A service-oriented component-based PaaS for managing portability, provisioning, elasticity, and high availability across multiple clouds	Fawaz Paraiso, Philippe Merle, Lionel Seinturier	Computing				2014			http://dx.doi.org/10.1007/s00607-014-0421-x
38	CloudMF: Applying MDE to Tame the Complexity of Managing Multi-Cloud Applications	Nicolas Ferry, Hui Song, Alessandro Rossini, Franck Chauvel, Arnor Solberg	UCC 2014	Electronic ISBN: 978-1-4799-7881-6	IEEE		2014			DOI: 10.1109/UCC.2014.36
39	SRL: A Scalability Rule Language for Multi-Cloud Environments	Kyriakos Kritikos, Jörg Domaschka, Alessandro Rossini	CloudCom 2014	Electronic ISBN: 978-1-4799-4093-6	IEEE		2014			DOI: 10.1109/CloudCom.2014.170
40	Event Pattern Discovery for Cross-Layer	Chrysostomos Zeginis, Kyriakos	ESOCC'14	Online ISBN: 978-	Springer Berlin		2014	138-147		http://dx.doi.org/10.1007/978-3-662-44879-

	Adaptation of Multi-Cloud Applications	Kritikos, and Dimitris Plexousakis		3-662-44879-3	Heidelberg					3_10
41	Executing Model-Based Applications	Lutz Schubert, Jörg Domaschka, Daniel Baur	CGW14				2014			
42	Design and Implementation of a Social Network for Model-driven Software Engineering	Christos Papoulas, Antonis Papaioannou and Kostas Magoutis	JISA				2015			DOI: 10.1186/s13174-015-0033-5
43	Cross-layer Management of Distributed Applications on Multi-Clouds	Antonis Papaioannou, Damianos Metalidis, and Kostas Magoutis	M'15	IEEE	Electronic ISBN: 978-1-4799-8241-7		2015			DOI: 10.1109/INM.2015.7140336
44	Vision on CLOUD and the IMpact on European R&D	Keith Jeffery, Lutz Schubery					2014			
45	GLUE Schema Unification using Hypergraph Approach	Michał Cieszyński, Szymon Łabuz, Bartosz Kryza, Renata Słota, Jacek Kitowski	CGW14				2014			
46	Executing Model-Based Applications	Lutz Schubert, Jörg Domaschka, Daniel Baur	CGW14				2014			
47	Towards Deployment and Autoscaling of Scientific Workflows with HyperFlow and PaaSage	Bartosz Balis, Marian Bubak, Kamil Figiela, Maciej Malawski, and Maciej Pawlik	CGW 2014				2014			http://ieeexplore.ieee.org/abstract/document/7431445/definitions?ctx=definitions
48	New Software Engineering Requirements in Clouds and Large-Scale Systems	Lutz Schubert, Keith Jeffrey	IEEECC	Online ISSN: 2325-	IEEE		2015			DOI: 10.1109/MCC.2015.13

				6095						
49	A Cross-Cloud, Multi-Tenant Deployment and Runtime Engine	Daniel Baur, Daniel Seybold, Frank Griesinger	SummerSOC 2015				2015			
50	Axe: A Novel Approach for Generic, Flexible, and Comprehensive Monitoring and Adaptation of Cross-Cloud Applications	Daniel Baur, Daniel Seybold, Frank Griesinger	SAMM15	Online ISBN: 978-3-319-33313-7			2015	184-196		DOI: 10.1007/978-3-319-33313-7_14
51	Cloud Orchestration Features -- Are Tools Fit for Purpose?	Daniel Baur, Daniel Seybold, Frank Griesinger, Jörg Domaschka	UCC 2015				2015			
52	A Comparison of Two-Level and Multi-level Modelling for Cloud-Based Applications	Alessandro Rossini, Juan de Lara, Esther Guerra, Nikolay Nikolov	ECMFA 2015	Online ISBN: 978-3-319-21151-0	Springer International Publishing		2015	18-32		http://dx.doi.org/10.1007/978-3-319-21151-0_2
53	Distributed Predictive Performance Anomaly Detection for Virtualized Clouds	Ali Jehangiri, Ramin Yahyapour, Philipp Wieder, Edwin Yaqub	IJHPCN				2015			
54	Distributed Predictive Performance Anomaly Detection for Virtualized Clouds	Ali Jehangiri, Ramin Yahyapour, Philipp Wieder, Edwin Yaqub	BDCloud 2015				2015			
55	Security-Based Adaptation of Multi-Cloud Applications	Kyriakos Kritikos, philippe massonet	QASA 2015	Online ISBN 978-3-319-29883-2	Springer International Publishin		2015	47-64		DOI: 10.1007/978-3-319-29883-2_4

