European Commission

Software & Services, Cloud Computing Concertation Meeting

Shaping Europe's future for software, services and cloud

11 September, 2014

Brussels, Belgium

Published & edited by CloudWATCH

CloudWatch
A European Cloud observatory supporting cloud policies, standards profiles & services

www.cloudwatchhub.eu

Table of Contents

CloudWATCH: Welcome & Shaping Europe's future for software, Services & cloud

We welcome you to the 2nd Concertation Meeting, hosted by the Software & Services, Cloud Computing, DG Connect. The objective is to contribute to the preparation of the H2020 LEIT ICT WP2016 -2017 in the area of cloud computing and software & services. With 45 position papers already collected and in this document, we already have a set of recommendations to work from during the meeting.

The major outputs from the previous Concertation meeting back in March highlighted the importance of **clustering & convergence** between projects. This can ensure European leadership and provide a European market landscape of products, services and applications. The fruits of that meeting show in today's position papers which describe clear, promising & tangible relations between a number of related projects.

An update on CloudWATCH

CloudWATCH Global Use Case Identification. Through the Reference Model Framework Report, CloudWatch has defined a coherent information model based on the ISO/IEC & ITU-T cloud computing reference architecture (upcoming Recommendation/International Standard Y.3502/17789). The model has let us create of a set of coherent case-studies from three key stakeholders classes of user community: academic, enterprise and public sector. From these we have then selected 15, which will be used as the foundations of our use cases. With detailed information captured through multiple telephone interviews, the use cases are all hosted within a living document, which is updated and versioned upon each follow up. From these disparate case-studies a single set of user stories that can then be used to inform activities on standards profiling.

CloudWATCH & Standardisation. Numerous Cloud Standards profile workshops have been organised together with a number of EC funded projects, initiatives and foundations, and SDOs in Europe and worldwide, such as OCEAN, Cloud for Europe, EGI-InSPIRE, OpenNebula, CompatibleOne, and many more. The ultimate goal of these workshops is to elicit the need, and the scope of profiling one or more Cloud standard in order to facilitate interoperability of Cloud services and infrastructures in Europe. CloudWATCH calls on any project involved in Cloud standards implementation or development to action participation at, or co-organising, regular, open and authoritative interoperability testing events across Europe. Cloud standards, and eventually profiles, are nothing without demonstrating interoperability.

European Cloud Scout hosted on CloudWATCHhub.eu. On 10 July the European Commissioner for the Digital Agenda and Vice President of the European Commission Neelie Kroes launched the European Cloud Scout (ECS). An elaboration of a concept originally developed in Germany, this easy-to-use online tool provides European SMEs with tailored information and recommendations regarding their use of cloud services.

The EU-wide version of Cloud Scout results from a collaboration between DsiN, DIGITALEUROPE, national trade associations around Europe, and CloudWATCH. It is hosted on the CloudWATCHHub. 7 country versions are available as well as an EU-wide version in English: Belgium (Flemish & French),

Finland, Germany, Hungary, Spain and the United Kingdom. This enables Cloud Scout to serve the needs of an impressive number of SMEs across Europe.

Access the EU Cloud-Scout from http://www.cloudwatchhub.eu/cloudscout and see the implications of Cloud Computing on your own business.

CloudWATCH Main Achievements

- Initiated clustering among projects & received endorsement from VP Neelie Kroes
- Cloud certification guidelines contributing to C-SIG group
- More than 15 use cases analysed to identify user requirements and best practices for interoperability
- Identifying best practices and landscape of standards & supported by ISO
- Cloud standards profiles, collaboration with SDOs & supporting cloud plugfests
- Legal Tips & Guidelines
- 400 contents on CloudWatchHUB and launch of the EUWide Cloud Scout
- Growing community: 55 projects, 14 events, Over 600 social media followers

WP 2016 – 2017 First List of Future Innovations

We have a packed agenda for the event with over 100 participants, and 50 projects represented. Watch out for the 2 prizes that we'll be giving away for the best minute-madness presentation and the best position paper. Vote now! @cloudwatchhub

The table below provides a high-level overview, specifically looking at the recommendations and topics for consideration within the new Work Programme. These have been extracted from the 45 position papers and will be expanded on in the break-out sessions

Cloud computing	Software Engineering	Open Source
Trust in the cloud: Security; Privacy; Resilience and Disaster Recovery; Identity Management; Crypto; Certification; Assurance **Current and emerging trends in the cloud:** Internet of things; Big Data; Mobile **Optimal resource utilization:** Energy efficiency; Metering; SLA management; Contract management **Business Model and market strategies:** Decision making supporting tools; Hybrid cloud; Cloud brokering ; Consumerization	**Exploit results Better to advance innovation:** cluster, reuse, share and build from results, evolve existing tools. **Cloud Shared Services platform:** simplified exchange of shared service offerings; delivery and provision of national data sets to EU level. **Security:** industry standards; risk relevancy for adopters to formalise SaaS decision-making process. **Federated Clouds:** Use cases requiring multiple facilities or providers, possibly geographically dispersed, possibly supporting different business models.	**Technical and SLA standardisation:** improve scalability through standard protocols and cloud deployability in IaaS; machine-readable representations of services, quality, negotiations; Standardise SLAs to simplify relations, clarify service levels. **Standards for:** Web accessibility, Mobile payment, Linked Open Data, Intelligent Transport Systems and Advanced Manufacturing, and Internet/Web of Things for Smart Grids, Metering, and Cities. **Open Source Software ecosystems**: Risk analysis, detection and mitigation methods for cloud-based virtualised infrastructure and services; ecosystem monitoring tools. **Open Data and Cloud:** cloud as an

		enabler for open data management: a network of open data nodes to enable a distributed management.

We would like to thank all the European Commission representatives for their support and all project participants for their prompt, interesting & active contributions.

Organised by the CloudWATCH project consortium

Further details: www.cloudwatchhub.eu, info@cloudwatchhub.eu, @CloudWatchhub

ARTIST: A tool-supported method to cloudify applications

Name: Clara Pezuela

Organisation: ATOS

Topics recommended for the 2016-2017 Work Programme

As far as (Model Driven) Reverse Engineering is concerned, the most challenging aspects are still related to the comprehension of the existing system (e.g. notably via understanding techniques such as model slicing, as currently experienced in ARTIST): how to provide sufficient views to different stakeholders (having different roles and expertise) so that they can get a better knowledge of their existing systems for various purposes (migration is one, but there are also others). Work has been focusing notably on static/structural constructs so far and less so on more dynamic/behavioural ones.

With respect to cloud, further research can be done in cloud provider SLA compliance verification processes and benchmark tools in terms of the offered SLA metrics, including abstracted tools and methods. Also, in PaaS level abstractions in terms of deployment and cloud service usage. With regards to software evolution, the tracking of the lifecycle of software development artefacts by collecting data from different systems (e.g. SCM, BugTrackers), and subsequently correlate and mine the information to support software evolution are also items to be further researched in the near future.

Projects major results

The main result of the project is a tool-supported methodology to assist software companies in the migration of their applications to the cloud. This methodology is customised to the specific migration project and consists of three main phases: feasibility analysis, migration and verification. In the feasibility analysis phase, ARTIST provides several tools that analyse the technical and business viability of migration, helping in the decision of whether the migration is worthwhile or not. Then, in the migration phase, through reverse engineering, we gather the knowledge and understanding of the application.

Through forward engineering, we refactor and optimise the existing application and adjust the business model so it can be delivered on the cloud. Furthermore, we help users in selecting the best cloud provider for their needs. Finally, during the verification phase, we validate the behavioural equivalence of the migrated application and the fulfilment of the optimisation requirement. Moreover, we certify that it is cloud-complaint.

Potential exploitation strategy:

The project exploitation is based on a methodology that combines standard business strategy tools e.g. the Osterwalder business model canvas, among others. The process started with an analysis of the market and business requirements. This was followed up by an exploration of the project

activities to elucidate the value chain. From the value chain, business scenarios were derived and then further analysed. Once the pros and cons of each model are well understood, partner buy-in is required and a plan to set it in motion can be adopted.
The joint exploitation path selected is based on forming a perpetual ARTIST consortium, based initially on the project consortium, but open to new members, through which to fund, in cash and in kind, the sustainability of the results. Then, every partner has analysed internally their business interest in the common scenario and has decided the role they want to play and the investment, which is ready to be provided. We have identified the activities in this scenario that will generate revenue and cost. We have estimated the initial required investment and expected income, firstly to keep the project alive, but further to make it profitable. A legal agreement and ownership model will be also developed.

An update since the last Concertation meeting (March 2014)

Since our presence at the last Concertation meeting, ARTIST has intensified its collaboration with ModaClouds and PaSaage in the definition of a common cloud modelling language (CloudML). Several cross-project meetings have been organised and a group of people in each project is working on this joint task force.
The development of the language is progressing by creating model transformations that allow the understanding of the three different instances of the language in the three projects, since each project is focused on different aspects of cloud modelling (at infrastructure, platform or application levels).
A joint paper is planned once the work was completed to publicly publish the performed work. Fruitful collaboration is also happening with many support actions by defining cross-shared dissemination mechanisms (events, newsletters, social networks, etc.). These CSAs are: OCEAN, PROSE, CloudingSMEs, CloudWATCH, CloudCatalyst and SUCRE. Finally, ARTIST is cooperating with the Modisco project in Eclipse about model discovery and with the cloud working group of SPEC RG about benchmarking cloud providers and profiling applications.

Links and references

Project web site: www.artist-project.eu
Publications: http://www.artist-project.eu/publications
Project presentation: http://www.slideshare.net/hspais/artist-project-general-presentation-31984547?ref=http://www.artist-project.eu/mediacenter
Project brochure: http://www.artist-project.eu/sites/default/files/ARTIST%20Brochure.pdf

ASCETiC: Adapting Service lifecycle towards Efficient Clouds

Name: Karim Djemame

Organisation: University of Leeds

Topics recommended for the 2016-2017 Work Programme

ASCETiC addresses the issue of energy efficient computing, specifically in the design, construction, deployment and operation of Cloud services. Its approach focuses firstly on the identification of the missing functionalities to support energy efficiency across all cloud layers, and secondly on the definition and integration of explicit measures of energy requirements into the design and development process for software to be executed on a Cloud platform. Therefore, in this context, there are several topics that can be considered in future work programmes:

- Energy efficiency and big data: by reduction of data movement, by eliminating redundancies without affecting fault tolerance, and by re-using big data analysis results.
- Cloud sustainability: investigation on heterogeneous hardware usage in clouds and new virtualisation means (aka containers) to improve Cloud sustainability, as well as extending sustainability aspects to consider network operation energy consumption.

Projects major results

ASCETiC important results so far include:

- Year 1 architecture reference implementation prototype: this milestone concentrates on delivering energy awareness in all cloud system components. Monitoring and metrics information can be measured at IaaS level and propagated through the other layers of the Cloud stack (PaaS, SaaS) considering static energy profiles.
- Early evaluation of ASCETiC cloud stack focusing on software re-engineering by means of the SaaS layer modelling tools. For example, the use case News as an Asset (NewsAsset), which is an end-to end multimedia cross-channel solution for evolving News Agencies, Broadcasters and Publishers, is able to express energy monitoring requirements with augmented UML in diagrams from process or logical view.

Year 1 ASCETiC architecture reference implementation is expected to be launched in November 2014.

Potential exploitation strategy

The project has already identified a number of toolkit components that will lead to direct exploitation, e.g. SaaS modelling tool, energy modeller, Virtual Machine Contextualizer. The exploitation of the entire ASCETiC cloud architecture and its reference implementation (SaaS, PaaS, IaaS) is envisaged through the creation of a spin-off company by the end of the project.

An update since the last Concertation meeting (March 2014)

Some collaboration has started with Eco2Clouds, which addresses the ecological implications of Cloud infrastructures. Communications are taking place to jointly investigate strategies that can ensure effective application deployment and execution on Cloud infrastructures while aiming to reducing energy consumption. More specifically, how Eco-metrics collection, cloud application deployment, and evaluation mechanisms/optimisation algorithms are considered in both projects.
Early contacts have taken place with CloudScale whose aim is to make cloud systems scalable by design so that they can exploit the elasticity of the cloud. This is of interest to ASCETiC as it aims to enable the analysis of software in terms of energy efficiency at the various stages of the cloud service lifecycle: construction, deployment, and operation.

Links and references
Website: http://www.ascetic-project.eu/
LinkedIn: http://www.linkedin.com/groups/ASCETiC-Project
Twitter: @asceticproject

BigFoot: Big Data Analytics of Digital Footprints

Name: Matteo Dell'Amico

Organisation: EURECOM

Projects major results

The Bigfoot project, launched in the last quarter of 2012, represents a 3-year engagement to design and implement an all-in-one, optimised and efficient solution to the storage and analysis of large volumes of data. Using existing technologies (e.g. Hadoop and Spark) and new (e.g. NoDB) open-source projects, Bigfoot targets automatic and self-tuned deployments of storage and processing engines, enriched by several components aiming at optimising operations and at an efficient resource utilisation (OpenStack).

The current and expected impact is to heavily contribute with BigFoot components to the open-source community, both in the context of cloud computing (BigFoot is an active contributor of the Sahara project within OpenStack), and BigData (BigFoot contributes with modules that can be used to patch existing Hadoop deployments).

From the scientific perspective, BigFoot is pushing novel architectures and mechanisms for an efficient utilisation of cloud resources, new storage engines and a new integrated system that supports both batch and interactive analytics.

Potential exploitation strategy

BigFoot is central to the interests of the project's industrial partners: Symantec and GridPocket. BigFoot will accelerate data analysis in business units at Symantec: we cite Symantec Security Response, a team of experts providing 24/7 security data analysis, the Brightmail department analysing spam and the Deepsight team analysing real-time data collected by sensors distributed on the Internet. This will enable better protection for customers running Symantec's security software.

GridPocket operates on SmartGrids, a rapidly growing business. It will use BigFoot's results to enhance its platform's scalability, and to design and deploy novel techniques for metering data analysis. Applications developed within BigFoot will help utility companies and other customers of GridPocket understand and possibly reduce their energy consumption, and to increase grid reliability through the use of failure detection, consumption forecasting and load analysis.

Bigfoot's academic partners benefit by transferring knowledge via lectures and laboratory sessions, and industrial relations with partners that are interested in the project's outputs. Furthermore, BigFoot's openStack-based platform is already in use by several fellow researchers.

An update since the last Concertation meeting (March 2014)

BigFoot intends to deliver an enhanced Hadoop ecosystem, with contributions that can be packaged together as a custom distribution of Hadoop, with additional components that are not present in current stable version, in addition to several improvements to existing components of the original software.

Interoperability and compatibility is ensured by maintaining complete API compatibility with the vanilla version of Hadoop. BigFoot's contributions can also be exploited by users of other Hadoop distributions, since they are released as a set of patches to upstream software. These patches are continuously evolving, and they are released on the project's GitHub and BitBucket pages.

Within the several components that have been recently updated, we highlight a module for OpenStack's Sahara which allows automated deployment of Spark clusters; OSMeF, a measurement framework for OpenStack; an implementation of decision trees for Spark; schedism, a simulation software deployed to test and evaluate scheduling techniques; and a patched version of Hadoop which enables more effective preemption within the system.

Links and references:

Website: http://bigfootproject.eu
Twitter: https://twitter.com/bigfoot_project
GitHub: https://github.com/bigfootproject/
BitBucket: https://github.com/bigfootproject/

Broker@Cloud: Enabling Continuous Quality Assurance and Optimization in Future Enterprise Cloud Service Brokers

Name: Kahina Hamadache, Stelios Pantelopoulos
Organisation: Singular Logic

Topics recommended for the 2016-2017 Work Programme

Highly Dynamic Cloud Infrastructure: as the Cloud becomes more and more omnipresent and intent to serve an even larger audience, targeting large companies as well as isolated user, the necessity for Cloud Services to frequently adapt and provide a large spectrum of performances is arising. In this perspective, it will be necessary for cloud service providers to have fine-grained and highly dynamic, fast down- and up-scalable capacities.

Cloud Services providers already have the ability to scale up or down their services, but in order to remain competitive in a more fiercer and fiercer market, they will need such highly dynamic mechanisms. Cloud and Internet-of-Things: the forthcoming and foretold explosion of the Internet-of-Things will bring new challenges for the cloud. Indeed, if current cloud services consumers are already numerous, the Internet of Things aims at connecting everyday life's objects to the web and making them smarter through data sharing and adaptation. This will inevitably imply an important increase on the consumption and needs for computing power, storage, networking capabilities, etc.

In this perspective it would be interesting to consider how this overhead will impact the cloud service providers, and if a differentiation needs to be made between the human and "things" use of cloud services. Certification mechanisms for Cloud Services: as the Cloud Market is growing bigger, the number of actors is also growing, making it even more difficult for services consumers to choose a service or another. To tackle this issue, research on cloud computing has considered the existing reputation techniques and is exploring specific trails.

A potential perspective for the evolution of the cloud would be to deliver certificates to trusted cloud vendors and brokers, based on their sustained performance, respect of SLAs and overall efficiency. Such approach is perfectly in line with the outcomes of the Broker@Cloud project, as project's results should be used as a basis to allow such sustainability and meet certifications requirements.

Projects major results

Multiple dissemination activities are scheduled: participation to ESOCC 2014 workshop [3], CrossCloud Brokers [4], CloudCom 2014 [5] The project has reached an important implementation milestone as the different cloud service description and interfaces components of the framework are now available (internally). The continuous quality assurance and optimisation mechanism are currently being finalised and will be soon available. On this aspect, the main contribution will be the

release of the PuLSar (Preference cLoud Service Recommender) component built on the top of SMICloud [6] and using Analytic Hierarchy Process to reach decisions.
Failure prevention and recovery asset is also under implementation. The DiVA reasoner [7], at the centre of the module has been re-implemented in Java for portability, simplicity and performance concerns.

Potential exploitation strategy

The consortium has worked for a long time now on the elaboration of a consistent and efficient exploitation strategy, allowing a wide diffusion and an optimisation of profits, spin-offs and else. In this perspective, it was decided some months ago to propose the implemented results as open sources modules under Apache [1] Dual License. Hence, organisations wishing to reuse Broker@Cloud's code and publish their developments as Open Source will be granted a free license, while organisations using Broker@Cloud code within a proprietary non-open source software will have to pay for a non-free license.
The consortium is currently working on two distinct topics: harmonising the individual efforts and exploitation strategies of the different partners; and in a second time it is working on the release schedule of the different modules implemented through the project. This last point is especially important as they proposed modules may face competition from similar projects. In this perspective the consortium is also actively working on the actions to undertake in case a similar product is released before the readiness of Broker@Cloud assets.