					g					
56	Multi-cloud Application Design through Cloud Service Composition	Kyriakos Kritikos, Dimitris Plexousakis	CLOUD 2015				2015			http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7214106
57	A PaaSage to Multi-Site Security for Clouds	Tom Kirkham, Kyriakos Kritikos, Bartosz Kryza, Philippe Massonet, Franky Vanraes	SummerSOC 2015				2015			
58	Cloud Application Modelling and Execution Language (CAMEL) and the PaaSage Workflow	Alessandro Rossini	ESOCC 2015				2015			
59	Beyond Mere Application Structure: Thoughts on the Future of Cloud Orchestration Tools	Jörg Domaschka, Frank Griesinger, Daniel Baur, Alessandro Rossini	CF 2015				2015			
60	Integration of DSLs and Migration of Models: a case study in the cloud computing domain	Nikolay Nikolov, Alessandro Rossini, Kyriakos Kritikos	CF 2015				2015			
61	Security Enforcement for Multi-Cloud Platforms - The Case of PaaSage	Tom Kirkham, Kyriakos Kritikos, Bartosz Kryza, Philippe Massonet	CF 2015				2015			
62	Privacy Aware On-Demand Resource Provisioning for IoT Data Processing	T. Kirkham, A. Sinha, N. Parlavantzas, B. Kryza and P. Fremantle	IoTaaS 2015				2015			https://hal.inria.fr/hal-01198381/file/IoTdata processing-submitted.pdf
63	The Evolution of CloudML and its Manifestations	Alexander Bergmayr, Alessandro Rossini, Nicolas Ferry, Geir	CloudMDE 2015				2015			http://ceur-ws.org/Vol-

		Horn, Leire Orue-Echevarria, Arnor Solberg, Manuel Wimmer								1563/paper3.pdf
64	SmartyCo: Managing Cyber-Physical Systems for Smart Environments	Daniel Romero, Clément Quinton, Laurence Duchien, Lionel Seinturier, Carolina Valdez	ECSA 2015	Online ISBN: 978-3-319-23727-5			2015	294-302		DOI: 10.1007/978-3-319-23727-5_25
65	Adaptive Multi-level Workflow Scheduling with Uncertain Task Estimates	Tomasz Dziok, Kamil Figiela, Maciej Malawski	PPAM 2015	Online ISBN: 978-3-319-32152-3	Springer International Publishing		2015	90-100		DOI: 10.1007/978-3-319-32152-3_9
66	A lightweight approach for deployment of scientific workflows in cloud infrastructures	Bartosz Balis, Kamil Figiela, Maciej Malawski, Maciej Pawlik, and Marian Bubak	PPAM 2015	Online ISBN: 978-3-319-32149-3	Springer International Publishing		2015	281-290		
67	Support for Scientific Workflows in a Model-based Cloud Platform	Maciej Malawski, Bartosz Balis, Kamil Figiela, Maciej Pawlik, Marian Bubak	UCC 2015	Electronic ISBN: 978-0-7695-5697-0	IEEE		2015			DOI: 10.1109/UCC.2015.70
68	Business-Oriented Evaluation of the PaaSage Platform	Achilleas P. Achilleos, Georgia M. Kapitsaki, Eleni Constantinou, Geir Horn, George A. Papadopoulos	UCC 2015				2015			
69	Challenges emerging from future cloud application scenarios	Keith Jeffery, George Kousiouris, Dimosthenis Kyriazis, Jörn Altmann, Augusto Ciuffoletti ,	CF 2015				2015			

		Ilias Maglogiannis, Paolo Nesi , Bojan Suzic, Zhiming Zhao								
70	Parameter studies on heterogeneous computing infrastructures with the Scalarm platform	D. Król, R. Słota, J. Kitowski	HPCS 2016		IEEE		2016			http://ieeexplore.ieee.org/document/7568310/
71	Workflow Performance Profiles: Development and Analysis	D. Król, R. Ferreira da Silva, E. Deelman, V. Lynch	EUROPAR 2016				2016			https://scitech.isi.edu/wordpress/wp-content/papercite-data/pdf/krol-heteropar-2016.pdf
72	Science Automation in Practice - Performance Data Farming in Workflows	D. Król, E. Deelman, R. F. Da Silva, G. Juve, J. Kitowski, M. Rynge and K, Vahi	ETFA 2016							https://scitech.isi.edu/wordpress/wp-content/papercite-data/pdf/krol-etfa-2016.pdf
73	Efficient Heuristics for Placing Large-Scale Distributed Applications on Multiple Clouds	Pedro Silva, Christian Pérez, Frédéric Desprez	CCGRID 2016		IEEE		2016			DOI: 10.1109/CCGrid.2016.77
74	Efficient Communication Aware Heuristics for Placing Large-Scale Distributed Applications on Multiple Clouds	Pedro Silva, Christian Pérez, Frédéric Desprez	IPDPS 2017				2016			
75	Towards Knowledge-Based Assisted IaaS Selection	Kyriakos Kritikos, Kostas Magoutis, Dimitris Plexousakis	CloudCom 2016				2017			http://hipore.com/stcc/2015/IJCC-Vol3-No3-2015c.pdf
76	An Integrated Meta-Model for Cloud Application Security Modelling	Kyriakos Kritikos, Philippe Massonet	CF 2016				2016			https://www.researchgate.net/publication/309196984_An_Integrated_Meta-

										model_for_Cloud_Application_Security_Modelling
	Towards Semantic-Based Cloud Application Management	Kyriakos Kritikos, Dimitris Plexousakis	STCC				2015			http://hipore.com/stcc/2015/IJCC-Vol3-No3-2015c.pdf
77	Experiences of Models@run-time with EMF and CDO	Daniel Seybold, Jörg Domaschka, Alessandro Rossini, Christopher B. Hauser, Frank Griesinger, Athanasios Tsitsipas	SLE 2016	ISBN: 978-1-4503-4447-0	ACM New York, NY, USA ©2016		2016	46-56		DOI: 10.1145/2997364.2997380
78	Experiences from Building a Cross-cloud Orchestration Tool	Daniel Baur, Jörg Domaschka	CrossCloud '16	ISBN: 978-1-4503-4294-0	ACM New York, NY, USA ©2016		2016			Doi: 10.1145/2904111.2904116
79	Cloudiator – Enacting deployments and adaptation in PaaSage	Daniel Baur, Jörg Domaschka	ESOCC 2016				2016			
80	Towards a Security-Enhanced PaaS Platform for Multi-Cloud Applications	Kyriakos Kritikos, Tom Kirkham, Bartosz Kryza, Philippe Massonet	FGCS				2016			http://hipore.com/stcc/2015/IJCC-Vol3-No3-2015c.pdf
81	Molecular Dynamics with HyperFlow and Scalarm on the PaaSage Platform	Maciej Malawski, Bartosz Balis, Kamil Figiela, Maciej Pawlik, Marian Bubak, Dariusz Król, Renata Słota, Michal	ESOCC 2016				2016			
82	Storage-aware Algorithms for Scheduling of Workflow Ensembles in Clouds	Piotr Bryk, Maciej Malawski, Ewa Deelman, Gideon Juve	JGC				2016			DOI: 10.1007/s10723-015-9355-6

83	Science Automation in Practice - Performance Data Farming in Workflows	D. Król, E. Deelman, R. F. Da Silva, G. Juve, J. Kitowski, M. Rynge and K, Vahi					2016			https://scitech.isi.edu/wordpress/wp-content/papercite-data/pdf/krol-etfa-2016.pdf
84	SALOON: a platform for selecting and configuring cloud environments	Clément Quinton, Daniel Romero, Laurence Duchien SP&E					2016			https://hal.inria.fr/hal-01103560/en
85	In Search for a Scalable and Reactive Architecture of a Cloud Application: CQRS and Event Sourcing Case Study	Andrzej Dębski, Bartłomiej Szczepanik, Maciej Malawski, Stefan Spahr, Dirk Muthig	Software				2016			http://www.icsr.agh.edu.pl/~malawski/DebskiSzczepanik-CQRS-IEEE-Software.pdf

TABLE A2: LIST OF DISSEMINATION ACTIVITIES						
NO.	Type of activities[18]	Title	Date/Period	Place	Type of audience[19]	Countries addressed
1	Conference	7th Middleware for Next-Generation Internet Computing Workshop at the 13th International Middleware Conference (Middelware 2012)	3-7 December 2012	Montreal, Canada	Scientific Community, Industry,	Canada
2	Conference	International Workshop on Multi-Cloud Applications and Federated Clouds	21-24 April 2013	Prague, Czech Republic	Scientific Community, Industry,	Czech Republic
3	Conference	13th IFIP International Conference on Distributed Applications and Interoperable Systems	3-6 June 2013	Firenze, Italia	Scientific Community, Industry,	Italia
4	Conference	6th International Conference on Cloud Computing	27 Jun - 02 July 2013	Santa Clara Marriott	Scientific Community, Industry,	USA
5	Conference	25th International Conference on Advanced Information Systems Engineering	17-21 June 2013	Valencia, Spain	Scientific Community, Industry,	Spain
6	Conference	17th International Software Product Line Conference	26-30 August 2013	Tokyo, Japan	Scientific Community, Industry,	Japan
7	Conference	2nd Nordic Symposium on Cloud Computing and Internet Technologies	2-3 September 2013	Oslo, Norway	Scientific Community,	Norway

[18] A drop down list allows choosing the dissemination activity: publications, conferences, workshops, web, press releases, flyers, articles published in the popular press, videos, media briefings, presentations, exhibitions, thesis, interviews, films, TV clips, posters, Other.