An update since the last Concertation meeting (March 2014)

A list of projects closely related to Broker@Cloud interests has been created to help considering the potential bridges and collaborations with B@C. Among the final list of six projects (CAST, DIVA, FIWARE, MODAClouds, MOTIVE and PaaSage), 2 of them have been identified as being sufficiently close to foresee a collaboration as they share some of B@C concerns (cloud service description, cloud service failure prevention and recovery, optimisation of cloud services): MODAClouds and PaaSage.
It is expected that the bridging between the different projects will be eased by the participation of one of the consortium member (SINTEF) in all of them, smoothly ensuring the liaison and coordinated efforts. In this perspective a shared workshop between PaaSage and Broker@Cloud will be held at Cloud Services Brokerage 2014 [2] in order to refine the cooperation points between the projects and plan the future actions.
CSB 2014 will also be a valuable opportunity for the consortium to make contact with additional projects as well as consider the latest evolution on Cloud Brokerage research and estimate the impact of the project.

Links and references

[1] http://www.apache.org/licenses/LICENSE-2.0.html
[2] http://csb2014.modelbased.net/
[3] http://esocc2014.cs.manchester.ac.uk/
[4] http://www.comp.lancs.ac.uk/~elkhatib/crosscloud/
[5] http://2014.cloudcom.org/

[6] Garg, S. K., Versteeg, S., & Buyya, R. (2011, December). Smicloud: A framework for comparing and ranking cloud services. In Utility and Cloud Computing (UCC), 2011 Fourth IEEE International Conference on (pp. 210-218). IEEE.
[7] https://sites.google.com/site/divawebsite/divastudio/diva-reasoning-framework

CACTOS: Context-Aware Cloud Topology Optimisation and Simulation

Name: Jörg Domaschka

Organisation: Ulm University

Topics recommended for the 2016-2017 Work Programme

Cloud data centre resource optimisation is a difficult and important problem. The scale of large cloud data centres alone motivates the need for research into improving service performance and resource utilisation in these environments. From the experiences in the CACTOS project gained so far, we have identified the following areas that need further research: With respect to cloud simulation, more detailed yet manageable and meaningful infrastructure models for data centres are needed that take into account the complexity of the real-world data centres' virtualisation infrastructure.

From a modelling and operator point of view prediction capabilities for non-functional properties of complex distributed software systems are just as necessary as approaches for merging Cloud data centre operation and architecture.

Finally, we claim that a much tighter integration of data centre facilities with virtualisation is required. For instance, it is currently unknown how network bandwidth could be assigned dynamically to virtual machines based on their dynamic needs.

Projects major results

The overarching goal of the CACTOS project is to provide autonomic and interactive tools to support and improve the efficiency of cloud data centre operations. The three CACTOS tools CactoOpt, CactoSim and CactoScale build the foundation to craft, evaluate, and improve data centre design, operation, and placement optimisation algorithms.

CactoOpt is designed to facilitate development of advanced optimisation mechanisms capable of both resource-level scheduling optimisation as well as holistic data centre-level optimisation. CactoScale provides data filtering and correlation analysis tools which will run on vast volumes of data generated from large data sets from both physical nodes and virtual machines. CactoSim enables the evaluation of optimisation strategies at design time based on various load and data centre models.

All tools are integrated in the CACTOS Cloud data centre optimisation and management toolkit and their interplay shall improve energy efficiency, overall decision making, and allow accurate but low effort blueprint-based analyses. The functionality will be validated in three real world use cases ranging from operating cloud platforms over an enterprise application to the needs of scientific computing.

At the current stage after year 1, we have established a model comprising physical and virtual components, as well as load patterns. We further enabled that CactoOpt is fed by the correct monitoring data including CPU load, memory usage. First mining algorithms have been established to trace the I/O behaviour of applications. CactoSim has been integrated with the live data centre topology optimisation prototype CactoOpt enabling the evaluation of arbitrary runtime

optimisations at design time. A first initial public release of CactoSim is due at the end of September 2014.

Potential exploitation strategy

CACTOS will engage a coherent and converged exploitation strategy at the heart of which is the concept of proving concrete impact from the project results. Critical to the strategy of providing clear impact from project activity and sustainability beyond the project lifetime is identifying the target markets for CACTOS and then generating a community with this market in mind.

A primary goal then, is building a community around the CACTOS tools in order to maximise impact and minimise effort for each member. With that in mind, the consortium currently evaluates the parts of the code to release as open source software, producing a set of exploitable tools mapped to use case validation and potential target users. This eases the community building and spreading the knowledge about CACTOS with the potential target groups including Cloud operators, Cloud infrastructure providers, data centre operators, cloud middleware developers and any company and scientific institution running a reasonably sized private cloud.

In the period immediately after the project's completion, the results will be used for educational purposes by all academic partners, but also for commercial purposes: Flexiant is going to integrate exploitable entities in their commercial offerings.

The University of Ulm will immediately apply the results and insights from CACTOS to their campus-wide data centre to ease its operation and offer automatised support for its scientific customers. Several partners envision a close future cooperation with other partners. Finally, partners from both academic and commercial domain are considering potential spin offs.

An update since the last Concertation meeting (March 2014)

A cooperation including research and tools exchange with the CloudScale project has been established. Further, we are investigating the re-use of PaaSage tools for multi-cloud and cross-cloud deployments.

Links and references

http://www.cactosfp7.eu/

CELAR: Automatic, Multi-Grained Elasticity-Provisioning for the Cloud

Name: Ioannis Konstantinou

Organisation: ATHENA/IMIS

Topics recommended for the 2016-2017 Work Programme

In CELAR, various interesting topics which appear to have significant impact along with interesting research challenges have been identified:

- Complex cloud-enabled multi-module systems present performance degradations, which cannot be easily identified and attributed to specific parts. Automated application profiling to allow bottleneck detection of specific modules, metrics of interest, etc, would be interesting.
- When quick and specific identification and healing is required, things become more interesting and challenging. Healing actions should focus not only on infrastructural changes, but they should consider application specific configuration changes.
- Towards generic applicability, the automation and abstraction of the procedure of application topology and elasticity actions definition in order to support a wider range of cloud applications is required. These procedures need to be as application agnostic as possible.
- The aforementioned topics can contribute to the vision of Software Defined Everything [1] (data-centres, storage, network, etc.), by enabling the programmatic definition of resource elements, their interaction and their management rules.

Projects major results

CELAR will offer a fully automated fine grained elasticity mechanism that can be utilised by cloud applications to adjust their resources according to user-defined policies when resource needs change. CELAR consists of tools for describing the application structure using TOSCA semantics (c-Eclipse [2]), real time monitoring application performance (JCatascopia [3] and MELA [4]) and a decision making module consisting of an elasticity controller [5] and a definition language [6]. These modules are integrated following an approach described in [7] and they provide the CELAR elasticity platform.

The platform is integrated with SlipStream [8], a tool that automates deployment and materialises elasticity commands into two different cloud platforms, namely the okeanos [9] and Flexiant's FCO [10] using open-sourced cloud interoperability libraries such as libclouds and jcloud. All the aforementioned are open-sourced [11] and offered as an easy to install packages through CELAR's repository. Two applications that will utilise CELAR's outcomes are currently being under development from PlayGen [12] and the Cancer Research UK Manchester Institute [13].

Potential exploitation strategy

CELAR will offer its results as an open-source package that can be downloaded and utilised. An active community that will consist of both users that utilise CELAR to manage their applications and

developers that enhance the already existing code is required, so that there will be a critical mass that will make the software sustainable and self-evolving.
Participating IaaS cloud providers can utilize the knowledge gained during the course of the project in order to configure deploy and offer, for instance, IaaS elasticity as a service (through GRNET's okeanos and Flexiant's FCO platforms). Software companies that build and deploy cloud applications with varying resource needs such as Playgen can utilise CELAR to offer highly elastic PaaS applications that can be deployed on top of any cloud, since CELAR has connectors for most popular cloud APIs.

An update since the last Concertation meeting (March 2014)
We have identified the relevant target markets that could utilise the CELAR offering. We have published several highly-detailed blogs [14] on the CELAR project results and we have utilized social media channels to share project activity with relevant audiences. We had actively participated at several events, including CloudScape IV, GeNeDiS and FIA where we presented on the CELAR solution.
The following collaboration and clustering activities are taking place:

- CELAR is considered for utilisation in the cloud platform developed in the GR-funded Modissense [15] project to handle unpredictable workload spikes.
- JCatascopia [3] is currently modified and will be incorporated in the PaaSport project [16].
- COMOT [5] is exploited together with Pacific Controls Lab to creating elastic machine-to-machine applications and with the Faculty of Electrical Engineering of Zagreb, Croatia for "elasticising" a SaaS cloud application used in Croatian market.
- Along with APICe research laboratory of University of Bologna we are exploiting rSYBL [6] for coordination based elasticity control mechanisms.
- Together with ISA research group of University of Seville, we are exploiting rSYBL [7] and MELA [4] for SLA management through elasticity.

Links and references
[1] http://www.wired.com/2013/05/are-you-ready-for-software-defined-everything/
[2] "c-Eclipse: An Open-Source Management Framework for Cloud Applications", C. Sofokleous and N. Loulloudes and D. Trihinas and G. Pallis and M. Dikaiakos, (EuroPar 2014), Porto, Portugal 2014.
[3] "JCatascopia: Monitoring Elastically Adaptive Applications in the Cloud", D. Trihinas and G. Pallis and M. D. Dikaiakos, "14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing" (CCGRID 2014), Chicago, IL, USA 2014.
[4] "MELA: Monitoring and Analyzing Elasticity of Cloud Services ", D. Moldovan, G. Copil, H. Truong, S. Dustdar, "5'th International Conference on Cloud Computing (CloudCom)", Bristol, UK, 2-5 December, 2013.
[5] H. Truong, S. Dustdar, G. Copil, A. Gambi, W. Hummer, Duc-Hung Le, D. Moldovan, "CoMoT - A Platform-as-a-Service for Elasticity in the Cloud", IEEE International Workshop on the Future of PaaS. Colocated with IC2E, Boston, Massachusetts, USA, 10-14 March 2014.
[6] G. Copil, D. Moldovan, H. Truong, S. Dustdar, "SYBL: an Extensible Language for Controlling Elasticity in Cloud Applications", 13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid), May 14-16, 2013, Delft, the Netherlands.

[7] Tsoumakos, D., Konstantinou, I., Boumpouka, C., Sioutas, S., & Koziris, N. (2013, May). Automated, Elastic Resource Provisioning for NoSQL Clusters Using TIRAMOLA. In Cluster, Cloud and Grid Computing (CCGrid), 2013 13th IEEE/ACM International Symposium on (pp. 34-41). IEEE.
[8] http://sixsq.com/products/slipstream.html
[9] Koukis, V., Venetsanopoulos, C., & Koziris, N. (2013). ~ okeanos: Building a Cloud, Cluster by Cluster. IEEE internet computing, 17(3), pp. 67-71.
[10] http://www.flexiant.com/flexiant-cloud-orchestrator/
[11] https://github.com/CELAR
[12] http://playgen.com/
[13] http://www.cruk.manchester.ac.uk/
[14] http://www.celarcloud.eu/newsroom-cat/blog
[15] http://www.modissense.gr/index.html
[16] EU-funded PaaSport project, http://paasport-project.eu/

CLIPS: Cloud approach for Innovation in Public Services

Name: Ianfranco Marasso

Organisation: Engineering Ingegneria Informatica

Topics recommended for the 2016-2017 Work Programme

EC to support joint exploitation strategies/plans for open source software - Many European research and industrial organisations have demonstrated their ability to produce world-class open source software. Moreover, the open source "market" could represent one of the leverages to support cost reduction needed in Public Administration domain, while maintaining high level of quality in service delivery.

Open Data and Cloud: cloud computing could be seen as one of the enablers for open data management. In particular, thinking at a federated network of open data nodes to enable a distributed management, service description, comparability, metrics, and validation of service levels. More research is required on high level languages to define SLAs and describe services to provide basis for evaluating cloud service equivalence. More research on tools and methods for monitoring applications in the cloud and conformance of services to SLAs are required, together with accounting and billing capabilities.

Open processes: instruments and approach supporting the collaboration of all the possible stakeholders in the definition, the approval and the monitoring of public administration process toward a more efficient and transparent management mechanism. In particular, the last three points should be driven following the open innovation approach supporting the shift towards a new paradigm in PA domain; Government as service (GasS).

Projects major results

The major result of the project will be the enablement of a new approach for cloud based Public Service provision, involving representation from a range of Community stakeholders (based around a PPPP – Public-Private-People-Partnership). This could be achieved through both technical and methodological solutions.

- The definition and use of a visual editor in order to provide an easy data mash-up capability to all the actors involved.
- The integration and use of a eID system, STORK in order to enable cross border capability.
- The definition and implementation of an Open Data nodes network be able to manage cites open data in a federated way.
- The analysis of privacy issues related to cross border aspects within countries involved in the project (Germany, Italy, Serbia, Spain, and the United Kingdom).

At the moment the overall CLIPS methodology has been defined together with a common understanding about actors, roles and main software components, moreover a first version of security architecture has been released.

Potential exploitation strategy

Following the end of the CLIPS project, revenue will be generated when both public and private organisations utilise the CLIPS platform to host applications/services:

- Public Administrations (PAs) who move existing services to the cloud during the project will have a cost reduction in having their services implemented with the involvement of third parties.
- SMEs internal/external to the project will generate revenue when they develop new services in response to public sector tender requirements to host on the CLIPS platform.
- From Large Enterprise point of view, the aim is to consolidate a set of tools in order to provide to potential customers all over the Europe (mainly PAs and SMEs) a complete framework for service delivery, consolidating the role of cloud service provider.
- The Consortium will generate revenue when SMEs external to the project pay a license fee to host applications they have developed in response to public sector requirements on the CLIPS platform. The CLIPS platform is attractive to SMEs as it provides access to open source data sets, the CLIPS user community and the CLIPS instruments.

The most important factor in ensuring this sustain the project outputs is the commitment of local authorities to make available in the CLIPS cloud their digital services and open data, and allow local authorities to use these services as a template for their own new services, and allow the use of open data by SMEs to create new services.

An update since the last Concertation meeting (March 2014)

CLIPS has some contact points with CLOUT in terms of cloud architecture and instruments to be provided to users. Even if former pays attention more on cross border aspects while the latter on the integration between cloud, the project both are addressing Public Administration as main final user. Possible synergies are being identified. At the moment, project month six has just passed (so still at an initial stage).

We are concentrating on the scenarios to be piloted within the projects together with architecture specifications be able to support those scenarios. With reference to the architecture we have been investigating on following technologies to be used/integrated: STORK as eID system, Talend as platform for services integration, CKAN as open data management system, OpenStack as cloud computing platform.

Links and references

CLIPS: http://www.clips-project.eu
CLOUT: http://clout-project.eu
STORK: http://www.eid-stork.eu
CKAN: http://ckan.org
OpenStack: http://www.openstack.org
Talend: http://www.talend.com

Cloud for Europe: Cloud adoption in the public sector

Name: Linda Strick
Organisation: Fraunhofer FOKUS

Topics recommended for the 2016-2017 Work Programme

Especially relevant from the perspective of cloud adoption in the public sector it is important to address a European-wide cloud infrastructure and service delivery platform as a "European Cloud Shared Services" platform. This simplifies the exchange of Shared Service offerings in the public sector and potential other sectors as well.

As a good practice example the Open Data efforts can be seen, where national offerings of data sets are delivered and provided to the European level as well. This seems a relevant European cloud application area. In the area of "European Cloud Shared Services", legal and organisational issues are equally relevant, accompanying the technical topics. A good example can be seen in the UK's "G-Cloud", but on a European level.

On the wish list is as well a better cooperation between the different DG Units. For example, it is interesting to investigate which parts of the initiatives in the Customs Multi-Annual Strategic Plan (MASP) can be moved to the cloud for the European Customs organisations, which is under control of DG TAXUD. Important for a trusted cloud computing, which is required by the public sector are as well the following topics:

- Privacy enhancing technologies for the cloud.
- Transparency technologies: such as log management and event monitoring and correlation tools.
- Service selection tools based on risk assessment and asset evaluations.
- Identity and access control technologies for both users and service provider personnel.
- Energy awareness in virtualised environments (IaaS, hypervisors, OS components).
- Resource awareness and resource usage adaptation in virtualised environments (IaaS).
- Flexible pricing and SLA models for IaaS with dynamic differentiated pricing (e.g., auction-based, depreciation-based, etc.).

A very relevant topic is the standardisation of SLAs in cloud computing according to the type of cloud it is running on. Terms and conditions in the SLA depend on the complexity of control variables that the provider gives to the consumers. In that context we highlight the SaaS SLAs as a possible concern to the main topic:

- **Standardise SLAs for cloud computing**: Reference the importance to standardise SLAs in multi-cloud environment. When terminologies are not standardised, the definition of a terminology that a partner (a requesting service provider) uses may differ from the definition that another partner (a substitute service provider or vendor) uses. The expanding

value chain for other services has made SLAs important for a myriad of often complex relationships between partnerships.

- **SaaS SLAs:** Some considerations for SaaS service-level agreement outlines the service levels and brings diverse customer-focused into a single guide for IT and business leaders who are considering cloud adoption.
- **Cloud Computing SaaS Security issues and risk:** The issue of risk relevancy is significant if adopters are to formalise the SaaS decision-making process. The results provide evidence that certainty about some elements of security, business continuity, and integration significantly influences the adopting organisation's level of satisfaction with its overall SaaS experience.

Projects major results

Cloud for Europe adjusts public sector requirements and establishes suitable contractual terms for future cloud procurements. The project's main result focuses on the adoption of cloud computing solutions in the public sector.

At first, obstacles for the use of cloud computing in the public sector are addressed by analysing requirements of the European public sector, the current market situation and the relevant standardisation landscape.

To overcome these obstacles challenges will be identified, which are of joint interest between several Member State public procurers, and which are subject to R&D to from the basis for cloud adoption. A pre-commercial procurement tender process will address industry, and especially SME's, to participate as bidders for those new cloud solutions.

Potential exploitation strategy

As we are currently in defining the tender package, therefore not much to elaborate, but the IPR aspects will be kept with the bidders. The PCP process allows for selecting the best results for each of the tender phases, so that at the end the best solutions will be published and shared with other member states, public procurers and industry.

An update since the last Concertation meeting (March 2014):

- Partnering with Cloud projects, such as: SUCRE; CLOUD CATALYST; OPEN DAI; STORM CLOUDS; COC CLOUDS and Cloud WATCH - latter with common workshop on Legal issues in Cloud computing in June 2014 in London.
- Cooperation with ENISA for Cloud certification.
- Partnering with PCP projects, such as: CHARM, Silver, PRACE-IP3 in terms of lessons learned for the PCP process.

Links and references

www.cloudforeurope.eu

CloudCatalyst: Tools to accelerate Cloud adoption in Europe

Name: Andreia Jesus

Organisation: Portugal Telecom

Topics recommended for the 2016-2017 Work Programme

More measures should be taken to improve European businesses' position in the Cloud Computing market, fostering the emergence of a strong and enthusiastic community of cloud adopters and supporters in Europe. Define together with EU startups (or their representatives) actions to support cloud business model innovation and go-to-market strategies joining large and small companies.

Future projects should have a very strong market orientation and be built upon business and technological experience of top industrial and research partners that are seen as a reference in the area of Cloud Computing technologies. New topics should be focused on the creation of real businesses, i.e. new companies, with real products and services, to boost EU businesses competitiveness.

Most important results

- Strategic planning for cloud innovation: Investigate business issues in Cloud Computing. The work developed will take into account socio-economic and technical issues, as well as consumer‟s requirements and provider‟s main challenges, especially SMEs.
- Creation of a cloud accelerator toolbox: Promoting active exchange of knowledge and requirements and common interests in implementing collaboration models for the software community and entrepreneurs. The main output will be a Cloud accelerator toolbox with guidelines and best practices for the developers‟ community, entrepreneurs, technical transfer units, start-up incubators and other stakeholders that can have an important role in stimulating Cloud Computing uptake.
- Development of a Go-to-the-Cloud support service: CloudCatalyst will implement a support service for information sharing and dissemination, coaching and consulting to EU SMEs and other key stakeholders interested in the development and implementation of Cloud solutions. The project will study the impact of using open-sourced and proprietary support technologies and strategies for enhancing the full capacities of both models. The Go-To-The-Cloud support service will also have a special area dedicated to consumers/end-users.

Identification of enablers for Cloud innovation by studying the current cloud trends in different EU countries in order to identify the major potential for advanced uses of Cloud solutions:

- Benchmarking exercise completed, identifying the main technologies, support policies, funding available and EU incubators to assess the remaining gaps, strengths, weaknesses and opportunities for cloud computing.
- Cloud Computing Trends survey has been done (more than 260 answers collected so far), analysing the main drivers and barriers to cloud adoption, the main market players, and important data security issues.

- Definition of critical success factors to overcome most common challenges to cloud expansion. This report has been created using the insights of coordinators of EU funded projects related to cloud computing.

Potential exploitation strategy

CloudCatalyst vision to enable sustainability based on the project results will produce a holistic strategy for Cloud Computing stakeholders with realistic and innovative measures to implement the envisaged business and sustainability models in key EU industries. A set of recommendations will also be produced and issued to European/national/regional authorities and private companies with the purpose of stimulating the local stakeholders to adopt Cloud solutions, leading into future successful (e.g. value for money, in budget, on time, sustainable) deployments.

In January 2015, we will set indicators to evaluate the success of the proposed action plan and helping establish the framework for future sustainability. In September 2014, we will define required actions and conditions under which the project outcomes might become a basis for development of cohesive set of longstanding activities.

Specific measures for ensuring that CloudCatalyst is a sustainable long-term initiative includes seeking funds from the industry to complement European Commission support use the Cloud accelerator toolbox and go-to-the-cloud platform to attract new partners.

An update since the last Concertation meeting (March 2014)

Joint workshop will be held in Slovenia on 26 September 2014 within the Class 2014 event organised by the SUCRE, OCEAN, CloudCatalyst, and CloudForEurope projects. We will address topics related to the adoption of Cloud Computing in Europe and the challenges that still have to be addressed.

The event will examine how Open Clouds can serve, or are already serving, the societal needs across Europe, through their uptake by key European sectors, such as public administration and the healthcare industry.

Focus will be given on horizontal issues such as security, trust and privacy. In the same context, the workshop will also serve as a vehicle to deliver interesting opportunities to the new generation of European business startups and SMEs.

Also we approached OpenNebula community in respect to bridging IaaS challenges, MODAClouds, PaaSage and ARTIST on possible collaboration activities between projects on building common eco system/community, develop combined sustainability strategy.

We have developed a survey to establish the extent to which existing European research meets encountered and envisioned technical challenges for using cloud technologies, to which responses have come from the coordinators of PROSE, ORBIT, STORM CLOUDS, CLOUDSCALE, COMPOSE, S-CASE, VENUS-C, CLOUDWATCH, CUMULONIMBO, ASCETiC, OPTIMIS, MO-BIZZ, U-QASAR, CLOUDSPACES, BROKER@CLOUD, HOST, SUCRE, ENVISAGE, CLOUD TM, CLOUDWAVE, BIGFOOT, ARTIST, BETAAS, PAASAGE, PANACEA, MODACLOUDS, I2WEB, CLOUDINGSMES, HTML5APPS. Based on this approach a deliverable was produced on critical success factors for cloud expansion.

Links and references

Newsletter: http://www.cloudcatalyst.eu/wp-content/uploads/2013/10/D.5.4.-CloudCatalyst-Press-Release-1.pdf | Press Release: http://us3.campaign-archive2.com/?u=e7a78fb9bfb1de1658bd3abb9&id=a23662d328&e=

CloudScale: Resolving scalability

Name: Gunnar Brataas
Organisation: SINTEF ICT

Topics recommended for the 2016-2017 Work Programme

- Improve the scalability engineering method so that it becomes less costly and easier to use and can be used on a broader scale.
- Improve the tools required for scalability engineering in the areas of spotting scalability problems in existing code, but also in analysing models as well as in extracting models from existing code and in capturing resource demands parameters from existing systems.
- More case studies demonstrating the benefit, but also cost of scalability engineering. This will increase the knowledge of scalability engineering, so that scalability engineering can be used more widely.

Projects major results

CloudScale will make a tool suite enabling:

- Improving the scalability and cost-efficiency of existing systems by spotting bottlenecks and anti-patterns in the code, both by analysing the running system (dynamic spotting) as well as by analysing the source code (static spotting).
- Analysing the scalability and elasticity at design time.
- Extracting models from source code to enable analysis.

We are now busy finishing almost complete tools in all these areas and are also working on validating these tools. In the next and final year we will focus on improving these tools using the input from the validation. Dissemination and exploitation will also get more attention.

Potential exploitation strategy

Effective exploitation of the results of the CloudScale project depends on the impact the project can have in the domain of scalability in the Cloud. We have identified four main factors in order to have a sustainable practical impact:

- A successful internal exploitation of the results by the consortium partners in order to enhance their own products, provide better services, and advance their research contributions, exemplifies the benefits of applying the CloudScale results.
- A sustainable exploitation is only possible if the main scientific results and tools developed in the project are open source, easily accessible and available for commercial use.
- Presentation of convincing demonstration and value propositions will demonstrate the applicability of CloudScale results to real business scenarios.

- Raising awareness for CloudScale together with contributions to international standardisation in this domain play a crucial role in sustainable exploitation of the results of the project.

An update since the last Concertation meeting (March 2014)

Since the last meeting CloudScale has been active in the following project collaboration and standards related groups:

- LIMBO by Universitat Wurzburg is used as part of our usage evolution tool.
- Open Cloud Directory where we have added CloudScale's showcase CloudStore: www.ocdirectory.org/organisation/cloudscale
- SPEC RG Cloud Group: where we presented a live demonstration of CloudStore at their annual meeting on 26-27 March in Dublin. The RG showed interest and considered using CloudStore as well.
- Palladio Community where we promote our open source strategy using Git as well as CloudStore.
- FZI: ("Forschungszentrum Informatik"; German technology-transfer institute). We convinced the FZI to build on top of CloudScale's Analyser in the context of a European project related to energy-efficiency ("CACTOS – Context-Aware Cloud Topology Optimisation and Simulation").

Links and references

CloudScale web site: http://www.cloudscale-project.eu Limbo web site: http://se2.informatik.uni-wuerzburg.de/mediawiki-se/index.php/Tools

SPEC RG Cloud Group: research.spec.org/working-groups/rg-cloud-working-group.html

CloudSpaces: From application-centric to person-centric models

Name: Pedro Garcia López

Organisation: Universitat Rovira i Virgili

Topics recommended for the 2016-2017 Work Programme

We consider that Software-Defined Technologies will have a relevant role in the next years. Software will mandate the autonomic management of cloud services (network, storage, computing) simplifying their massive adoption. Novel programming abstractions and middleware systems will be required to interact with several autonomic software layers.

Another important topic is about decentralised overlay technologies and decentralised software systems linking the next generation Internet. Decentralisation is at the heart of the Internet, and the massive deployments of fibre technologies will considerably facilitate edge-to-edge approaches. In particular, since the network layer is ossified, novel overlays combined with Software defined technologies will facilitate the interoperation of heterogeneous resources. A good example is the transparent interconnection of data centres and home computing resources across Europe for creating novel applications.

Finally, cyber-physical systems are going to be of paramount importance. Mixed reality environments will require advanced distributed computing platforms with extreme scalability. Personal Information Systems will drive our interactions with Digital Information and social networks. Thanks to immersive devices and 3D interfaces, a next generation of Personal Systems will emerge capturing the best of Personal Clouds, Social Networks, and Virtual Desktops.

Projects major results

CloudSpace aims to provide an advanced scalable data management agent offering advanced storage, sync and share mechanisms. On top of which we devised novel adaptive replication algorithms providing dynamic membership reconfiguration of untrusted repositories as well as advanced consistency mechanisms.

At this moment, we already have a working prototype called StackSync. StackSync is an open source personal cloud built on top of OpenStack Swift that includes novel sharing and interoperability capabilities. It was presented officially to the OpenStack community in the last Summit in Atlanta [1]. Due to the impact of the presentation we had several deployments of StackSync around the world in countries such as Spain, France, Vietnam, Czech Republic or Switzerland among others. The URV is also evaluating StackSync with a subset of their students in order to consider it as an official synchronisation solution and offer it to the whole community. Moreover, we also released a key-

value hybrid cloud storage system that robustly replicates data over untrusted public clouds while keeping metadata on trusted private premises called Hybris.
We planned to implement techniques ensuring secure, trust-worthy and privacy friendly interactions within the CloudSpaces environment. To this end, we proposed a set of privacy-aware data sharing mechanisms that will employ hybrid technical approaches like obfuscation, anonymisation, encryption, digital signatures, and information hiding.
We have already released a prototype called PrivySeal, a software that provides easily usable privacy technologies to end users who want to share their data on any cloud service. It is based on published research and employs crowdsourcing and a psychologically grounded model for risk estimation.
We avoided vendor lock-in thanks to both semantic and syntactic interoperability techniques. We will define a new standard to share and exchange resources between heterogeneous Personal Clouds. CloudSpaces has already defined a freely-implementable and generic specification of a protocol for allowing Personal Cloud interoperability.

Potential exploitation strategy

There is a big opportunity for those IaaS providers who want to offer an innovative Personal Cloud to their customers. Cloud storage is a real need for the foreseeable future, so organisations will demand efficient and secure solutions. That's why an IaaS provider can take advantage of StackSync by providing clients with the following services: Public or Hybrid clouds based on the StackSync technology; and reselling cloud storage services as a "white label" cloud.
StackSync will help IaaS providers to deploy an enterprise Personal Cloud for thousands of users at low cost. They will be able to update their services portfolio and engage customers with a Personal Cloud specially designed for organisations. Software providers will improve their solutions building them on top of the StackSync platform and offering innovative services to their clients. They can take advantage of cloud storage features and bundle their software with a Personal Cloud, focused on the organisation's needs: security, scalability and openness.
If a software provider wants to partner with an IaaS provider, StackSync can be deployed by the IaaS partner so the software provider would only need to care about its solution development. StackSync proposes two deployment scenarios focused on SMEs in order to boost their storage in the cloud: StackSync-P (Private Cloud), and StackSync-H (Hybrid Cloud). StackSync-P can be deployed on-premise, so that companies can finally retake the control of their data.
SMEs benefit from OpenStack Swift as a proven open-source cloud technology that can be adapted to their storage needs on the go. StackSync-H keeps metadata of all outsourced data on-premise and provides robustness by replicating data across multiple commodity clouds, in a fully transparent way.
SMEs will benefit from the best of both Public and Private Clouds. StackSync-H can be tuned to optimize reliability, cost and performance. Public institutions like governments, hospitals, or universities have strong security and privacy requirements, so it's very important for them to use a Personal Cloud storage that meets these restrictive needs.
StackSync keeps citizens data in a cloud storage platform that let administrators manage this big amount of information without breaches in their data security, nor big investments on IT infrastructure. They can choose between private, public and hybrid platforms, depending on their needs.

An update since the last Concertation meeting (March 2014)

CloudSpaces is already collaborating with other European projects. We are working with the Intertrust project in attribute-based encryption (ABE) and on the definition of privacy policies. There is also an ongoing partnership with the Leads project about elastic synchronisation and with the Confine project in decentralised network technologies and bittorrent.
We are working together with IBM Haifa on generating compressible workloads for storage benchmarks. The work resulted in a system that can generate data that is compressible to a prescribed extent and whose compression takes a certain amount of effort.
There is also a collaboration with our partner NEC to implement the specifications of the CloudSpaces' interoperability protocol that would allow NEC's personal cloud and StackSync to share and exchange resources as a proof-of-concept.

Links and references

[1] StackSync at the OpenStack Summit Atlanta 2014
https://www.openstack.org/summit/openstack-summit-atlanta-2014/session-videos/presentation/stacksync-a-dropbox-like-personal-cloud-for-openstack-swift

CloudWave: Agile Service Engineering for the Future Internet

Name: Eliot Salant
Organisation: IBM

Topics recommended for the 2016-2017 Work Programme

CloudWave is concentrating on improving the efficiency and resiliency of the execution of software as a service (SaaS) both through the adaptation of Cloud and/or the service itself, as well as through the use of DevOps principles to feedback recommendations to the next iteration of the software service.

We believe that enabling the agile adaptation of the Cloud and its hosted services is an important one, and there remains a large number of research issues that fall into this category, such as improved distributed placement optimisation based on a workload awareness. Additionally, as more and more applications start to shift to the Cloud model, we expect serious issues to surface, affecting both the security and privacy of the data used by these applications for reasons such as:

- Typically, legacy applications were not developed by design with the thought of how to ensure multi-tenancy. Porting this code to a SaaS environment might suddenly expose flaws in keeping data isolated between users of the Cloud.
- Newly developed applications, especially those developed under the time pressure of short development cycles, may accidentally contain security holes. Even worse, applications such as in a mobile marketplace may have been maliciously developed with backdoor security breaches.
- Brute force methods such as simply obtaining a full separation between users through full isolation of execution instances (e.g. use of containers) is not an efficient solution for SaaS, as will it not allow for resource sharing.

We recommend the funding of efforts that will extend CloudWave's work to look into ways of strengthening data security and privacy in SaaS and PaaS environments. Along with this, as Cloud Computing becomes more popular, new uses for the Cloud (e.g. IoT), new trends in the industry (e.g. big data) and new programming models (e.g. Apache Spark, Google Cloud Dataflow) have begun to emerge.

To meet these changes, new approaches to design and implementation of cloud infrastructure - such as highly distributed clouds, high-performance lightweight virtualisation, computation-to-data affinity, complex in-memory data processing, etc., are required.

Projects major results:

The goal of CloudWave is to create a synergy between a Cloud and hosted applications such that both the cloud and the applications can dynamically change to create an optimal running environment. Better monitoring of both cloud and application performance while executing real workloads will allow coordinated adaptation both of the applications and of the cloud configuration.

CloudWave also aims to improve the feasibility of the emerging DevOps paradigm by giving developers comprehensive analysis of the behaviour of their application running in a cloud.
CloudWave is now completing its first year. We have defined the end-to-end reference architecture as well as low level component definitions. An initial end-to-end deployment of CloudWave along with use cases from industry are in their final stages.

Potential exploitation strategy

CloudWave's exploitation strategy is based around the belief that CloudWave's legacy will be in its scientific impact (through publications) and in its impact at the component level through commercialisation of its results and spinoff technologies.
The CloudMore use cases have been picked by its industrial partners to represent real business needs within their companies. Successful demonstration of the use cases serve to help drive CloudWave technologies both within the partner organisations, and as the basis for commercial offerings.

Links and references

www.cloudwave-fp7.eu

ClouT: Cloud of Things for empowering the citizen clout in smart cities

Name: Isabel Matranga

Organisation: Engineering Ingegneria Informatica SpA

Topics recommended for the 2016-2017 Work Programme

ClouT overarching objective is to provide enhanced solutions for smarter cities by using cloud computing to overcome some of the current challenges and limitations of the Internet of Things (IoT) domain.

Some of the advantages identified are related to cloud scalability, reliability and availability features, which give smart cities the opportunity to take advantage as much as possible of the data produced by billions of networked devices and millions of people.

Further research in the use of cloud computing in combination with the internet of things and/or big data technologies would give smart cities the opportunity to build new and enhanced services. Interest towards this topic was also expressed during the Japan-EU workshop (Tokyo - 31st July 2014), co-organised by OCEAN, ClouT and SUCRE projects.

Topics related to cloud elasticity, dynamicity and multi-tenancy are also considered of great interest to provide solutions for an enhanced management of resources, while further research in identity, authorisation federation and access delegation would give the opportunity to take advantage of cloud Federation models with gained scalability and flexibility.

Projects major results

To reach its objectives ClouT identifies, prototypes and validates the ClouT Reference Architecture (ClouT-RA) which sets a common ground of objects, definitions and rules, mapping the IoT and cloud advantages into a unique context.

As an EU-Japan co-funded project, ClouT sees a collaboration between six EU and seven Japanese organisations, with an overarching objective to shape a common baseline for future research in the IoT + cloud context.

To date, ClouT has defined its ClouT-RA leveraging work by established IoT and cloud research communities in Europe and Japan and has identified a set of reusable components for the implementation of the proposed reference architecture.

Thanks to the collaboration with four cities – Santander, Genova, Mitaka and Fujisawa - ClouT-RA is strongly linked to real-world examples and application. Early prototypes of its application in city contexts have been implemented and are being further developed.