[19] A drop down list allows choosing the type of public: Scientific Community (higher education, Research), Industry, Civil Society, Policy makers, Medias, Other ('multiple choices' is possible).

					Industry,	
8	Conference	2nd European Conference on Cloud and Service Computing	11-13 September 2013	Malaga, Spain	Scientific Community, Industry,	Spain
9	Conference	5th IEEE International Conference on Cloud Computing Technology and Science	2-5 December 2013	Bristol, UK	Scientific Community, Industry,	UK
10	Conference	European Conference on Cloud and Service Computing	11-13 September 2013	Malaga, Spain	Scientific Community, Industry,	Spain
11	Conference	36th International Conference on Software Engineering	31 May - 7 Jun 2014	Hyderabad, India	Scientific Community, Industry,	India
12	Conference	9th International Symposium on Software Engineering for Adaptive and Self-Managing Systems @ ICSE 2014	31 May - 7 Jun 2014	Hyderabad, India	Scientific Community, Industry,	India
13	Conference	The International Conference on Computational Science	10-12 June 2014	Cairns, Australia	Scientific Community, Industry,	Australia
14	Conference	6th ACM/IEEE International Conference on Utility and Cloud Computing	9-12 December 2013	Dresden, Germany	Scientific Community, Industry,	Germany
15	Conference	4th International Conference on Cloud Computing and Services Science	3-5 April 2014	Barcelona, Spain	Scientific Community, Industry,	Spain
16	Conference	11th IEEE International Conference on Services Computing	27 June - 2 July 2014	Alaska, USA	Scientific Community, Industry,	USA
17	Conference	11th International Conference on Security and Cryptography	28-30 August 2014	Vienna, Austria	Scientific Community, Industry,	Austria
18	Conference	7th International Conference on Cloud Computing	27 June - 2 July 2014	Alaska, USA	Scientific Community, Industry,	USA

19	Conference	7th IEEE/ACM International Conference on Utility and Cloud Computing	8-11 December 2014	London, UK	Scientific Community, Industry,	UK
20	Conference	33rd IEEE Symposium on Reliable Distributed Systems	6-9 October 2014	Nara, Japan	Scientific Community, Industry,	Japan
21	Conference	IEEE International Conference on Internet of Things (iThings 2012)	20-23 November 2012	Besancon, France	Scientific Community, Industry,	France
22	Conference	18th International Software Product Line Conference	15-19 September 2014	Florence, Italy	Scientific Community, Industry,	Italy
23	Conference	2nd International Workshop on Cloud Service Brokerage	2 September 2014	Manchester, UK	Scientific Community, Industry,	UK
24	Conference	The 2nd International Conference on Future Internet of Things and Cloud (FiCloud-2014)	27-29 August 2014	Barcelona	Scientific Community, Industry,	Spain
25	Conference	International Workshop on Clouds and (eScience) Applications Management (CloudAM 2014)	8-11 December 2014	London, UK	Scientific Community, Industry,	UK
26	Conference	Cracow Grid Workshops	27-29 October 2014	Kraków, Poland	Scientific Community, Industry,	Poland
27	Conference	International Symposium on Engineering Secure Software and Systems	4-6 March 2014	Milan, Italy	Scientific Community, Industry,	Italy
28	Conference	European Conference on Cloud and Service Computing	2-4 September 2014	Manchester, UK	Scientific Community, Industry,	UK
29	Conference	IFIP/IEEE International Conference on Integrated Network Management	11-15 May 2015	Ottawa, Canada	Scientific Community, Industry,	Canada
30	Conference	CLOUD Assisted Services in Europe	25-26 September 2014	Bled, Slovenia	Scientific Community,	Slovenia