Potential exploitation strategy

ClouT is implementing an exploitation strategy which sees the partners both in Europe and in Japan bring visibility to the project and transferring its results to target groups, end-users and stakeholders.
Exploitation activities are strongly supported by the collaboration with four cities two European and two Japanese, members of ClouT consortium: Santander, Genova, Mitaka and Fujisawa. By including both European and Japanese cities the project covers a wide portfolio of users and requirements associated to real user needs and city challenges in both regions. At the same time, ClouT ensures wider adoption of results thanks to their application in real city contexts.

An update since the last Concertation meeting (March 2014)
ClouT has established a fruitful collaboration with OCEAN and SUCRE projects. The collaboration has seen its practical application in the organisation of two EU-Japan workshops which took place on 16th May 2014 in Brussels and on 31st July 2014 in Tokyo.
Workshop participants learned about cloud computing research projects and had the chance to compare results from the EU and Japan, paving the way for future opportunities between the two regions. The workshop outcomes will contribute to discussions between the EU and the Japanese National Institute of Information and Communications Technology and the Japanese Ministry of Internal Affairs and Communications. These discussions aim to identify future cloud research topics for the upcoming calls with Japan under Horizon 2020.
In the ClouT project, cloud storage standards like CDMI (Cloud Data Management Interface) are being used. ClouT also follows SNIA (Storage Networking Industry Association) and is inspired by standard reference architectures defined by relevant international organisations like NIST (National Institute for Standards and Technology) and ETSI (European Telecommunications Standards Institute). The objective is of re-using existing standards and best practices, extending them when needed.

Links and references
http://clout-project.eu/

CoherentPaaS: A Coherent and Rich PaaS with a Common Programming Model

Name: Ricardo Jimenez-Peris
Organisation: Universidad Politecnica Madrid

Topics recommended for the 2016-2017 Work Programme

Data centre disaster tolerance: Providing cloud infrastructure that is able to survive disasters destroying or rendering unavailable whole data centres in order to provide highly available cloud infrastructure so businesses can survive despite this kind of disasters (as it happened recently with the hurricanes in US and the floods in NY).

Scalable cloud platforms for machine-to-machine and IoT applications: M2M and IoT applications are ideal to run in the cloud. However, they are pushing the needs for scalability by 1-2 orders of magnitude compared to current cloud applications what will require new cloud platforms able to address these scalability challenges.

Projects major results

Currently, companies are using a combination of cloud SQL and NoSQL technologies to deal with their data management needs. This is what is called polyglot persistence. However, polyglot persistence is bringing new problems.

First, there is no consistency across different data stores. Second, queries involving multiple data stores imply to program the query manually and optimize it manually as well. CoherentPaaS is solving the issues of polyglot persistence by first integrating an ultra-scalable transactional processing that provides holistic transactions across all data stores NoSQL (key-value data store, graph DB, document-oriented data store) and SQL. This provides full ACID transactions for updates across data stores and therefore, guaranteeing full data consistency in the advent of failures and concurrent accesses.

Secondly, CoherentPaaS is creating a query engine enabling to perform SQL queries to enable to correlate and aggregate data across data stores. This will enable for instance to perform joins between different data stores. In order to keep the power of the APIs/query languages of each data store, the SQL enables to embed sub-queries written in the APIs/query language of the underlying data stores.

In this way, CoherentPaaS gets the best of the two worlds, the power of local query languages/APIs and the ease of use of SQL.

Potential exploitation strategy

A startup, LeanXcale, has been created by two partners of the project for commercialising some of the core results of the project. There is a large fraction of the consortium partners that are data store providers that will integrate their data stores with the ultra-scalable transactional processing. The exploitation plan lies in signing bilateral commercial agreements between LeanXcale and each of these partners to commercialise the integrated systems.

The consortium also brings end user partners that are potential clients of the resulting platform. The startup is being constitute legally and will be created by end of September 2014.

An update since the last Concertation meeting (March 2014)

CoherentPaaS has partnered with the LeanBigData project also led by UPM on an ultra-scalable big data platform. These two projects will integrate their results what will result in a cloud data management platform providing not only an integrated data management platform for OLTP with SQL and NoSQL data stores, but also with OLAP capabilities to perform analytics on top of it.

Links and references

http://coherentpaas.eu

http://leanbigdata.eu

COMPOSE: Collaborative Open Market to Place Objects at your Service

Name: Benny Mandler

Organisation: IBM

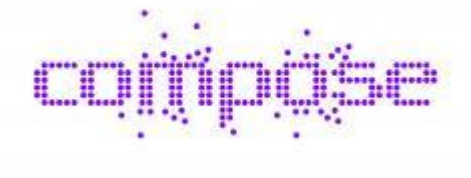

Topics recommended for the 2016-2017 Work Programme

- IoT - Back-end as a Service. Bringing the IoT from the consumer to the enterprise market
- IoT security / privacy (including data control / ownership / access control)
- Standardisation and interoperability in the smart objects platforms arena

Projects major results

Results focus mostly open source platform to ease the task of a developer creating applications based on internet connected smart objects. Secure platform covering the entire development life cycle.

COMPOSE also supports different stakeholders: application developers, smart object producers and owners, and end-users.

Currently prototypes of all technological aspects exist, and a first version of an integrated platform will be available in 4Q 2014.

Potential exploitation strategy

Key factors are:

- The identification of marketable products (ready to enter the market just after the project ends).
- Engage early adopters (potential customers) and/or investors and cooperate with them to create feasible business cases.
- No real need to create big (e.g. whole consortium) exploitation paths, but focus on establishing partnerships of a few interested partners on specific products/technologies.

The combination of these 3 factors will facilitate the creation of sustainable start-ups and spin-offs.

In addition to this, the identification of further funding opportunities (regional, national and European; public and private), as well as training on business, innovation and IPR subjects should be part of the exploitation work. This will help/stimulate partners and/or young researchers in creating business ventures, based on their project outcomes.

Finally, pilot activities are of paramount importance for exploitation. Having a compelling demo of the project outcomes, as well as concrete KPIs for measuring the actual impact of the proposed solutions is necessary for successful exploitation strategies.

An update since the last Concertation meeting (March 2014)

- Standardisation - held a first successful meeting of the Web of Things Community Group (W3C Workshop on the Web of Things - Enablers and services for an open Web of Devices, 25–26 June 2014, Berlin, Germany).
- Work with another EU project who maintains large amounts of smart objects data to expose it via COMPOSE (COMPOSE as a data hub).
- Work with another EU project on architectural aspects of mutual interest and providing them with access to COMPOSE pilots data.
- Work with another EU project on having them use COMPOSE run-time.

Links and references

http://www.compose-project.eu/

CUMULUS: Certification infrastrUcture for MUlti-LAyer cloUd Services

Name: George Spanoudakis
Organisation: CITY UNIVERSITY LONDON

Topics recommended for the 2016-2017 Work Programme:

Integrated security development environments supporting security, privacy and resilience-by-design
Autonomic security
Automated security certification models and continuously verifiable certification processes
Joint calls with third countries

Projects major results:

(a) New automated and evidenced based models and processes for certifying the security of cloud services.
(b) Test-, monitoring- and TPM-based tools to support the certification of security of cloud services.
(c) Processes and tools to support the use of cloud service security certificates in engineering systems using these services.
(d) Integrated security certification framework able to orchestrate or trigger the use under demand of the aforementioned processes
(e) (a)-(d) are expected to reach at least TRL 5 (according to NASA scale)

Potential exploitation strategy:

CUMULUS has identified an initial set of five potential customer segments of its outcomes (i.e., Cloud Service Providers (CSP), Certification Authorities (CA), Auditors (AU), Application Developers (AD) and Service Consumers (SC)) and is exploring a range of different value propositions (i.e., bundles of project outcomes) for each of these segments.
Project outcomes address the following segments of the ICT security markets: Trusted platforms, Cloud service security, Security Engineering, Security as a Service, and Security certification.

An update since the last Concertation meeting (March 2014):

CUMULUS has collaborated with SECCRIT and A4Cloud.

Links and references:

http://cumulus-project.eu/

ECO2Clouds: Experimental Awareness of CO2 in Federated Cloud Sourcing

Name: Pierluigi Plebani
Organisation: Politecnico di Milano

Topics recommended for the 2016-2017 Work Programme:

Based on the experience gained in this two-year project, we have realised that a set of models and tools to properly design clouds-based applications is crucial. Moreover, the definition of a set of common and standard metrics for measuring the impact of an application on the environment is also fundamental.

On this basis, we suggest the following topics:

- Definition of methods and tools for modelling eco-aware applications living in the Cloud.
- Definition of common metrics to measure the impact of the applications considering the sustainability as a mix of standpoints (energy consumption, energy sources, risk associated to an energy source, co2 emissions, ...).
- Definition of architectures able to monitor and adapt applications to make them more sustainable.

Projects major results:

The ECO2Clouds project develops techniques and mechanisms for CO2 aware deployment and management of workloads on federated Clouds. Relevant results of the ECO2Clouds project are:

- Definition of a layered set of metrics able to measure the impact on the environment of physical infrastructures, virtualisation layer, and applications in terms of CO2 emissions.
- Architecture for collecting the aforementioned metrics in a federated cloud environment.
- Optimised deployment strategies based on the CO2 emissions.
- Adaptive application modeling aiming to increase the energy awareness of the developers.

Potential exploitation strategy:

The project has produced a set of tools and methods that can be directly exploited by the community. Among the others we mention:

- A set of metrics to evaluate and measure the sustainability of a cloud federation as well as the CO2 emissions of the application running on it.
- Monitoring facilities to compute the eco-metrics.
- The eco-driven scheduler to optimise the deployment of virtual machines in a federated cloud.
- The ECO2Clouds project has also organised different conferences and workshop to foster the adoption of the project results in both the academic/research and industrial community.

An update since the last Concertation meeting (March 2014):
ECO2Clouds did not participate to the previous Concertation meeting

Links and references:
http://eco2clouds.eu

FELIX: Federated Test-beds for Large-scale Experiments

Name: Bartosz Belter

Organisation: Poznan Supercomputing and Networking Center

Topics recommended for the 2016-2017 Work Programme

FELIX is creating a common framework where users can request, monitor and manage slices of resources (computing, storage, networking, etc.) provisioned over distributed and distant Future Internet experimental facilities in Europe and Japan. Relevant topics emerged in FELIX research which require further investigation and community support are:

- Joint control and management of both network and computing resources, with more and more function virtualisation techniques to allow multi-tenancy on shared hardware resources, flexibility and highly dynamic network and data centre reconfigurations.
  Federation of resources from multiple facilities/owners, more and more geographically dispersed across the world, to increase the scale of the experimental facilities, spread the access to new technologies to society at large, optimise investments/costs of infrastructures.
- Enhanced monitoring of the virtualised resources (network + computing), to allow more fine grain control of the virtual resources and networks end-to-end.
- Solid International collaborations to foster the appropriate impact of Future Internet solutions and enablers beyond the borders.

Projects major results

FELIX is creating an integrated Europe-Japan SDN test-bed to:

- Utilise functionality of OpenFlow and NSI (Network Service Interface) for creation of Federated SDN Services.
- Introduce new APIs and logic for globally distributed heterogeneous SDN and IT islands.
- Enable interchange of resources information, share overall resource pools.
- Provide dynamic network interconnectivity between and within islands.

To this purpose, FELIX has designed and is implementing a control framework for the integration of different resources (e.g. transit network, data centre network, computing and storage, etc.) residing in multi-domain heterogeneous facilities.

The inter-domain dynamic connection provisioning mechanism used in FELIX to build inter-continental resource slices is based on the Network Service Interface specified by OGF, and currently under deployment by many R&E networks around the world.

Potential exploitation strategy

The main area of applicability for the FELIX results is the domain of federated infrastructures for Future Internet experimentation and research. Nevertheless, the FELIX control framework and its tools for orchestration, control and policy enforcement on top of virtualised resources have the potential to create a broader impact on multiple stakeholders.

The key FELIX assets on federation and NSI can impact the business of network operators, particularly those in the R&E business/area, who can benefit from these solutions to enhance their current service offer in production environments, and thus allow a better use of their network and computing infrastructures.

Also, the FELIX research can open business potentials for the technology providers (both SMEs and research centres) who can leverage on the developed foreground on federation and virtualisation to configure as solution experts towards other network operators and facility owners.

An update since the last Concertation meeting (March 2014)

This is the first position paper from the FELIX project. FELIX leverages on OFELIA and RISE projects, run in Europe and Japan, respectively. The project enhances the existing tools and infrastructures developed in these initiatives and deploy a new federation framework, targeting EU and JP research communities as potential users of the common infrastructure.

FELIX is closely cooperating with the GÉANT project, which in one of its activities addresses the problem of federation of network and compute infrastructures to create dynamic experimental facilities for GÉANT users. FELIX is currently investigating possibilities of federating European and Japanese testing facilities with other regions of the world, i.e. other Asia countries and the United States.

Links and references

http://www.ict-felix.eu/

GENiC – Globally optimised Energy Efficient Data Centres

Name: Elizabeth Massey

Organisation: United Technologies Research Centre

Topics recommended for the 2016-2017 Work Programme

- Data security compliance and data mobility on the cloud with respect to customer requirements (SLAs) and energy usage in optimised DC's.
- Energy optimised orchestration and workload provisioning in the cloud.
- Sensor (HW & SW) based prediction algorithms used for load prediction and cloud energy optimisation.

Projects major results

GENIC's main results so far are as follows:

- Cloud based energy optimisation services in geographically dispersed DCs. Genic will focus around finding the right balance between energy consumption at CSP (Cloud Service Provider) level while maintaining SLAs from customer workloads perspective.
- Highest available uptime. GENiC addresses the '99.99% uptime' cost problem from an energy consumption perspective while minimisiing impact on user SLAs.
- Holistic DC management platform. GENiC will develop a management platform with open interfaces and a common data format to provide control and optimisation functions and decision support tools to balance IT workload, power supply capacity and thermal cooling needs with energy consumption reduction requirements and customer SLAs.

Potential exploitation strategy

Results from the GENiC project will provide opportunities to expand portfolios for industry and stakeholders alike. The main benefit to the data centre sector will be achieved through the key technologies and services enabling integrated data centre-wide energy management providing business growth opportunities for the different value chain stakeholder including HVAC, IT and energy systems manufacturers.

Links and references

GENiC official website: Http://projectgenic.eu/

http://idpc.gov.mt/dbfile.aspx/WP196.pdf, article 29 Data protection working party, 01037/12/EN, WP 196, July 1, 2012.

http://www.consilium.europa.eu/uedocs/cms_data/docs/pressdata/en/ec/139197.pdf, EUCO 169/13, 24/25 October 2013.

HARNESS: Hardware- and Network-Enhanced Software Systems for Cloud Computing

Name: Alexander Wolf
Organisation: Imperial College London

Topics recommended for the 2016-2017 Work Programme

Support for multi-data centre cloud computing - Commercial or governmental entities will increasingly operate multiple data centres. They will do this for various reasons, including to:

- Lower network access latency.
- Improve fault resilience.
- Increase available resources.
- Adhere to jurisdictional regulations.

While multiple data centres are operated today by many vendors of cloud services, they are treated as independent enterprises for which there is little or no support for flexibly and dynamically deploying or executing applications across them. There is a need for fundamental research into how this can be accomplished, along the following lines:

- Application deployment and management: physical placement and run-time adaptation.
- Data management and dependability: consistency models that can trade resiliency against scalability.
- Policy-based regulatory compliance: multiple domains, flexible security, guaranteed disposal.
- Programming models and software engineering: annotations, architectures, tools, and environments.

We expect the following major project outcomes:

- The architectural principles that allow heterogeneous hardware and network resources, such as GPUs, FPGAs, SSDs, software routers, and the like, to be seamlessly incorporated into a cloud platform.
- The development of a data-centre infrastructure that manages the execution of applications on heterogeneous resources.
- The design of low-level management algorithms that optimise the allocation of different types of resources.
- The implementation of virtualisation technologies that allow traditionally single-tenant resources to be efficiently shared among multiple tenants.
- The design of programming models allowing applications to make effective use of HARNESS-managed clouds.

The following results have been achieved:

- A generic IaaS/PaaS architecture that integrates various prototypical heterogeneous resources.
- An extended version of an open-source PaaS implementation that mediates between applications and heterogeneous resources.
- The development of computation, network and storage resource management systems.
- The virtualisation of data-flow computation engines.
- The design of aspect-oriented and dataflow programming models for heterogeneous platforms.

Potential exploitation strategy

HARNESS is exploring largely new territory, and developing a wide range of novel technologies. Thus, the project does not plan to exploit its outcomes as a single entity (e.g. a start-up).

Instead, individual technology components have their own exploitation plan, falling into 3 main categories:

- Commercialisation as part of existing product offerings: Maxeler and SAP (the two HARNESS industrial partners) plan to evaluate commercialisation of HARNESS results within their products. In the Maxeler case, innovations developed in HARNESS are already part of 2014 product releases.
- Release as part of open source platforms: HARNESS is contributing to several open source projects, most especially ConPaaS and XtreemFS. Developed technology will add to the capabilities of these platforms.
- Publication and use for future teaching and research: Outcomes not immediately suitable for commercialisation will be disseminated as the basis for future research and teaching by the academic partners.

An update since the last Concertation meeting (March 2014)

The consortium has selected a collaboration strategy whereby collaboration with other projects is effected by those partners participating in both projects. This provides significant benefits from collaboration under minimal overhead, which is important given the wide range of HARNESS technology areas. Results from such collaborations include:

- Integration of HARNESS code within ConPaas and XtreemFS in conjunction with FP7 Contrail.
- Three patents filed by Maxeler as a result of joint HARNESS and SAVE project development.
- Collaboration with FP7 FASTER on one of the use case applications, where both Maxeler and Imperial are partners in both projects, providing enhanced use cases for study.

Moving into the third year, the consortium is also considering developing collaborations with earlier-stage projects that may be interested in building upon some of the HARNESS outcomes. One example is with the FP7 POLCA project (where Maxeler is a partner), which may wish to build upon programming tools developed by Imperial during HARNESS.