					Industry,	
31	Conference	9th Symposium and Summer School On Service-Oriented Computing	28 June - 3 July 2015	Heraklion, Crete, Greece	Scientific Community, Industry,	Greece
32	Conference	Cloud Forward	6-8 October 2015	Pisa, Italy	Scientific Community, Industry,	Italy
33	Conference	Seamless Adaptive Multi-cloud Management of Service-based Applications	15-17 September 2015	Taormina, Italy	Scientific Community, Industry,	Italy
34	Conference	8th IEEE/ACM International Conference on Utility and Cloud Computing	7-10 December 2015	Limassol, Cyprus	Scientific Community, Industry,	Cyprus
35	Conference	11th European Conference on Modelling Foundations and Applications	20-24 July 2015	L'Aquila, Italy	Scientific Community, Industry,	Italy
36	Conference	5th IEEE International Conference on Big Data and Cloud Computing	26-29 August 2015	Dalian, China	Scientific Community, Industry,	China
37	Conference	4th International Workshop on Quantitative Aspects of Security Assurance	21-25 September 2015	Vienna, Austria	Scientific Community, Industry,	Austria
38	Conference	8th International Conference on Cloud Computing	27 June - 2 July 2015	New York, USA	Scientific Community, Industry,	USA
39	Conference	2nd EAI International Conference on IoT as a Service	26-27 October 2015	Rome, Italy	Scientific Community, Industry,	Italy
40	Conference	4th European Conference on Service-Oriented and Cloud Computing	15-17 September 2015	Messina, Italy	Scientific Community, Industry,	Italy
41	Conference	3rd International Workshop on Model-Driven Engineering on and for the Cloud	29/09/2015	Ottawa, Canada	Scientific Community, Industry,	Canada

42	Conference	9th European Conference on Software Architecture	7-11 September 2015	Dubrovnik/Cavtat, Croatia	Scientific Community, Industry,	Croatia
43	Conference	11th International Conference on Parallel Processing and Applied Mathematics	6-9 September 2015	Kraków, Poland	Scientific Community, Industry,	Poland
44	Conference	Cloud Forward	6-8 October 2015	Pisa, Italy	Scientific Community, Industry,	Italy
45	Conference	Seamless Adaptive Multi-cloud Management of Service-based Applications	15-17 September 2015	Taormina, Italy	Scientific Community, Industry,	Italy
46	Conference	8th IEEE/ACM International Conference on Utility and Cloud Computing	7-10 December 2015	Limassol, Cyprus	Scientific Community, Industry,	Cyprus
47	Conference	11th European Conference on Modelling Foundations and Applications	20-24 July 2015	L'Aquila, Italy	Scientific Community, Industry,	Italy
48	Conference	5th IEEE International Conference on Big Data and Cloud Computing	26-29 August 2015	Dalian, China	Scientific Community, Industry,	China
49	Conference	4th International Workshop on Quantitative Aspects of Security Assurance	21-25 September 2015	Vienna, Austria	Scientific Community, Industry,	Austria
50	Conference	8th International Conference on Cloud Computing	27 June - 2 July 2015	New York, USA	Scientific Community, Industry,	USA
51	Conference	2nd EAI International Conference on IoT as a Service	26-27 October 2015	Rome, Italy	Scientific Community, Industry,	Italy
52	Conference	4th European Conference on Service-Oriented and Cloud Computing	15-17 September 2015	Messina, Italy	Scientific Community, Industry,	Italy
53	Conference	3rd International Workshop on Model-Driven	29 September 2015	Ottawa, Canada	Scientific Community,	Canada

		Engineering on and for the Cloud			Industry,	
54	Conference	9th European Conference on Software Architecture	7-11 September 2015	Dubrovnik/Cavtat, Croatia	Scientific Community, Industry,	Croatia
55	Conference	11th International Conference on Parallel Processing and Applied Mathematics	6-9 September 2015	Kraków, Poland	Scientific Community, Industry,	Poland
56	Conference	2016 International Conference on High Performance Computing & Simulation	18 – 22 July 2016	Innsbruck, Austria	Scientific Community, Industry,	Austria
57	Conference	the Fourteenth International Workshop Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms	23 August 2016	Grenoble, France	Scientific Community, Industry,	France
58	Conference	21st IEEE International Conference on Emerging Technologies and Factory Automation	6 – 9 September 2016	Berlin, Germany	Scientific Community, Industry,	Germany
59	Conference	16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing	16 – 19 May 2016	Cartagena, Colombia	Scientific Community, Industry,	Colombia
60	Conference	21st IEEE International Conference on Emerging Technologies and Factory Automation	May 29 – June 2, 2017	Orlando, Florida USA	Scientific Community, Industry,	USA
61	Conference	16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing	12 -15 December 2015	Luxembourg, Luxembourg	Scientific Community, Industry,	Luxembourg
62	Conference	31st IEEE International Parallel & Distributed Processing Symposium	October 18 - October 20	Madrid, Spain	Scientific Community, Industry,	Spain
63	Conference	8th IEEE International Conference on Cloud Computing Technology and Science	30 October-1 November 2016	Amsterdam, The Netherlands	Scientific Community, Industry,	The Netherlands

64	Conference	Cloud Forward Conference 2016: From Distributed to Complete Computing	18 April 2016	London, UK	Scientific Community, Industry,	UK
65	Conference	9th ACM SIGPLAN International Conference on Software Language Engineering (SLE)	5-7 September 2014	Vienna, Austria	Scientific Community, Industry,	Austria
66	Conference	3rd Workshop on CrossCloud Infrastructures & Platforms	15-17 September 2015	Taormina, Italy	Scientific Community, Industry,	Italy
67	Conference	5th European Conference on Service-Oriented and Cloud Computing	7-10 December 2015	Limassol, Cyprus	Scientific Community, Industry,	Cyprus

3.6. Section B – Tables

TABLE B1: LIST OF APPLICATIONS FOR PATENTS, TRADEMARKS, REGISTERED DESIGNS, ETC.					
Type of IP Rights[20]:	Confidential Click on YES/NO	Foreseen embargo date dd/mm/yyyy	Application reference(s) (e.g. EP123456)	Subject or title of application	Applicant (s) (as on the application)
N/A					

[20] A drop down list allows choosing the type of IP rights: Patents, Trademarks, Registered designs, Utility models, Others.

Table B2: Exploitable Foreground

Type of Exploitable Foreground[21]	Description of exploitable foreground	Confidential Click on YES/NO	Foreseen embargo date dd/mm/yyyy	Exploitable product(s) or measure(s)	Sector(s) of application[22]	Timetable, commercial or any other use	Patents or other IPR exploitation (licences)	Owner & Other Beneficiary(s) involved
N/A								

[19] A drop down list allows choosing the type of foreground: General advancement of knowledge, Commercial exploitation of R&D results, Exploitation of R&D results via standards, exploitation of results through EU policies, exploitation of results through (social) innovation.

[22] A drop down list allows choosing the type sector (NACE nomenclature) : http://ec.europa.eu/competition/mergers/cases/index/nace_all.html

4. Report on societal implications

A General Information *(completed automatically when **Grant Agreement number** is entered.*	
Grant Agreement Number:	**317715**
Title of Project:	**PAASAGE MODEL BASED CLOUD PLATFORM UPPERWARE**
Name and Title of Coordinator:	**Philippe Rohou**

B Ethics	
1. Did your project undergo an Ethics Review (and/or Screening)? • If Yes: have you described the progress of compliance with the relevant Ethics Review/Screening Requirements in the frame of the periodic/final project reports? Special Reminder: the progress of compliance with the Ethics Review/Screening Requirements should be described in the Period/Final Project Reports under the Section 3.2.2 *'Work Progress and Achievements'*	**No**
2. Please indicate whether your project involved any of the following issues (tick box) :	**No**
RESEARCH ON HUMANS	
• Did the project involve children?	No
• Did the project involve patients?	No
• Did the project involve persons not able to give consent?	No
• Did the project involve adult healthy volunteers?	No
• Did the project involve Human genetic material?	No
• Did the project involve Human biological samples?	No
• Did the project involve Human data collection?	No
RESEARCH ON HUMAN EMBRYO/FOETUS	
• Did the project involve Human Embryos?	No
• Did the project involve Human Foetal Tissue / Cells?	No
• Did the project involve Human Embryonic Stem Cells (hESCs)?	No
• Did the project on human Embryonic Stem Cells involve cells in culture?	No
• Did the project on human Embryonic Stem Cells involve the derivation of cells from Embryos?	No

PRIVACY	
• Did the project involve processing of genetic information or personal data (eg. health, sexual lifestyle, ethnicity, political opinion, religious or philosophical conviction)?	No
• Did the project involve tracking the location or observation of people?	No
RESEARCH ON ANIMALS	
• Did the project involve research on animals?	No
• Were those animals transgenic small laboratory animals?	No
• Were those animals transgenic farm animals?	No
• Were those animals cloned farm animals?	No
• Were those animals non-human primates?	No
RESEARCH INVOLVING DEVELOPING COUNTRIES	
• Did the project involve the use of local resources (genetic, animal, plant etc)?	No
• Was the project of benefit to local community (capacity building, access to healthcare, education etc)?	No
DUAL USE	
• Research having direct military use	No
• Research having the potential for terrorist abuse	No

C Workforce Statistics

3. Workforce statistics for the project: Please indicate in the table below the number of people who worked on the project (on a headcount basis).