Links and references

HARNESS website: www.harness-project.eu

Open source projects:

- www.conpaas.eu
- www.xtreemfs.org

Notable collaborating/potentially collaborating projects:

- www.contrail-project.eu
- www.fp7-save.eu
- www.fp7-faster.eu
- www.polca-project.eu

HTML5Apps: Closing the gap between native and HTML5 apps through standardisation

Name: Daniel Dardailler
Organisation: W3C

Topics recommended for the 2016-2017 Work Programme

We think the following areas, which are also part of the EC Rolling Plan for ICT standardisation, should be covered in the upcoming program:

- Web accessibility.
- Mobile payment.
- Linked Open Data.
- Intelligent Transport Systems and Advanced Manufacturing.
- Internet/Web of Things for Smart Grids, Metering, and Cities.

Projects major results

The goal of the HTML5Apps project is to close the gap between native and HTML5 apps through the standardisation of missing HTML5 functionalities. Progress has been made in all areas presented in the DoW, in particular, new OS level APIs for HTML5 apps are being standardised in the W3C SysApps WG (which met f2f in China to resolve issues with the security and execution API models) and we ran a first Web Payments workshop with more than 100 participants. We also created an industry charter for the group, a Web site, etc.

Potential exploitation strategy

In terms of exploitation, W3C standards are successfully deployed because they have a large community of players involved from the very beginning of the process. With resources from this project, we have recruited representatives from 20 W3C Member organisations in the W3C Web and Mobile Interest Group, and attracted more than 130 subscribers to its public mailing list.
We also held 5 teleconferences, 1 face-to-face meeting, 10 task forces, 3 of which are led by the HTML5Apps project, and cover work linked to high-level analysis of the needs of HTML5apps.

An update since the last Concertation meeting (March 2014)

Since the last concertation meeting, and in addition to the above, we started our coordination for pre-standardisation coming from the Unit projects and updated our popular roadmap document on Standards for Web Applications on Mobile. We've also done a lot our outreach with the press and at conferences.

Links and references

Website: http://html5apps-project.eu

http://www.w3.org/Mobile/mobile-web-app-state/

http://html5appsproject.files.wordpress.com/2013/10/html5apps_factsheet.pdf

LEADS: Large-Scale Elastic Architecture for Data-as-a-Service

Name: Etienne Riviere

Organisation: Université de Neuchâtel

Topics recommended for the 2016-2017 Work Programme

The topics of interests for the LEADS project and its members can be classified along two main lines: research on novel cloud service models and research on novel cloud infrastructures. Research on cloud services should address the challenges and reluctances (e.g. as identified by ENISA) associated to the shift to cloud computing for companies. Cloud services should address concerns on security and privacy. This is a difficult problem especially when clouds support computations and applications logic and not only blind storage.

Research on privacy-preserving processing is required and due to the difficulty and importance of the task, should be given specific funding, oriented at fundamental research. It seems too early to consider privacy-preservation and security in clouds only as part of innovation and integration projects. The potential of exploitation of cloud services offering such guarantees is very important. It should not be neglected.

Research on novel cloud services should also consider the importance for companies to link existing infrastructures and software solutions, and to link existing data, to data and systems available in the cloud. Finally, research on cloud infrastructures should consider the opportunity of geographically distributed cloud services to integrate with the energy distribution network and with consumers and producers of data. This is complementary to making cloud elements energy-efficient, but has a higher potential for environmental footprint reduction.

Projects major results

The LEADS project works towards answering the demand of companies to exploit the wealth of public data available on today's Internet, to combine it with private company data and to apply business-specific processing on masses of historical and real-time data. Companies do not always have the will or capacity to exploit in-house the large computing facilities required for these operations.

The LEADS answer is to build a shared Data-as-a-Service platform running on an innovative infrastructure formed by a collection of geographically distributed micro-clouds. The project has contributed new techniques and software integrations for crawling, storing, querying and analysing in real-time publicly available Web data, and for operating the service on top of the complex, hierarchical and dynamic infrastructure formed by the micro-clouds.

Specificities of the infrastructure characteristics led to the addition of multi-site and hybrid fault and consistency models to the Infinispan open source data grid, partitioning techniques for the ZooKeeper coordination service, a fully distributed and locality-aware version of the Apache Nutch Web crawler, and query and data placement schedulers that take into account the nature and cost associated to each micro-clouds.

A query language supporting both declarative statements and real-time/streaming operations, with the ability to register specific business operations on data as part of (real-time) queries and interfaces to defined queries (based on Apatar) or to exploit results visually, were also developed and are currently integrated.
Finally, the project features an example business case and application provided by Adidas that combines the use of all components of the project and drives the evaluation of the platform over a collection of deployed micro-clouds.

Potential exploitation strategy

The Data-as-a-Service (DaaS) model itself has potential for long-term exploitation by being provided either by a collaboration between client companies or by a specific third-party commercial provider. The opportunity to establish such a provider company or to bootstrap new projects towards the exploitation of the DaaS principle is still being discussed in the project.
Individual components of the platform have a high potential for exploitation, or are already exploited. Many of the project outcomes are open source. The storage layer, which also supports the base querying capabilities is based on, and extends Infinispan, an open source product and the base for the JBoss data grid technology.
Several additions from the project are already exploited in this context, and further additions towards, e.g., the interaction with the Apache Hadoop software stack, are immediate exploitation assets. System support for decentralised cloud infrastructures have potential for exploitation for energy-aware and local cloud offerings that are blossoming thanks to high-speed networks and environmental considerations.

An update since the last Concertation meeting (March 2014)

The LEADS project is seeking to establish collaboration activities with the CloudSpaces FP7 project (http://cloudspaces.eu/) both technically and by means of joint events. CloudSpaces investigate innovative solutions for personal clouds and share technical challenges with LEADS in terms of service partitioning, elasticity and platform awareness.
LEADS is also in contact with the BigFoot project (http://bigfootproject.eu/) in this context.
LEADS organised the CloudDP 2014 workshop (http://clouddp14.unine.ch/) on its key scientific interests, data management and cloud infrastructures, at the EuroSys 2014 conference in Amsterdam in April 2014. The organisation of the workshop was a joint work with the LinkedDesign FP7 project (http://www.linkeddesign.eu/). LEADS will propose the organisation of the next edition of Cloud 2015 with the next edition of EuroSys, in April 2015 in Bordeaux.
We would like to involve the HARNESS FP7 project (http://www.harness-project.eu/) that works on new generation cloud computing platforms.

Links and references:

LEADS project Web site: http://www.leads-project.eu
Publications of the project: http://www.leads-project.eu/publications

MIDAS - Model and Inference Driven Automated testing of Services architectures

Name: Libero Maesano

Organisation: Simple Engineering France

Topics recommended for the 2016-2017 Work Programme

The 2016-2017 Work Programme should focus on:

- Radically new and native cloud applications as services.
- Test-driven Software Development Life Cycle (SDLC) methods and tools for these applications.

We mean by "radically new" those applications in all relevant sectors (Health, Logistics, Energy…) whose functionalities are made available firstly through simple and powerful service APIs. These applications cannot even be conceived without:

- "Radical" elasticity – massive upsize and downsize (e.g. ten thousand servers) in seconds and automatically.
- "Radical" pay per use - if you do nothing, you pay nothing, and costs are decreasing on quantity.
- "Radical" availability – 0.999999 – no service interruption for maintenance.
- "Radical" fault-tolerance – the unavailability of even critical data and algorithms is treated at the functional level.
- "Radical" security – no disruption, no leaking, extreme privacy of data and processes.

Test-driven SDLC tools must sustain the shift from "specify, design and implement components" towards "select, evaluate and integrate services through APIs". They must support:

- Both agile and contract-driven (model-driven) service design interweaved from the beginning with automated (model-based) test of large services architectures.
- Massive collaborative editing of very large scale services architectures models.
- Incremental continuous integration of services architectures.
- Continuous delivery without service interruption in large scale distributed architectures.

The aim of this recommendation is to foster the EU presence at the hearth of the digital service economy.

Projects major results

A SOA/API testing facility whose key innovation points are:

- Extreme automation of services architecture testing (grey-box testing of service compositions). Functional, security and usage-based testing. Model-based automation of all testing tasks - test case/oracle generation, test execution, arbitration, reporting, scheduling and planning. Model checking for test case/oracle generation. Probabilistic (Bayesian) reasoning for test planning and scheduling (failure seeking and troubleshooting).

- Testing “as a service” on cloud. Self-provisioning, elasticity of resource allocation, pay per use, security of access, integrity and privacy of data. On-line learning. Coping with multi-owner services architecture. User community tools.
- Programmable testing facility through APIs. Functionalities fully available through APIs on the Internet. Easy integration of MIDAS services by external modelling tools, IDEs, ALM tools, CAST tools, SOA governance tools, API management tools, for automated, background testing of services architectures within the Test Driven SDLC.
- Evolutionary registry/repository of test methods. Test researchers and practitioners can easily package their implemented test methods for uploading, registering and deploying - after certification - on the MIDAS registry/repository. The new registered test methods are proposed as such to the MIDAS users through the MIDAS APIs. Test method developer community tools.

Potential exploitation strategy

The MIDAS testing facility will be offered in the Software-as-a-Service modality, probably driven by a dedicated start-up integrating partners’ spin-offs. The service pricing will be “extreme” pay per use: the user opens an account with a credit card and if s/he does nothing, s/he pays nothing.

Research users – e.g. MIDAS other partners, European projects - will be billed for the mere cost of utilised underlying cloud resources. Industrial customers will be billed for the cost of utilised cloud resources plus a transparent margin.

Fair accounting tools and cost simulation tools - that allows estimating test campaign costs - will be supplied.

Lower costs of the cloud resources will be passed to prices.

A collaboration policy with researchers and practitioners for new enhanced test methods (e.g. with European projects) will be put in practice for an evolutionary service testing facility. The first target business sectors will be healthcare and logistics (experience within project pilots), with dedicated test methods and tools.

An update since the last Concertation meeting (March 2014)

At the end of the second year:

<u>Collaboration and clustering:</u>

- EC funded projects: Contrail (Open computing infrastructure for elastic services); Cassandra (Common Assessment and analysis of risk in global supply chains); e-Freight (European e-Freight capabilities for Co-modal transport).
- Italian Smart City initiatives: ADAPT (Accessible Data for Accessible Proto-Types in Social Sector) and (SMART AGING (Personal mobile services platform for data acquisition and processing for prolongation of active life, wellness, disease prevention and care in aging population everywhere).
- Proposed as testing framework in two proposals: FOLLOWER - Proposal submitted to the call PHC 26 of Horizon 2020, proposal submitted to the call for Tender in the framework of the FP7 DECIPHER PCP Project (<u>http://decipherpcp.eu</u>).

<u>Standardisation</u>

- Contribution to the OMG UML Testing Profile (UTP) 2.0 standard.

- Contribution to the TTCN-3 standard: ETSI standardisation project (STFs) on TTCN-3, e.g. (i) Extension of Conformance tests for TTCN-3 tools (09/13 – 05/14) and (ii) Conformance test suite for using XML schema with TTCN-3 (01/14 – 12/14).
- Contribution to the standardisation on ETSI Test Description Language (TDL): MTS Specialist Task Force (STF) for Test Description Language (TDL) – Phase 2.

Links and references

http://www.midas-project.eu/

MODAClouds: Model-Driven Approach for design and execution of applications on multiple Clouds

Name: Dana Petcu

Organisation: Institute e-Austria Timisoara

Topics recommended for the 2016-2017 Work Programme

- Model-driven engineering of Clouds; - Decision support systems for Cloud migration.
- Automated deployments and redeployments in Clouds.
- Data management and synchronisation between multiple Clouds.
- Combine automation with customisation of Cloud services.
- Define re-engineering process for Cloud migration.
- Establish standard metrics for quality of service.
- Introduce Modelling-as-a-Service.
- Portability and interoperability.
- Self-organisation in Federation and Multi-Clouds.
- Benchmarks for Clouds.
- Enhance the number of Cloud-related standard implementations.
- Establish standards for machine-readable representations of services, quality, and negotiations.

Projects major results

MOdel-Driven Approach for design and execution of applications on multiple Clouds (MODAClouds) provides methods, a decision support system, an open source Integrated Development Environment and a run-time environment for the high-level design, early prototyping, semi-automatic code generation, and automatic deployment of applications on Multi-Clouds, with guaranteed quality of service.

Model-driven development combined with model-driven risk analysis is introduced by the project to the Cloud computing application cycle. The project addresses the application developers that will be able to specify in a vendor agnostic manner the models of Cloud services in which they are interested, as well as to enrich these models with quality parameters.

In order to close the current gap between the design and run-time stages of an application, both performing quality predictions, as well as run-time monitoring and optimisation, are providing information to the design time environment.

Potential exploitation strategy

The project partners are committed to maintain and enhance, beyond the project end, the software engineering tools developed in the frame of the project. Moreover, the commercial partners are already integrating these tools in their products and services. For example, we mention here the triggers technology introduced by FLEXIANT in their platform as a service product and the ADONIS: cloud by BOC for business process management.

An update since the last Concertation meeting (March 2014)

The four proof-of-concept applications are currently providing the feedback on the usefulness of the tools that were developed.

The project team is working with PaaSage and ARTIST teams on the CloudML development and improvement following the feedback from use cases.

Moreover, we take part in various initiatives of other projects for common dissemination and training activities.

The Cloud application portability is a topic that is relevant for MODAClouds as targeting the easy deployment in Multi-Clouds. In this direction the team is enhancing the facilities provided by the middleware developed in the frame of earlier projects mOSAIC, Cloud4SOA and REMICS. A particular direction of improvement is related to data models, communication and synchronisation for and between multiple Clouds.

Links and references

Project web site: http://wwww.modaclouds.eu

Publications (public deliverables, research papers, training material, videos, white paper, interviews, presentations: http://www.modaclouds.eu/publications/

Software repositories, documentation, video with demos: http://www.modaclouds.eu/software/

MONDO: Scalable modelling and model management on the cloud

Name: Dimtris Kolovos
Organisation: University of York

MONDO
Scalable Modelling and Model Management on the Cloud

Topics recommended for the 2016-2017 Work Programme

- Hybrid modelling: Investigate how new and existing heterogeneous modelling notations (e.g. graphical, textual) can be seamlessly combined, with the involvement of domain experts, to address domain-specific modelling problems.
- Online modelling: Online Integrated Development Environments such as Eclipse Orion are gaining traction. Additional research is required to support the development of multi-user online modelling tools that can contribute software and systems modelling capabilities to such IDEs.

Projects major results

The aim of MONDO is to tackle the increasingly important challenge of scalability in MDE by developing the theoretical foundations and an open-source cloud-based platform for scalable modelling and model management.

Achieving scalability in modelling and MDE involves being able to construct large models and domain specific languages in a systematic manner, enabling teams of modellers to construct and refine large models in a collaborative manner, advancing the state-of-the-art in model querying and transformations tools so that they can cope with large models (of the scale of millions of model elements), and providing an infrastructure for efficient storage, indexing and retrieval of large models.

Potential exploitation strategy

Most of the project results and all of the core technologies will be made available as open source and the partners will be working with additional tool and platform providers beyond those in the project to encourage their adoption of the MONDO technologies for their own tools and model persistence technologies.

Both commercial and open source tool and platform providers will be targeted through a series of dissemination actions that will highlight the quantitative data collected from the industrial applications evaluations carried out in the project verifying the performance improvements, improved scalability and advances in multi-developer based modelling for larger teams.

The open source approach will encourage a wide range of tool and platform providers to exploit the results in order to address already existing market demands.

The project partners will complement and reinforce this approach with standardisation proposals and industry consensus actions to ensure interoperability of updated modelling tools and model persistence platforms.

An update since the last Concertation meeting (March 2014)

MONDO is currently at the end of its first year. At this stage, requirements have been collected, prioritised and assigned to technical work plan, and prototypes for scalable model persistence, domain-specific language development and model transformation and querying on the cloud are under development.

During the second year of the project, these prototypes will be fully developed and integrated so that they can be evaluated through the industrial use-cases of the project between months 24-30.

Links and references

Project website: www.mondo-project.org

Project start/end: November 2013 - April 2016.

OCEAN Project: Mapping the Open Cloud Collaborative R&D Projects Landscape

Name: Olivier Bouzereau
Organisation: OW2

Topics recommended for the 2016-2017 Work Programme

The following topics were proposed by the community at the organised by OCEAN workshops:

- Support for energy efficient and green Cloud-based IT
- Methods for cloud cost assessment
- Service description, comparability, metrics, and validation of service levels
- Big data and Cloud.
- Internet of Things and Cloud
- Security and data protection.
- High availability / fault tolerance for Cloud systems
- Interoperability to prevent data/functional/management lock-ins
- Open Source Software governance

Projects major results

The OCEAN Open Cloud Directory (www.ocdirectory.org) provides an online catalogue of active and recently completed Open Cloud research projects and their results.

Potential exploitation strategy

The OCEAN consortium will continue to develop the Open Cloud Directory as a centre of knowledge about collaborative R&D cloud projects and their FLOSS outcomes.

An update since the last Concertation meeting (March 2014)

With 137 cloud assets registered within 70 projects, the OCEAN Open Cloud Directory is still growing. The OCEAN team is working in collaboration with other projects to position their results in the OCEAN Cloud Interoperability Framework.

Links and references

OCEAN Project: http://www.ocean-project.eu
Open Cloud Directory: http://www.ocdirectory.org/

Open-DAI: Opening Data Architectures and Infrastructures

Name: Luca Gioppo
Organisation: CSI-Piemonte

Topics recommended for the 2016-2017 Work Programme

- Data integration, data quality and automatic cleaning of data - linking real time data with big data tools to mix the data services with the capability to do predictive analysis or complex correlation.
- More orientation towards the integration of "out of the box tools" that offer a rich feature set instead of single product that try to cover all the needs without reaching the needed completeness, integration of these components into a scalable cloud environment. The aim should be to foster standard protocols and cloud deployability in particular in the IaaS context where a more standard approach is obtainable.
- Funding evolution of existing tools and not the creation of new ones.

Projects major results

The level of integration of the different open source components allowed to deliver a complete and modular platform, wrapped with easy interfaces that allow both skilled technical users to benefit from "state of the art" software and non-technical ones to deploy new API services with little effort.

Potential exploitation strategy

The project will become part of the commercial offering of two partners of the consortium, merged into a more global service offering for Open Data. The factors that sustain the project are the ease of installation of the platform and the ease of use of the API production pipeline even in consideration of the richness of features that the overall components offer. The project managed to wrap the complexity of different open source components under an umbrella of integrated functions.

An update since the last Concertation meeting (March 2014)

The Open-DAI project is becoming the starting point for the CloudOpting project as a model in the deployment of the platform. The project integrated with the HOMER project for the indexing of the metadata of API services.