Type of Position	Number of Women	Number of Men
Scientific Coordinator		1
Work package leaders		10
Experienced researchers (i.e. PhD holders)	6	44
PhD Students		
Other	16	71

4. How many additional researchers (in companies and universities) were recruited specifically for this project?	0
Of which, indicate the number of men:	

D Gender Aspects		
5. Did you carry out specific Gender Equality Actions under the project?		No

6. Which of the following actions did you carry out and how effective were they?

		Not at all effective				Very effective
❑	Design and implement an equal opportunity policy	○	○	○	○	○
❑	Set targets to achieve a gender balance in the workforce	○	○	○	○	○
❑	Organise conferences and workshops on gender	○	○	○	○	○
❑	Actions to improve work-life balance	○	○	○	○	○
X	Other:	N/A				

7. Was there a gender dimension associated with the research content – i.e. wherever people were the focus of the research as, for example, consumers, users, patients or in trials, was the issue of gender considered and addressed?

○ ~~Yes- please specify~~

X No

E Synergies with Science Education			
8. Did your project involve working with students and/or school pupils (e.g. open days, participation in science festivals and events, prizes/competitions or joint projects)?			
X No			
9. Did the project generate any science education material (e.g. kits, websites, explanatory booklets, DVDs)?			
X No			
F Interdisciplinarity			
10. Which disciplines (see list below) are involved in your project?			
X Main discipline[23]:			
○ Associated discipline[23]:	○	Associated discipline[23]:	
G Engaging with Civil society and policy makers			
11a Did your project engage with societal actors beyond the research community? *(if 'No', go to Question 14)*	X ○	Yes	

[23] Insert number from list below (Frascati Manual).

11b If yes, did you engage with citizens (citizens' panels / juries) or organised civil society (NGOs, patients' groups etc.)?

- X **No**
- ○ Yes- in determining what research should be performed
- ○ Yes - in implementing the research
- ○ Yes, in communicating /disseminating / using the results of the project

11c In doing so, did your project involve actors whose role is mainly to organise the dialogue with citizens and organised civil society (e.g. professional mediator; communication company, science museums)?	○ X	No

12. Did you engage with government / public bodies or policy makers (including international organisations)

- ○ No
- X **Yes- in framing the research agenda**
- ○ Yes - in implementing the research agenda
- ○ Yes, in communicating /disseminating / using the results of the project

13a Will the project generate outputs (expertise or scientific advice) which could be used by policy makers?

- X Yes – as a **primary** objective (please indicate areas below- multiple answers possible)
- ○ Yes – as a **secondary** objective (please indicate areas below - multiple answer possible)
- ○ No

13b If Yes, in which fields?

Agriculture		Energy		Human rights	
Audiovisual and Media		Enlargement		Information Society	
Budget		Enterprise		Institutional affairs	
Competition		Environment		Internal Market	
Consumers		External Relations		Justice, freedom and security	
Culture		External Trade		Public Health	
Customs		Fisheries and Maritime Affairs		Regional Policy	
Development Economic and Monetary Affairs		Food Safety		**Research and Innovation**	
Education, Training, Youth		Foreign and Security Policy		Space	
Employment and Social Affairs		Fraud		Taxation	
		Humanitarian aid		Transport	

13c If Yes, at which level? ○ Local / regional levels ○ National level X European level ○ International level		
H Use and dissemination		
14. How many Articles were published/accepted for publication in peer-reviewed journals?		85
To how many of these is open access[24] provided?		58
How many of these are published in open access journals?		
How many of these are published in open repositories?		58
To how many of these is open access not provided?		27
Please check all applicable reasons for not providing open access:		
X publisher's licensing agreement would not permit publishing in a repository ❑ no suitable repository available ❑ no suitable open access journal available ❑ no funds available to publish in an open access journal ❑ lack of time and resources ❑ lack of information on open access ❑ other[25]: ……		
15. How many new patent applications ('priority filings') have been made? *("Technologically unique": multiple applications for the same invention in different jurisdictions should be counted as just one application of grant).*		N/A
16. Indicate how many of the following Intellectual Property Rights were applied for (give number in each box).	Trademark	N/A
	Registered design	N/A
	Other	N/A
17. How many spin-off companies were created / are planned as a direct result of the project?		0
Indicate the approximate number of additional jobs in these companies:		0

[24] Open Access is defined as free of charge access for anyone via Internet.
[25] For instance: classification for security project.