Links and references

http://www.open-dai.eu

OPENi: Open-Source, Web-Based, Framework for Integrating Applications with Cloud-based Services and Personal Cloudlets

Name: Sinead Quealy
Organisation: TSSG

Topics recommended for the 2016-2017 Work Programme

Our OPENi consortium partners are continuing to explore data and privacy in all its guises. Recognising the data generated, gathered and stored within city spaces, we have representatives on the Smart Cities Stakeholder Platform. This EU networking initiative will enable us to identify problems to which we can provide a data and privacy solution.

Projects major results

In the last year, the Telecommunications Software & Systems Group (TSSG) has successfully secured both national funding and industry investment in two follow up projects based on the OPENi platform and ethos of enabling users control of their data.

In the research and development space, our follow on project is working to develop a cloud based solution with a seamless user interface to enable individuals to unobtrusively store, control and realise the benefit from their personal data. It will be a personal data locker in the cloud (cloudlet), with a user interface that allows users to customise their preferences selecting what data they share and with whom they share the selected information.

Building brokerage technology, it will collect and organise data in a highly structured multi-dimensional format according to user permissions and market it as dynamically composed packages according to enterprise needs.

The project's core value proposition is the creation of a virtual marketplace that facilitates the trade of Volunteered Personal Information (VPI) and involves consumers as data producers and beneficiaries in that transaction.

Potential exploitation strategy

A key activity is continuing to demonstrate, through commercial and scientific research methods, the advantages for society of citizen centric data control and use. In parallel to exploring wider applications, we are working hard to apply, in commercial enterprise settings, applications based on our original OPENi platform.

In one enterprise application, the core technology being developed is a telematics personalisation and management platform for consumers, which allows users (drivers in this particular case) to have instant access to their driving performance profile in an easy to use mobile web app.

The web app will be supported by cutting edge data analytics software deployed on a backend platform – which can be used to provide data to consumers and stakeholders such as insurance companies alike.

An update since the last Concertation meeting (March 2014)

- The OPENi team created a whitepaper on our project's aims, objectives and influences. It can be found at: http://www.openi-ict.eu.
- H2020 participation - with contacts met and nurtured through CloudWATCH Hub March 2014, we submitted a proposal in H2020 ICT 7. One of the use cases of the proposal aims to apply data interoperability to local authorities for improved and targeted provision of services in the community.

Links and references

http://www.openi-ict.eu/

ORBIT: Business Continuity as a Service

Name: Dimosthenis Kyriazis
Organisation: National Technical University of Athens

Topics recommended for the 2016-2017 Work Programme

- **Topic 1:** Service runtime management exploring relationships between processing, data and access properties of services. The complete service lifecycle should be dynamically adapted according to these relationships to enable high-performance, ubiquitous, optimised, and cost-effective service provisioning.
- **Topic 2:** Composite service optimisation in large-scale deployments. Dynamic interaction protocols between atomic services in composite cases, taking into consideration data and service parameters. Orchestration and deployment dynamic patterns in large-scale scenarios.
- **Topic 3:** Adaptive application programmable models. Service engineering and provision of applications on top of future internet infrastructures, which encompass cloud and IoT resources, addressing challenges related to heterogeneity and availability of resources as well as the high rates of data generation, processing and distribution.

Projects major results

ORBIT aims at ensuring business continuity by introducing a new cost-effective approach in virtualised infrastructures for providing application-agnostic high availability. Key to this direction is the consolidation of virtualised memory and I/O resources emerging from multiple physical hosts.
To this end, the following main results are envisioned:

- Highly Available Consolidation of Virtualised Resources, building atop the concept of server virtualisation to enable guest VMs to consume remote memory and I/O resources in a consolidated manner.
- Application Transparent Virtual Machine Fault Tolerance, improving the current solutions, which are either application specific, limited to Uni Processor workloads, lack required performance targets, or require propriety hardware, by providing a software-only solution that can be widely deployed on commodity hardware.
- Metro-Area Zero Downtime Disaster Recovery, to improve business continuity by geographically distributed over a Metropolitan Area Network, whilst maintaining the desired KPIs of instantaneous fail-over with near-zero downtime, supporting business continuity even in light of major faults downing an entire site.

Potential exploitation strategy

All project outcomes will be released as open-source. The corresponding tree has been created on GitHub as a placeholder for all project up-to-date results (https://github.com/orbitfp7/ORBIT). Libvirt, QEMU, and OpenStack extensions will be released.

An update since the last Concertation meeting (March 2014)

- Communication with QEMU community through distribution and RFC of initial project developments (i.e. post-copy live migration, and kernel user fault).
- Follow-up of relevant standards related to CIM profiles, Virtualisation Management (VMAN), and Open Cloud Computing Interface (OCCI).

Links and references

http://www.orbitproject.eu
https://github.com/orbitfp7/ORBIT

OSSMETER: Automated Measurement and Analysis of Open Source Software

Name: Dimitris Kolovos

Organisation: University of York

Topics recommended for the 2016-2017 Work Programme

- Analysis of social aspects of open-source software development: Understand the routes through which developers and users interact with open source projects. Investigate mechanisms through which users of open source software can be encouraged to engage with the developers of open-source software (often, developers of open-source software receive very little feedback about successful use of their software in industrial applications).
- Domain-specific open-source software analytics: Through our interaction with a wide range of open source projects in OSSMETER, we have identified that software targeting particular application domains can benefit from diverse types of analysis (e.g. responsiveness of OSS projects to changes in related standards is essential for some domains and of little importance for others).

Projects major results

The OSSMETER consortium has designed and implemented a modular and scalable platform for multi-aspect analysis and measurement of open-source software projects.

The OSSMETER platform can analyse source code, bug reports, newsgroup and forum threads, stored in a variety of technical infrastructures (e.g. Subversion, Git, NNTP, Bugzilla) and enables users to discover and compare different open-source projects using a wide range of criteria using intuitive visual metaphors such as bar charts, radar charts and timelines.

Potential exploitation strategy

The results from the OSSMETER project will be made available as open source and substantial actions for dissemination will create awareness of the new and more comprehensive capabilities to analyse open source technologies including the evolution and level of activities amongst the contributing and participating communities that surround an open source technology.

The project partners anticipate that a combination of public (not-for-profit) services will be established that exploit the OSSMETER analysis platform both by project partners and third parties, as well as commercial services providers that may potentially focus on specific open source communities such as Eclipse, or vertical market segments such as Aerospace.

In addition to making the project results available in open source, the partners will also publish specifications and propose new or revised standards that to allow others to develop interoperable value-adding components on top of the OSSMETER platform and to provide a consistent set of measures and analysis methods that ensures results from different instantiations can be compared and relied upon by European software developers.

An update since the last Concertation meeting (March 2014)

Since the last concertation meeting, the OSSMETER partners have fully developed the backend and frontend of the OSSMETER platform and plan to deliver it to the use-case providers for evaluation at the end of September 2014.

Links and references

Project website: www.ossmeter.com

Project start/end: October 2012 - March 2015.

PaaSage: Model-based Cloud Platform Upperware

Name: philippe massonet
Organisation: CETIC

Topics recommended for the 2016-2017 Work Programme

The following research topics have been identified during the PaaSage project, but will not be addressed within the project. The research topics could very well be addressed in further research projects.

- Optimised Cloud data partitioning and replication: The PaaSage platform automates the deployment of applications based on a deployment model. A key decision in a cloud deployment is how to partition and deploy the data between different cloud providers. PaaSage provides no decision support for partitioning and replicating data.
- Optimising data partitioning and replication in the cloud is a challenging research topic. When replicating data across different cloud providers it is important that data consistency be taken into account.
- Industry standards for publishing monitoring data: The PaaSage platform supports multi cloud scenario where applications can be deployed on different clouds. However to be able to compare monitoring data across clouds, industry standards for monitored data should be available.
- Guidelines for developing cloud enabled applications: Porting legacy applications, e.g. client server, to the cloud is not straightforward. Guidelines to help developers develop cloud ready applications are needed.
- Security service level objectives: Security metrics are not widely available from cloud providers, so it is currently difficult to monitor cloud provider security. More research on security level objectives in the form of industry standards is needed.

Projects major results

The first version of the PaaSage platform provides the following features:

- Model based cloud deployment: Models are central to the PaaSage platform. The PaaSage platform takes as input a deployment model of the application to be deployed in the cloud. The deployment model is independent of any specific cloud platform.
- Optimised Cloud deployment per component: the PaaSage platform uses models of cloud providers to find the optimal cloud deployment on a selected set of cloud providers. The deployment decision is taken for each component individually. The resulting deployment may involve a single cloud provider or may involve several cloud providers.

- Cross cloud monitoring and adaptation: service level objectives are monitored by the PaaSage platform and can trigger adaptations to the deployment in order to respect agreed service level objectives.
- Secure and privacy aware cross cloud deployments: the deployment models include security models of cloud providers. The security model identifies the security controls and security level objectives that are supported by cloud providers. They are taken into account as constraints when calculating the optimal deployment.

Potential exploitation strategy

PaaSage aims to be a sustainable platform over time and to involve users that share interest in its continued evolution. To this end, all the software and tools of PaaSage are being delivered as open source.

The industrial partners of the PaaSage consortium are providing requirements based on their case studies. These requirements are driving the development of the PaaSage deployment model language and the platform.

The case studies cover different application domains in order to ensure that the PaaSage platform applies across different application domains. The industrial partners of the PaaSage consortium represent these different application domains. They are the initial exploitation partners of the PaaSage platform.

An update since the last Concertation meeting (March 2014)

Interoperability and portability across cloud providers are key objectives of the PaaSage platform.

The PaaSage platform supports a multi-cloud scenario, where deployed applications can be undeployed and redeployed from one cloud to another by the PaaSage platform. PaaSage deployment models are defined in a provider independent modelling language.

The same deployment model can thus be used by the PaaSage platform to deploy on different cloud providers.

Internally the PaaSage platform transforms the cloud provider independent deployment model into a cloud provider specific model. It uses the Cloudify library (http://getcloudify.org/) to create cloud specific deployments.

For the security model, PaaSage is using industry standards on security controls from the Cloud Security Alliance (https://cloudsecurityalliance.org/research/ccm/) so that cloud provider security can be matched with security controls that are required by an application.

Links and references

PaaSage Web site: http://www.paasage.eu

PaaSage use case descriptions: http://www.paasage.eu/success-stories

PaaSage deliverables: http://www.paasage.eu/documentation

PANACEA: Proactive autonomic management of cloud resources

Name: Dimiter Avresky
Organisation: IRIANC

Topics recommended for the 2016-2017 Work Programme

Panacea project creates Machine Learning Framework (MLF) for predicting anomalies in the cloud and a proactive autonomic management of cloud resources. In this research field, there are several issues that can be taken into consideration for future work programmes.

PANACEA will allow users several advanced possibilities, based on the ML framework and the autonomic principles:

- Proactive autonomic management of cloud resources.
- Proactive software migration within one cloud.
- Promising users anytime, anywhere access to their programs and data.
- Developing new applications at the developer site.
- Creating mission oriented distributed clouds with autonomic self* properties.
- For owners of existing applications, with SLA requirements, efficient use of cloud resources will be ensured.
- For owners of deployed applications to monitor their executions, by controlling and pro-actively managing them (VM migrations, proactive rejuvenation, predicting the time to crash).
- Utilisation of Software Defined Networks (SDN) networks for implementing autonomic self* properties of Internet services and applications.

Projects major results

Up to now PANACEA generated the following outcomes:

- Machine learning specification and training Version: v.1.0 for creating ML-based Prediction Framework for Internet services and applications anomalies. It is composed of several major steps: Initial Training Data Sets Collection, Lasso Regularisation Process, Training Process, and Prototype Implementation of the ML-based Prediction Framework, Impact of the Utilisation of ML Framework on applications availability, Availability enhancement by proactive machine rejuvenation. Fort coming steps: Resilient VM based topologies, dynamic reconfiguration, and experimental environment for validation.
- First prototype of experimental environment and software is implemented: TPC-W user, VM client (Apache Tomcat, TPC-W, MySQL, Feature Monitor client), VM Server (Feature Monitor Server, Online Monitor and Decision module, Machine Learning Framework). First

experimental results are obtained: server response time as a function of the number of users from 16 to 128. The ML framework is validated in run time. The first prototype complete realisation is expected in the middle of 2015.

Potential exploitation strategy

Panacea Use cases selection is made based on several requirements, including impact on the society, its relevance and potential exploitation in the nearest future.

Two Use cases have been accepted for implementation:

- Cloud Web Hosting with a ML framework for a proactive cloud management With this use case, we could validate that we are able correctly to predict the remaining time to crash and response time of web servers, even when they are running in a quite dynamic environment and under heavy workloads. In addition, this use case will be invaluable in demonstrating that we can significantly increase the availability of web applications by means of software rejuvenation.
- Second Use Case: Data Analytics as a service Data Analytics as a Service (DAaaS). The data analytics is based on Apache Hadoop7. It can use Map – Reduce algorithms written in Java. DAaaS represents an ideal application to be executed in a Cloud environment.

An update since the last Concertation meeting (March 2014)

Discussions for preliminary Cross collaboration had been made during the Concentration Meeting (March 2014), Brussels with the following projects:

- Portability and Interoperability in Clouds: contributions from the mOSAIC Project - http://www.mosaic-cloud.eu; CACTOS – Context-aware cloud typology optimisation and simulation Project – Leader University of ULM.

In addition, positive discussions for the future collaboration have been held with the leader of the project “Okeanos IAAS” - https://okeanos.grnet.gr/home/. This is GRNET's cloud service, for the Greek Research and Academic Community.

Links and references

Panacea website: http: www.panacea-cloud.eu

Project start/end: October 2013 - March 2016

PROSE: Promoting Open Source in European Projects

Name: Alfredo Matos
Organisation: Caixa Mágica Software

Topics recommended for the 2016-2017 Work Programme

We believe that Free/Libre and Open Source Software (FLOSS) is and will be one of the primary vehicles to convey (re)usable project results, and a key driver for R&D activities. Therefore, exploitation of project results through open source projects is key. With the right infrastructure and methodology, it is possible to increase the value and sustainability of results stemming from R&D projects as open source software.

Building on the necessary infrastructure, software and project metrics will be the main driver for innovation, ensuring measurable results and contributions that promote software sustainability. Pragmatic and tangible success metrics create the robust framework upon which to evaluate and promote project outcomes going into H2020.

These two key vectors can only occur through distributed and collaborative cloud solutions that can create synergies across projects. Reusing, sharing and building upon the same tools will allow easier collaboration, and increase the effectiveness of the distributed R&D funds.

Projects major results

The major project results that directly contributed to the promotion of open source software focus on methodology and training, through deliverables, webinars and dissemination, all brought together under a common platform: Open Source Projects EU (OSP) (http://opensourceprojects.eu).

Over the last few months PROSE has continuously worked to promote the OSP platform, which is one of the highlights coming from the PROSE Coordination Action. To achieve this, PROSE participated and organised several events, as means to disseminate and promote OSP, but also to raise the awareness around the value of FLOSS as a key innovation drive.

To complement the Open Source Projects EU platform, we have continued to expand and grow on the already existing achievements, by integrating meaningful metrics for hosted projects, which can be used to track progress and visibility, as well as the integration of FLOSS procurement methodologies and tools, that allow projects to undergo self-evaluation as means towards better understand the potential commercial impact of their FLOSS results and contributions.

These results are all being integrated into http://opensourceprojects.eu, which is also being worked on for increase reliability and usability.

Potential exploitation strategy

The most important factor for the longevity of the PROSE results focus on the adoption of the Open Source Projects platform as a key collaboration tool going into the new research cycles. This can be achieved through the sustainability models that are being proposed inside the project, which include

but are not limited to the creation of foundation or entity to continue the drive of the open source projects platform, or a support action that can build upon and grow the important results achieved in the PROSE coordination action.

An update since the last Concertation meeting (March 2014)

Since the last Concertation meeting PROSE has been participating and organising events towards the promoting of Open Source, such as the PROSE Webinars on FLOSS, Solutions Linux, or the upcoming LinuxCon Europe.

- On the Open Source Projects platform, PROSE has been steadily growing the platform's adoption, and providing more tools that can directly create value.
- Over the past year, PROSE focused on providing meaningful metrics for hosted projects, so they can track their own progress and visibility, as well as the integration of FLOSS procurement methodologies and tools, so that project can undergo self-evaluation to understand the future direction of their contributions.

These achievements strengthen the PROSE position towards promoting open source software, and provide a clear indication of the value that FLOSS can have for future iterations and Work Programmes.

Links and references

http://ict-prose.eu
http://opensourceprojects.eu

Prowess: Property-based testing for web-services

Name: John Derrick
Organisation: University of Sheffield

Topics recommended for the 2016-2017 Work Programme

- Testing for complex systems and big data Reverse engineering and refactoring testing and monitoring.

Projects major results

The Prowess project aims to provide the European software industry with efficient and effective testing tools and techniques for web services and internet applications. We address the challenge to reduce time spent on testing, whilst increasing software quality, in order to quickly launch new, or enhancements of existing, web services and internet applications.

To address this, we have been developing methods and techniques to improve property based testing (PBT). PBT provides a powerful, high-level, approach to testing; rather than focusing on individual test cases to encapsulate the behaviour of a system, in PBT, this behaviour is specified by properties, expressed in a logical form. The system is then tested by checking whether it has the required properties for randomly generated data, which may be inputs to functions, sequences of API calls, or other representations of test cases.

A number of tools have been developed within the project to further these aims, these include: QuickCheck, a tool for property-based testing, Wrangler, a tool for refactoring, WSToolkit, a testing toolkit for web services, Smother, an MC/DC code coverage tool, WEDL-DSL a DSL for automatic test data generation, and many more.

Potential exploitation strategy

All technology developed during this project are tested in practice in a number of industrial pilots that the academic partners perform together with our industrial partners and their customers. This evaluation feed backs into the project itself, and is used to further improve and extend the technology developed.

Prowess results allow the partners to increase their business attracting new customers on the web services and internet sector in general.

- The exposure to new companies is already achieved through Quviq consultancy on Volvo IT, Ericsson and Basho.
- Erlang Solutions is approaching potential customers outside of their usual Erlang users' network. Our academic partners are making research widely known and increasing the user base.

An update since the last Concertation meeting (March 2014)

Cross collaboration has started with several EU and national projects. These include the EU projects:

- FITTEST, MIDAS, IoT.est, Broker@Cloud, SafeCer and RELEASE.

We have shared speakers and technical results, and plan a joint mini-workshop with members of the Broker@Cloud consortium on common aspects of the testing of web and cloud services.

Links and references

www.prowessproject.eu

RISCOSS: Risks and Costs in Open Source Software Adoption

Name: Angelo Susi
Organisation: Fondazione Bruno Kessler (FBK) - Italy

MANAGING RISK AND COST IN OPEN SOURCE SOFTWARE PROJECTS

Topics recommended for the 2016-2017 Work Programme

relevant topics for the WP2016-2017 have emerged in recent years in the work related to risks analysis in OSS ecosystems where communities, companies and public administrations coexist and exchange values, knowledge and the possibility to share and collaboratively identify and treat risks:

- Risk analysis in the context of cloud in particular for the activities of data management, use of services and their inclusion in other services.
- Identification and analysis of privacy issues in the cloud especially in the context of public administration where the trade-off between the complete interoperability of the public services and the need of assuring the privacy of the citizens is becoming more and more crucial.
- New risk analysis and treatment techniques that can be used in contexts of distributed and collaborative environments that are emerging thanks to the cloud.

Projects major results

RISCOSS develops a risk management methodology to facilitate the adoption of open source code into mainstream products and services [1].

The methodology is supported by a software platform that integrates the whole decision-making chain. The project has delivered a number of methods and techniques as described for example in [2, [3].

We focused on the development of the methodology based on the modelling and analysis of OSS-based ecosystems statistical and logic based assessment and measurement techniques for the management of risk adapted to the specifics of open source software.

These techniques also include ontologies for OSS ecosystems and risks, patterns for modelling OSS ecosystems according to the business strategies adopted, risk-reasoning techniques based on risk models and goal-oriented models.

Finally, they have been integrated into the prototype of the platform that will be exploited for the evaluation of the proposed methods in contexts such as companies, OSS communities and public sector.

Potential exploitation strategy

The increasing adoption of OSS components calls for decision-support practices, platforms or on-line services, which help users understand the kind of risk underlying their choice.

All use-case partners, TEI, KPA, XWiki, Moodbile, Cenatic and OW2 have plans to incorporate RISCOSS into the conduct of their businesses. For example, TEI intend to introduce the platform into its software development process.

- Cenatic will test RISCOSS in enabling the dissemination of open source solutions in the public sector.
- XWiki and Moodbile.org open source projects will leverage RISCOSS to enhance the consistency of their user- and community-driven feature roadmaps, and put the focus on reliability, stability and support for backward compatibility.
- OW2 will leverage RISCOSS to complement its SQuAT (Software Quality Assurance and Trustworthiness) quality program.

Moreover, RISCOSS is in the phase of disseminating the product into some OSS communities for the validation of the platform and for the identification of exploitation opportunities.

An update since the last Concertation meeting (March 2014)

The project is in contact with OSSmeter project that is developing OSS communities and projects data gathering and analysis tools. This aspect is important for RISCOSS since these data can be the bases of the risk management approach proposed by our project. The objective here is to make it deep the interaction while the two projects evaluate their prototypes.

On the research side, several new publications have been accepted since the last meeting, some of them are listed in the reference section (for an exhaustive list you can visit the RISCOSS website http://www.riscoss.eu).

These publications mainly concern the description of the methodology and of the techniques we propose for the problem of risk modelling and analysis via the statistical and logic based methods.

In parallel, the first prototype of the platform, integrating some of the risk analysis techniques, has been released, ready to be evaluated in the next months.

Links and references

RISCOSS website is http://www.riscoss.eu/.

Project start/end date: November 2012 - October 2015

Publications

- Xavier Franch, Angelo Susi, Maria Carmela Annosi, Claudia P. Ayala, Ruediger Glott, Daniel Gross, Ron Kenett, Fabio Mancinelli, Pop Ramsamy, Cedric Thomas, David Ameller, Stijn Bannier, Nili Bergida, Yehuda Blumenfeld, Olivier Bouzereau, Dolors Costal, Manuel Dominguez, Kirsten Haaland, Lidia López, Mirko Morandini, Alberto Siena: Managing Risk in Open Source Software Adoption. ICSOFT 2013: 258-264
- Xavier Franch, Ron Kenett, Fabio Mancinelli, Angelo Susi, David Ameller, Ron Ben-Jacob, Alberto Siena. “A Layered Approach to Managing Risks in OSS Projects”. In Proceedings of OSS 2014.
- Ron S. Kenett, Xavier Franch, Angelo Susi, Nikolas Galanis, “Adoption of Free Libre Open Source Software (FLOSS): A Risk Management Perspective” In Proceedings of COMPSAC 2014.

S-CASE: Scaffolding Scalable Software Services

Name: Kyriakos Chatzidimitriou
Organisation: Aristotle University of Thessaloniki

Topics recommended for the 2016-2017 Work Programme

The main goal of S-CASE is to provide a cloud-based realm of services and tools for software developers to enable rapid software prototyping based on user requirements and system models, provided in multimodal formats.

S-CASE will allow the extraction of system specifications and low-level architecture, as well as the discovery and synthesis of composite workflows of software artefacts from distributed open source and proprietary resources.

In this context, based on the current work programme and probing on the 2016-17 work programme, the S-CASE partners are focusing their discussions against five main axes:

- Boosting public sector productivity and innovation through cloud computing services
- Collective Awareness Platforms for Sustainability and Social Innovation
- Big data
- Research
- Internet of Things and Platforms for Connected Smart Objects - User-driven web

Projects major results

The most important result of the project is the methodology defined and the engineering framework developed to (semi)-automatically generate RESTful web services out of multimodal requirements. The requirements may come in the form of textual requirement in semi-structured documents, UML diagrams and storyboards.

The developed methodology employs natural language processing (NLP) and automated software engineering (ASE) primitives in order to transform the provided information into a Computation Independent Model (CIM), which is further elaborated to a Platform independent Model (PIM), a Platform Specific Model (PSM), and finally, software code.

Up to this point the ASE and NLP methodology have been defined, while two developer tools have also been released: a tool for importing, annotating and incorporating 3rd party web services into S-CASE services, and a tool for collecting and (semi-)automatically annotating textual requirements from software projects. Additionally, a search engine for software assets has been presented.

Potential exploitation strategy

The project is still at its first year, thus a discussion on setting up start-ups and/or spin-offs would be too premature. However, all partners have identified a number of potential exploitable items, which will be periodically refined and revised throughout the course of the project.

Additionally, a set of exploitation KPIs will be defined by the end of year 1 (October 2014). The main exploitable items that have been identified, apart from the S-CASE platform as a whole, which aims to accelerate the overall agile software prototyping process, are the Automatic Semantic Annotation and developer tools, the Open Source Resource Finder and the Question Answering Mechanism. Next we plan on conducting a market research on the S-CASE niche, i.e. multi-modal requirements, model driven engineering and restful web-services.
From this analysis we plan on identifying key potential industrial actors, S-CASE customers and users.

An update since the last Concertation meeting (March 2014)
Being early in the project, we have probed for future collaborations and for adopting work already done in other projects.
Specifically, we have closely investigated work done in projects MARKOS, 4CaaSt, and MODAClouds and consider that the MARKOS project could potentially aid us in the development of the open source resource finder. We have established communication and have started working with the MARKOS API. This approach is expected to lead to tight collaboration.
Recently we started collaboration with project STORM CLOUDS, aiming to exchange know-how and expertise. S-CASE members have been participating in meetings and providing feedback whenever requested.
With respect to standards contribution, we have already decided to adopt the WSDL and WADL standards for describing S-CASE produced and consumed web services and we are aiming to contribute in that direction. Additionally, we are also working towards standardising work on UML profiles for RESTful application development.

Links and references
Website: http://www.scasefp7.eu
Twitter: http://twitter.com/scasefp7
Facebook: bit.ly/SCasefb
LinkedIn: bit.ly/SCasegrp

SeaClouds: Seamless adaptive multi-cloud management of service-based applications

Name: Francesco D'Andria
Organisation: ATOS Spain SA

Topics recommended for the 2016-2017 Work Programme

With cloud computing representing a major shift in the way information technology is used it is of very high importance to ensure that developers, app providers but the IT industry in general does not remain dominated by monopolies.

To support competition in the cloud computing market there should be different providers to choose from for any kind of application. Developers developing applications should be able to choose between different cloud services and should also be able to switch between cloud services of different providers whenever needed.

SeaClouds directly impacts on the way developers are going to build cloud apps without worrying about underlying execution of different PaaS or IaaS providers, relying on its service orchestration capabilities based on informed election among of the provider. In this regard, SeaClouds will allow developers and development companies take a much shorter and consolidated approach over different underlying technologies and stacks (proprietary and open sources).

At the same time, it offers application monitoring and adaptation based on different factors, cloud infrastructure and application behaviour, from design time, easing developers from this difficult task.

On this scenario, SeaClouds project can be seen as an enabler step towards fostering a much better and ready ecosystem for developers, especially those who do not have experience on cloud, or are biased by a particular technology stack.

By allowing developers to monitor and audit underlying PaaS and IaaS execution of their cloud applications, SeaClouds will contribute to the market by creating the ground base means for developers, from design time, simplify application monitoring and independently of the underlying vendor.

This could benefit applications developers or services provider to create self-adaptive applications that comply to more strict requirements and SLAs from their customers, and also they will be ready to offer to their clients and users possibilities of SLA definition and compliance, which would clearly affect how providers and clients will establish commercial relationships. Therefore, this capability will imply the possibility of having clearer and transparent contract relations from cloud providers and applications customers.

Projects major results

The SeaClouds approach is based on the concept of service orchestration and it is designed to fulfil functional and non-functional properties over the whole application. Applications will be dynamically

reconfigured by changing the orchestration of the services when the monitoring detects that such properties are not respected.

SeaClouds main result is the implementation of a novel platform which performs a seamless adaptive multi-cloud management of service-based applications. More specifically:

- Orchestration and Reconfiguration Mechanism: implementing application lifecycle management capability to dynamically deploy, migrate, replicate, and distribute modules compose applications among multiple Clouds, while checking both QoS violations and dynamic changes in the offer of the providers and the current demand. A reconfiguration process capable of preserving the soundness of the orchestration, by performing life-cycle management actions when required, by means of a unified management API, which will be a reference for major international efforts of standardisation such as CAMP and TOSCA.
- Application Monitoring Service: A range of Standardised and Unified metrics of different types (low level, container level, app level, etc) based on disparate underlying cloud providers that will allow the runtime monitoring of deployed services so as to assure the end-to-end QoS of the complex application, regardless of how it is distributed across different PaaS.
- Based on Standards by extending and incorporating CAMP into SeaClouds, SeaClouds covers all future CAMP-compliant providers or tools, allowing application developers to manage applications hosted on multiple Clouds environments. Application packaging will be implemented using the TOSCA specification for multi-cloud applications, and deployed being CAMP-compliant.

Potential exploitation strategy

SeaClouds impacts on the way developers build cloud apps without worrying about underlying execution of different PaaS or IaaS providers, relying on its service orchestration capabilities based on informed election among of the provider. SeaClouds will allow developers and ISVs take a much shorter and consolidated approach over different underlying technologies and stacks. It offers application monitoring and adaptation based on different factors, cloud infrastructure and application behaviour, from design time, easing developers from this difficult task.

Two exploitation approaches have been considered to ensure the sustainability of the SeaClouds solution:

- Open Source a Commercial Entity providing services - SeaClouds consortium already agreed at releasing the SeaClouds framework in the Cloud market under an Open Source approach, allowing its use by the developer community, and promoting it in an effort to create a reference framework that will contribute to the development of cloud applications with the main offering of eliminating vendor lock-in at and performing application monitoring PaaS and IaaS level.
- Commercial Entity providing services - The consortium will evaluate the option to create a new business entity by interested partners of the consortium that will evolve SeaClouds and test different business models around this scenario with the intention to obtain a sustainable business by selling services or technology.

An update since the last Concertation meeting (March 2014)

- Collaboration activities have been started with the other related FP7 projects. In particular, collaboration meetings have been held with Cloud4SOA, MODAClouds, Artist, and PaaSage.

- The collaboration with industry is also progressing through the close connection with Brooklyn Project.
- SeaClouds is also establishing a mutually beneficial relationship with the standard bodies. Since it will be developed according to emerging standards, SeaClouds will get constant feedback from their users and, at the same time, will contribute to their dissemination and evolution. The main standards on which SeaClouds will rely are the OASIS CAMP (Cloud Application Management for Platforms) and TOSCA Topology and Orchestration Specification for Cloud Applications).

Links and references

http://www.seaclouds-project.eu/
https://www.facebook.com/seacloudsproject
https://twitter.com/SeaClouds_EU
http://www.linkedin.com/groups/SEACLOUDS-PROJECT-7449431

SERSCIS: Semantically Enhanced Resilient and Secure Critical Infrastructure Services

Name: Mike Surridge
Organisation: University of Southampton IT Innovation Centre

Topics recommended for the 2016-2017 Work Programme

Cloud infrastructure enables agile, definition, composition and deployment of IT resources. These are used to meet in-house IT requirements by dynamic scaling of IT resources, and increasingly to create shared IT resources to support inter-enterprise applications, including social networking assets and Big Data systems.

- Conventional information security risk analysis approaches are not well suited to managing risks in such dynamic and evolving systems.
- Real-time threat analysis, detection and mitigation methods are needed, that can be applied during the design or operation of cloud-based virtualised infrastructure and services.

Projects major results

- The development of state-of-the-art risk management and modelling techniques suitable for dynamic ICT systems, where the run-time configuration of the system is unknown at design time.
- SERSCIS has developed a layered semantic modelling approach that can be adapted to any domain, and allows security experts and system designers to insert their knowledge, e.g. the security experts define generic assets, security threats and controls, while the system designer models their own system via intuitive and easy to use graphical user interfaces.
- Machine reasoning is then used to apply security expertise to analyse the system design and automatically identify security risks.
- The output from this analysis can also be used to analyse monitoring data from the system at run-time, using Bayesian and semantic inference to determine which threats are active and estimate the most likely root cause.
- The results of SERSCIS have been published in high quality conferences and journals.

Potential exploitation strategy

The semantic modelling work is being extended via further projects e.g. in FP7 OPTET which deals with socio-economics issues in trustworthy ICT systems, and the UK TRIFORM project which is investigating trust monitoring approaches in eHealth systems.

SERSCIS models are also being applied in other application domains, e.g. to model risks for long-term archive service operations.

The System Composer from SERSCIS is being developed into a risk management tool for system designers, so they can exploit these semantic knowledge bases to automatically identify potential security risks within their own system designs.

Links and references

- SERSCIS Project: http://www.serscis.eu/, Surridge, M., Nasser, B., Chen, X., Chakravarthy, A., Melas, P. (2013) Run-Time Risk Management in Adaptive ICT Systems. Forthcoming: Proceedings of the 8th International Conference on Availability, Reliability and Security, September 2013, University of Regensburg, Germany, (ARES, 2013) Surridge, M., Chakravarthy, A., Hall-May, M., Chen, X., Nasser, B., Nossal, R. (2012).
- SERSCIS: Semantic Modelling of Dynamic, Multi-Stakeholder Systems. In: Proceedings of the 2nd SESAR Innovation Days, November 2012, Braunschweig, Germany (EUROCONTROL) Kostopoulos, D., Leventakis, G. Tsoulkas, V., and Nikitakos, N. (2012).
- An Intelligent Fault Monitoring and Risk Management Tool for Complex Critical Infrastructures: The SERSCIS Approach in Air-traffic Surface Control. In: 14th International Conference on Computer Modelling and Simulation (UKSim), March 2012, Cambridge, UK (IEEE Computer Society) Surridge, M., Chakravarthy, A., Bashevoy, M., Wright, J., Hall-May, M., Nossal, R. (2012).
- SERSCIS-Ont: Evaluation of a Formal Metric Model using Airport Collaborative Decision Making. International Journal On Advances in Intelligent Systems, 4(3) (IARIA) Hall-May, M., Chakravarthy, A., Leonard, T., and Surridge, M. (2011).
- Semantic Modelling of Resource Dependability for SLA-based Service Governance. In: Handbook of Research on Service-Oriented Systems and Non-Functional Properties: Future Directions, eds. Reiff-Marganiec, S. and Tilly, M., ch. 18, pp. 401-441 (IGI Global) Surridge, M., Chakravarthy, A., Bashevoy, M. and Hall-May, M. (2010).
- Serscis-Ont: A Formal Metrics Model for Adaptive Service Oriented Frameworks. In: Second International Conference on Adaptive and Self-adaptive Systems and Applications (ADAPTIVE 2010), November 2010, Lisbon, Portugal. [Best Paper Award] Boniface, M., Surridge, M., Hall-May, M., Bertram, S. and Briscombe, N. (2010).
- On-Demand Dynamic Security for Risk-Based Secure Collaboration in Clouds. In: IEEE International Conference on Cloud Computing 2010, July 2010, Miami, FL, USA.

SPECS: Secure Provisioning of Cloud Services based on SLA management

Name: Massimiliano Rak
Organisation: Second University of Naples / CeRICT

Topics recommended for the 2016-2017 Work Programme

Cloud Security, Security Service Level Agreement and measurable Security:

While the economic and technological advantages of cloud computing are appealing, the migration of key sector applications onto the cloud has been partly limited by the lack of trust in the security assurance provided by Cloud Service Providers (CSPs).

The diffusion of the "as-a-service" approach, which delegate each kind of resources to remote cloud provider, needs the adoption of new security methodologies, able to introduce security mechanisms and controls in a modular way, offering security as a service.

Security Service Level Agreements looks a clear and standard way to clarify the responsibilities of the involved parties, Cloud Service Providers (CSPs) and Cloud Service Customers (CSCs), respect to the security threats.

At state of art there is a lack of:

- Clear definition of possible Service Level Objectives, i.e. the terms defined in the Security SLAs and their measured levels.
- Clear and shared techniques for measuring security parameters.

Projects major results

SPECS proposes an innovative Security Platform-as-a-Service that offers a solution for the SPECS' Security-as-a-Service approach, based on SLA life cycle management. Such platform enables the delivering of security services, described in details through Security SLAs. Cloud Service Customers are able to define at fine grain level the security features they need through a user-centric negotiation of Cloud SLA, that helps CSCs to negotiate Cloud SLA effectively with a set of CSPs, by understanding the resulting trade-offs.

In order to support CSCs to verify the correctness of the services offered by CSPs, SPECS offers innovative Solutions for Continuous Security Monitoring, which implements SLA monitoring solutions dedicated to continuously control the security offered by CSP and to help ensuring the granted security service level objectives.