18. Please indicate whether your project has a potential impact on employment, in comparison with the situation before your project:

❑	Increase in employment, or	❑	In small & medium-sized enterprises
❑	Safeguard employment, or	❑	In large companies
❑	Decrease in employment,	X	None of the above / not relevant to the project
❑	Difficult to estimate / not possible to quantify		

19. For your project partnership please estimate the employment effect resulting directly from your participation in Full Time Equivalent (*FTE = one person working fulltime for a year*) jobs:

Indicate figure:

Difficult to estimate / not possible to quantify ❑

I Media and Communication to the general public

20. As part of the project, were any of the beneficiaries professionals in communication or media relations?

○ Yes X No

21. As part of the project, have any beneficiaries received professional media / communication training / advice to improve communication with the general public?

○ Yes X No

22 Which of the following have been used to communicate information about your project to the general public, or have resulted from your project?

X	**Press Release**	X	**Coverage in specialist press**
❑	Media briefing	❑	Coverage in general (non-specialist) press
❑	TV coverage / report	❑	Coverage in national press
❑	Radio coverage / report	❑	Coverage in international press
X	**Brochures /posters / flyers**	X	**Website for the general public / internet**
❑	DVD /Film /Multimedia	X	**Event targeting general public (festival, conference, exhibition, science café)**

23 In which languages are the information products for the general public produced?

❑	Language of the coordinator	X	English
❑	Other language(s)		

***Question F-10*:** Classification of Scientific Disciplines according to the Frascati Manual 2002 (Proposed Standard Practice for Surveys on Research and Experimental Development, OECD 2002):

FIELDS OF SCIENCE AND TECHNOLOGY

1. NATURAL SCIENCES

1.1 Mathematics and computer sciences [mathematics and other allied fields: computer sciences and other allied subjects (software development only; hardware development should be classified in the engineering fields)]

1.2 Physical sciences (astronomy and space sciences, physics and other allied subjects)

1.3 Chemical sciences (chemistry, other allied subjects)

1.4 Earth and related environmental sciences (geology, geophysics, mineralogy, physical geography and other geosciences, meteorology and other atmospheric sciences including climatic research, oceanography, vulcanology, palaeoecology, other allied sciences)

1.5 Biological sciences (biology, botany, bacteriology, microbiology, zoology, entomology, genetics, biochemistry, biophysics, other allied sciences, excluding clinical and veterinary sciences)

2 ENGINEERING AND TECHNOLOGY

2.1 Civil engineering (architecture engineering, building science and engineering, construction engineering, municipal and structural engineering and other allied subjects)

2.2 Electrical engineering, electronics [electrical engineering, electronics, communication engineering and systems, computer engineering (hardware only) and other allied subjects]

2.3. Other engineering sciences (such as chemical, aeronautical and space, mechanical, metallurgical and materials engineering, and their specialised subdivisions; forest products; applied sciences such as geodesy, industrial chemistry, etc.; the science and technology of food production; specialised technologies of interdisciplinary fields, e.g. systems analysis, metallurgy, mining, textile technology and other applied subjects)

3. MEDICAL SCIENCES

3.1 Basic medicine (anatomy, cytology, physiology, genetics, pharmacy, pharmacology, toxicology, immunology and immunohaematology, clinical chemistry, clinical microbiology, pathology)

3.2 Clinical medicine (anaesthesiology, paediatrics, obstetrics and gynaecology, internal medicine, surgery, dentistry, neurology, psychiatry, radiology, therapeutics, otorhinolaryngology, ophthalmology)

3.3 Health sciences (public health services, social medicine, hygiene, nursing, epidemiology)

4. AGRICULTURAL SCIENCES

4.1 Agriculture, forestry, fisheries and allied sciences (agronomy, animal husbandry, fisheries, forestry, horticulture, other allied subjects)

4.2 Veterinary medicine

5. SOCIAL SCIENCES

5.1 Psychology

5.2 Economics

5.3 Educational sciences (education and training and other allied subjects)

5.4 Other social sciences [anthropology (social and cultural) and ethnology, demography, geography (human, economic and social), town and country planning, management, law, linguistics, political sciences, sociology, organisation and methods, miscellaneous social sciences and interdisciplinary , methodological and historical S1T activities relating to subjects in this group. Physical anthropology, physical geography and psychophysiology should normally be classified with the natural sciences].

6. HUMANITIES

6.1 History (history, prehistory and history, together with auxiliary historical disciplines such as archaeology, numismatics, palaeography, genealogy, etc.)

6.2 Languages and literature (ancient and modern)

6.3 Other humanities [philosophy (including the history of science and technology) arts, history of art, art criticism, painting, sculpture, musicology, dramatic art excluding artistic "research" of any kind, religion, theology, other fields and subjects pertaining to the humanities, methodological, historical and other S1T activities relating to the subjects in this group]