Moreover SPECS aims at developing Innovative Security Services to Enforce SLA, i.e. services dedicated to grant security features to CSCs in order to fulfil an agreed SLA.

Potential exploitation strategy

Standardisation effort: Security Service Level agreements needs a clear and shared view among CSPs and CSCs, in order to enable correct negotiation and definition of responsibilities and actions to be performed. In such a context standardisation bodies have a relevant role.

- SPECS participates actively to different standardisation bodies, contributing to definition of standards and following the existing activities in order to produce solutions aligned with the state of art.
- Dissemination: scientific results are diffused trough publications to relevant conferences and through cooperation with other research projects.
- Prototypes are open source and diffused through well known code repository.
- The exploitation plan of results after the end of the project and possible commercialisation of products is under discussion and will be publicly available after the first year of the project.

An update since the last Concertation meeting (March 2014)

- SPECS has collaborated with the A4Cloud project, organising a shared workshop and participated to the CIRRUS final event.
- SPECS shares technologies and innovative solution with other cloud-related projects like MODA-clouds and reusing results from closed projects like mOSAIC or SLA-at-SOI.
- SPECS members participate to the EU C-SIG on SLA and support standardisation activities on cloud and SLA by SLA and NIST.

Links and references

http://specs-project.eu/

STORM CLOUDS: Surfing Towards the Opportunity of Real Migration to cloud-based public services

Name: Agustin González-Quel
Organisation: Ariadna SI

Topics recommended for the 2016-2017 Work Programme

There are two interrelated issues with regards to the privacy of data and security:

- Balance between a private cloud that I could use for my personal work, record, data, etc. and the information/applications from this cloud that I want to share with a public cloud. For instance, some of my files may be shared with the Taxes system of my country for tax payment, while some other will be shared with my hospital, etc.

Projects major results

For the time being we are involved in the selection of public services for the 1st cycle of migration of these services to the cloud. So far, we have a list of criteria as an intermediate result, as well as the migration guidelines being prepared by the cloud providers.

Potential exploitation strategy

When the guidelines and experience are well settled down, these will be easily transferred as commercial project to Public Administration bodies.

An update since the last Concertation meeting (March 2014)

- The four cities involved in the project (AGUEDA, MANCHESTER, VALLADOLID, THESALONIKI) have prepared a list of services that can be transferred to the cloud.
- An open innovation process has started with the selection of stakeholders and the selection of which services are to have priority in the process.
- In parallel, the cloud is being prepared and the technical partners also expressed their views on which services can be moved more efficiently to the cloud.

Links and references

http://stormclouds.eu/

STRATEGIC: Service distribution network And Tools for interoperable programmable, and unified public Cloud services

Name: Nuria Rodríguez Domínguez
Organisation: Atos Spain

Topics recommended for the 2016-2017 Work Programme

STRATEGIC integrates leading edge technology developments completed in recent research projects in order to produce a novel framework that facilitates organisations (notably public bodies) in leveraging the benefits of public cloud services. Despite the benefits of cloud, there are still limited experiences in public organisations. In this context, a topic to be considered in the next work programme could be:

- Security and privacy management mechanisms (along with related policies) are not yet fully integrated into public cloud services.

Projects major results

STRATEGIC is still in an early stage of its lifecycle. First steps have been addressed to the requirements collection and analysis as well to the definition of the use cases. The main outcomes to date are:

- A methodology has implemented in order to collect, document and analyse the users requirements based on two main axes: the state of the art review in the area of e-Government/public cloud services and the development of a questionnaire supported by few selected interviews in order to get input from a diverse set of stakeholders, both cloud users and providers (including cloud developers, integrators, brokers, and software vendors).
- 11 applications were analysed belonging to different scenario types. For every application, the expected added-value was analysed, both for the applications-per-site and the STRATEGIC framework. In addition, cross-scenario links were established to potentially benefit from the synergy of the applications.

Nowadays, STRATEGIC is focusing its efforts on the specification of its architecture and on the technical specifications. Furthermore, a first prototype of the STRATEGIC cloud broker and marketplace is expected by the end of 2014.

Potential exploitation strategy

Overall, the wider implementation and use of the STRATEGIC solutions will be based on the following exploitation modalities:

- The adoption of the STRATEGIC migration, adaptation/localisation and application development solutions from a larger mass of public bodies (beyond the pilot sites (public bodies) of the consortium).
  The partners will focus on attracting customers (public bodies) from their territories/regions, while at the same time attempting to extend their activities to other countries and regions as well. In several cases, the partners will attempt to attract public bodies and deploy solutions for them on the basis of partnerships with other ISVs and service providers, especially for countries where they do not have established presence.
- The exploitation and wider use of the solutions and tools of the STRATEGIC framework by third-party integrators, cloud developers and ISVs.
  As part of this exploitation modality the consortium will endeavor to commercially exploit the STRATEGIC framework (including its tools and services), through selling licenses, charging pay-per-use fees, charging hosting and cloud provisioning services, as well as through offering supporting services such as training, maintenance and consulting services.

An update since the last Concertation meeting (March 2014)

- A list of about 155 stakeholders was created to disseminate the questionnaire as well as to regularly keep them informed about the project results.
- Liaisons have been established with other EU projects (e.g. Broker@Cloud).

Links and references

http://www.strategic-project.eu
https://twitter.com/strategic_eu

SUCRE: Supporting Cloud Research Exploitation

Name: Eleni Toli

Organisation: University of Athens

Topics recommended for the 2016-2017 Work Programme

The following topics on future research have emerged from events organised by SUCRE:

- Convergence of heterogeneous Cloud solutions: Driven by requirements in many application domains, data and service mobility lie among the most important challenges that need to be tackled. Such intercontinental-scale mobilisation of software and data entails both technical and policy issues.
- Network Interconnection and Orchestration: Collaboration between Europe and Japan, as well as other parts of the globe in the area of Cloud computing inevitably brings forth the issue of interconnectivity. Integration of the IT infrastructure available in Europe and other regions will produce a complex ecosystem offering services to users over shared federated Cloud environments adhering to potentially different architectural and deployment models. The support of dynamic provisioning and interconnection of resources belonging to independent heterogeneous Cloud infrastructures are issues that should become the focus of extensive joint research.
- Cloud of Things: With billions of physical devices connected and interacting, new deployment and monitoring models will be called for. The Cloud of Things also brings forth new challenges related to Big Open Data management. Voluminous, social and personal data produced by existing and future wearable devices as well as by other sources will require new privacy enhancing technologies to deal with trust and security issues.

Projects major results

SUCRE aims at investigating and suggesting ways to reinforce the adoption of Open Cloud solutions by key stakeholders. In doing so, SUCRE has engaged two user communities whose future computing requirements could be very well accommodated by Open Clouds, namely, the public sector and the healthcare service provisioning industry.

Through the formation and operation of an EU-Japan experts group, SUCRE is also providing support to Cloud interoperability and collaboration between the two regions. In order to fulfil its impact in this context, SUCRE has already contributed a number of significant results.

Three issues of the project's official CloudSource magazine have been released, while two videos have been prepared illustrating the benefits of Open Clouds for the two aforementioned domains. Moreover, a set of public reports effectively captures the current status of Open Cloud adoption by the two areas of interest. Enriched with insights given by pertinent stakeholders through qualitative interviews, these documents pinpoint the tangible benefits as well as the technical, economic and

political challenges that still have to be met in order to facilitate the wide adoption of Open Cloud technologies.
SUCRE has also released a public document that details the outcomes of the activities of the EU-Japan experts group in the form of a set of recommendations for strengthening Cloud interoperability and collaboration between the two regions.

Potential exploitation strategy
SUCRE is a Support Action, and therefore no technical implementations or technical products will be delivered. Nevertheless, the consortium will produce an End Product comprising mainly:

- The CloudSource magazine and video products addressing user communities, public administrations, SMEs, and the Industry, which showcase the developments in Open Clouds.
- High-impact reports, such as the Primer documents for the Public and Healthcare Provisioning Sectors, and the Recommendation Report of the EU-Japan Expert Group.
- The SUCRE portal targeting the communities of Cloud Computing and Open Source, and aspiring to serve as a one-stop-shop for new users and potential adopters in those areas.

The project recognises the creation of a community with a critical mass as the main driving factor that will help sustain its End Product. To this end, the SUCRE consortium has already engaged in a number of exploitation activities, the vast majority of which aim at promoting the project outcomes by means of presentations at major related events, establishing strong collaboration links with other projects, and seeking publications to pertinent Cloud and Open Source scientific journals, as well as online and paper magazines.

An update since the last Concertation meeting (March 2014)
SUCRE has established strong collaboration links with a large number of projects in the Unit, which are mainly manifested by the organisation of joint events. More specifically:

- SUCRE collaborated with the OCEAN and ClouT projects for the joint organisation of a EU-Japan workshop that took place last May in Brussels, and fostered an international dialogue on interoperability and collaboration between the two regions in the fields of Cloud Computing and Internet of Things. A sister event was co-organised in July 2014 in Tokyo by the same projects and carried on the initiated dialogue between stakeholders from Europe and Japan.
- Another SUCRE workshop, co-organised with the ARTIST, CELAR, and MODAClouds projects, took place last May in Athens, the frame of the FIA 2014 event. The workshop specifically focused on Cloud solutions by EU-funded projects that address problems related to elasticity, migration and interoperability.
- Finally, SUCRE is currently collaborating with the CloudCatalyst, CloudforEurope, and OCEAN projects in the joint organisation of its final workshop that will be held this September in Bled, Slovenia, in the frame of the CLASS 2014 conference.

Besides, the active collaboration between SUCRE and other projects was also streamlined through the online presence and exchange of information between the project portals, as well as by the submission and publication of various project-related articles in the SUCRE CloudSource magazine.

Links and references
http://www.sucreproject.eu/.

SyncFree: Large-scale computation without synchronisation

Name: Tyler Crain
Organisation: INRIA/LIP6

Topics recommended for the 2016-2017 Work Programme

- Conflict-Free Consistency for low-cost cloud computing:
  Traditional strong-consistency approaches do not scale to extreme-scale distributed applications, such as social networks or multiplayer games, and their huge quantities of frequently-changing geo-distributed shared data, for performance, reliability and scalability reasons.
  Cheaper, faster, more relaxed consistency levels, in which concurrent updates are allowed, are more appropriate, but have been so far reserved to highly-specialised programmers. Our SyncFree project [1] is developing tools (conflict-free replicated data types, programming tools, and platforms) that provably ensure eventual consistency, while still being easy to program with.
  This is an important topic, which should be continued and expanded (e.g. via standardisation) to make this approach available to the industry at large. This will lower the cost and the entry barriers to large-scale cloud computing, opening new opportunities for European industry and SMEs.
- Interplay between conflict-free consistency and strong consistency: Despite the scalability advantages of weak consistency, an application occasionally needs to issue a strong update, for instance to ensure some global invariant.
  The correct interplay between weak consistency and occasional strongly-consistent operations is essential to the design of dependable and highly-scalable applications. This topic has been little studied so far, and raises both theoretical and practical issues.
- Scalable security in conflict-free distributed systems: Existing approaches to security typically assume that a policy change is immediately visible, i.e, require strong consistency. In today's decentralised cloud architectures, this is a brittle and non-scalable assumption.
  It is an important topic to study how to ensure a satisfactory level of security under eventual consistency.

Projects major results

Current approaches for ensuring data consistency in modern clouds, made up of loosely coupled, widely-distributed and heterogeneous localised datacentres, require huge investment and highly-specialised expertise, available only to a few large monopolies.
SyncFree has already started to address this challenge thanks to a simple yet principled approach called Conflict-Free Replicated Data Types (CRDTs) [9]. CRDTs avoid the complexities of ad-hoc approaches, while maintaining the scalability advantage. Here is the insight.
By following a few simple mathematical principles, for example commutativity, distributed updates can occur without synchronisation, while guaranteeing eventual consistency. What's more, CRDTs ease development, by encapsulating the replication and concurrency properties of common shared objects, such as sets, maps, sequences, or graphs. These data types are easily combined to form robust, scalable, powerful applications.
The project partners have already demonstrated a number of CRDTs and deployed them in example applications. Together, the industrial and academic partners of SyncFree are currently formalising specifications for highly-scalable innovative applications, building programming and deployment platforms, mathematically proving that they are correct, and preparing for extreme-scale experiments on real-world crowd-source applications.
A natural follow-up of this work will be to develop and propose standards for these data types including libraries of open-source data structures. This will aim to be used by European industry looking to benefit from the cloud, thus providing scalable solutions to quickly develop cloud-based applications.

Potential exploitation strategy

The SyncFree project is advancing both the theory and practice of large-scale application architectures, and especially of CRDTs and related mechanisms.
With several SyncFree partners coming from European based enterprises that already have large user bases and feel the need for increased scalability in their applications, the project will include an extreme-scale crowd-sourced experiment, pushing the scalability needs of real world applications. Furthermore, an open-source cloud storage platform [2], including a library of CRDTs in addition to strongly consistent abstractions, to be used in future scalable distributed applications, will be made available, leaving a lasting and beneficial impact far beyond the end of the project.
Using these open source libraries, organisations will be able to create highly scalable programs more easily, meeting the strict consistency requirements present in today's highly connected services while improving user experience through low latency and fault tolerance.
These advantages will help extend the reach of the cloud into mainstream connected applications and services and provide the platform to create new and innovative cloud-based businesses.

An update since the last Concertation meeting (March 2014)

The collaborators within in the SyncFree project have been making significant progress on designing abstractions for extreme-scale applications, this work will be essential for defining future cloud standards as they define the system on which future cloud services and applications will be developed to run on top of.
Selected results include a designs for scalable dictionary data structure [3], methods for efficiently storing and exchanging consistent data-types across multiple datacentres [4], methods for providing

scalable invariants [7], abstractions for increasing the scalability of online collaborative editing programs[4], and methods for verifying the correctness of scalable applications [8].
To verify these and other results that are being exploded during the project, a state-of-the-art, industrially realistic, cloud storage platform is being developed with the support of Basho Technologies, a world leader in scalable could storage technology. This platform is provided as open-source code on GitHub [2].

Links and references

[1] SyncFree FP7 Project, http://syncfree.lip6.fr
[2] SyncFree GitHub Repository, https://github.com/SyncFree/
[3] Russell Brown, Sean Cribbs, Sam Elliott and Christopher Meiklejohn: Riak DT Map: A Composable, Convergent Replicated Dictionary. PaPEC 2014.
[4] Paulo Sérgio Almeida, Ali Shoker and Carlos Baquero: Efficient State-based CRDTs by Decomposition (Work in progress report). PaPec 2014.
[5] Carlos Baquero, Paulo Sérgio Almeida and Ali Shoker: Making operation based CRDTs operation based (Work in progress report). PaPec 2014.
[6] Mehdi Ahmed-Nacer, Pascal Urso, Nuno Preguica and Valter Balegas: Merging OT and CRDT Algorithms. PaPec 2014.
[7] Valter Balegas and Mahsa Najafzadeh: The Case for Fast and Invariant-Preserving Geo-Replication. PaPec 2014.
[8] Burcu Külahçıoğlu Özkan, Erdal Mutlu and Serdar Tasiran: Towards Verifying Eventually Consistent Applications. PaPec 2014.
[9] Strong Eventual Consistency and Conflict-free Replicated Data Types: Marc Shapiro, Nuno Preguiça, Carlos Baquero and Marek Zawirski, http://research.microsoft.com/apps/video/default.aspx?id=153540&r=1.

UQASAR: Universal Quality Assurance & Control Services for Internet Applications with Volatile Requirements and Contexts

Name: Aitor Elorriga
Organisation: Innopole

Topics recommended for the 2016-2017 Work Programme

UQASAR implements semantics in the field of Quality Assurance, whereas the techniques applied have to be light-weight or the performance of the platform can be seriously affected. This is especially critical in a big data context where the different monitoring services will be continuously gathering data from a mutating environment composed by personal IDEs of developers, designer tools, source code repositories, etc.

In this field, there are several issues that can be taken into consideration for future work programmes:

- Complex semantics in real time big data environments with acceptable time responses.
- Identification of good/bad software engineering patterns out of the monitoring of the development environments.

These patterns could be obtained not only by the metrics gathered directly from repositories and tools but also from the categorisation of the coders' and testers' personal styles. This is of special interest in highly social development/testing contexts.

Projects major results

Up to now UQASAR has generated 2 main outcomes:

- An early prototype to demonstrate the validity of the concept, which is basically the monitoring of development contexts to provide the end user with evidences of the good or bad track of a development project or product.

  Currently the prototype allows the user to create quality models with quality goals, indicators and metrics; to set up quality assurance projects based on the previous models and gather data from two selected tools (JIRA and SONARQube).

  The Full Prototype, that will be ready by the end of February 2015, will propose changes to the Quality Objectives, Indicators and Metrics based on previous experience, will allow to establish negotiation processes to define the values of quality (e.g. thresholds of metrics).
- A methodology that indicates end-users how to identify the adequate collection of quality goals, indications and metrics depending on the type of process followed (agile, traditional) or type of product to be developed. UQASAR is looking for potential alfa-users of the platform once it's launched on February/March 2015.

Potential exploitation strategy

The aim of UQASAR is to become a service in the cloud, though there are some restrictions at this moment, such as the dependency to specific life cycle tools, though the platform is open enough to incorporate new tools as the UQASAR community grows.

- Basic service can be offered at a very low price and charge for added value services such as smart reporting or statistical analysis.
- The setup of a SaaS orientation requires the creation a fully dedicated startup, which is one of the most promising options handled in our Exploitation Plan.

Main barriers are the non-existence of a community figure for the setup of the company, as well as the initial personnel costs required for exploitation.

An update since the last Concertation meeting (March 2014)

Cross collaboration has started with RISCOSS project:

- Several conversations have been maintained with different members of the consortium in order to define common needs.
- JIRA and SONARQube Wrappers developed within UQASAR are of interest for RISCOSS.
- For UQASAR the risk management strategy identified in RISCOSS is also of big interest to be integrated in the monitoring of management stage, which is part of our Monitoring Services.

Links and references

Website: http://www.uqasar.eu (blog, project overview, articles, conferences, deliverables)
Facebook: http://www.facebook.com/UQASAR
Twitter: @uqasar
LinkedIn: https://www.linkedin.com/groups/UQASAR-Project

CloudWATCH is funded by the European Commission
DG Connect Software & Services, Cloud. Contract No. 610994