

From **cloud computing** and software technologies research into **innovative services** for the European digital single market

A portfolio of offers for **trusted and secure services**



March 2015





Disclaimer

CloudWATCH (A European Cloud Observatory supporting cloud policies, standard profiles and services) is funded by the European Commission's Unit on Software and Services, Cloud Computing within DG Connect under the 7th Framework Programme.

The information, views and tips set out in this publication are those of the CloudWATCH Consortium and its pool of international experts and cannot be considered to reflect the views of the European Commission.

Index


A portfolio of offers for trusted and secure services	3
Cloud Computing, Internet of Services and Advanced Software Engineering	
Market ready in 2015- 2016.....	5
ARTIST - Advanced software-based seRvice provisioning and migraTion of legacy Software.....	6
BETaaS - Build the Environment for the Things as a Service.....	7
BIGFOOT - BigData Analytics of Digital Footprints.....	8
Broker@Cloud - Enabling continuous quality assurance and optimization in future enterprise cloud service brokers.....	9
CELAR - Automatic, multi-grained elasticity-provisioning for the Cloud.....	10
CloudScale - Scalability Management for Cloud Computing	11
CloudSpaces - Open Service Platform for the next generation of Personal Clouds	12
COMPOSE - Collaborative Open Market to Place Objects at your SService.....	13
LEADS - Large-Scale Elastic Architecture for Data as a Service.....	14
MARKOS - The MARKet for Open Source – An Intelligent Virtual Open Source Marketplace.....	15
MODAClouds - Model-Driven Approach for design and execution of applications on multiple Clouds	16
OPENi - Open-Source, Web-Based, Framework for Integrating Applications with Cloud-based Services and Personal Cloudlets	17
OSSMETER - Automated Measurement and Analysis of Open Source Software	18
PaaSage - Model Based Cloud Platform Upperware	19
PROSE - Promoting open source projects in European Projects.....	20
PROWESS - Property-based testing for web services.....	21
RISCOSS - Managing Risk and Costs in Open Source Software Adoption	22
U-QASAR - Universal Quality Assurance & Control Services for Internet Applications with Volatile Requirements and Contexts.....	23
Software Engineering, Services & Cloud Computing - Market ready in 2015- 2017.....	24
ASCETiC - Adapting Service lifeCycle towards Efficient Clouds	25
CACTOS - Context-Aware Cloud Topology Optimisation and Simulation.....	26
CloudCatalyst - Reenergize productivity, efficiency and competitiveness of European economy through Cloud Computing	27
CloudingSMEs.....	28
CloudWATCH - A European cloud observatory supporting cloud policies, standard profiles and services.....	29
CloudWave - Agile Service Engineering for the Future Internet	30
ClouT - Cloud of Things for empowering the citizen clout in smart cities.....	31
CoherentPaaS - Coherent and Rich PaaS with a Common Programming Model	32
ECIM - European Cloud Marketplace for Intelligent Mobility.....	33
Envisage - Engineering Virtualized Services	34
HEADS - Heterogeneous and Distributed Services for the Future Computing Continuum.....	35
HTML5Apps - HTML5 for Apps: Closing the Gaps	36
MONDO - Scalable Modelling and Model Management on the Cloud	37
ORBIT - Business Continuity as a Service	38
Panacea - Proactive Autonomic Management of Cloud Resources.....	39
S-CASE - Scaffolding Scalable Software Services.....	40
SeaClouds - Seamless adaptive multi-cloud management of service-based applications.....	41
SyncFree - Large-scale computation without synchronisation	42

Advanced Cloud Infrastructures and Services - Market ready in 2016 -2018.....	43
ALIGNED - Aligned, Quality-centric Software and Data Engineering	44
AppHub - The European Open Source Market Place.....	45
ARCADIA - A novel reconfigurable by design highly distributed applications development paradigm over programmable infrastructure	46
BEACON - Enabling Federated Cloud Networking	47
CHOREVOLUTION	48
CloudWATCH2 - Think Cloud Services for Government, Business & Research	49
CLARUS - A framework for user centred privacy and security in the cloud	50
CLOUDLIGHTNING – Self-Organising, Self-Managing Heterogeneous Cloud	51
CloudSocket - Business and IT-Cloud Alignment using a Smart Socket	52
CloudTeams - Collaborative Software Development Framework based on Trusted, Secure	
Cloud-based Pool of Users	53
CYCLONE - Complete Dynamic Multi-cloud Application Management.....	54
DICE - Developing Data-Intensive Cloud Applications with Iterative Quality Enhancements.....	55
ENTICE - dEcentralized repositories for traNsparent and efficienT vRtual maChine opErations.....	56
ESCUDO-CLOUD - Enforceable Security in the Cloud to Uphold Data Ownership.....	57
HOLA CLOUD - Effective collaboration for European RD and Innovation in software, services and	
Cloud computing: Knowledge discovery and Roadmapping.....	58
HyVar - Scalable Hybrid Variability for Distributed Evolving Software Systems.....	59
iKaaS - intelligent Knowledge-as-a-Service Platform.....	60
INPUT - In-Network Programmability for next-generation personal clouD service supporT	61
IOWStack - Software Defined Storage for Big Data.....	62
Mikelangelo - Improving Responsiveness and Agility of HPC Cloud Infrastructure.....	63
MUSA - Multi-cloud Secure Applications.....	64
PaaSword - A Holistic Data Privacy and Security by Design Platform-as-a Service Framework	
Introducing Distributed Encrypted Persistence in Cloud-based Applications	65
PICSE - Procurement innovation for cloud services in Europe.....	66
RAPID - Heterogeneous Secure Multi-level Remote Acceleration Service for Low-Power Integrated	
Systems and Devices.....	67
SERECA - Secure Enclaves for REactive Cloud Applications.....	68
SLALOM - Service Level Agreement - Legal and Open Model	69
SLA-Ready - Making Cloud SLAs readily usable in the EU private sector.....	70
SSICLOPS - Scalable and Secure Infrastructures for Cloud Operations.....	71
SUNFISH - SecUre iNformation SHaring in federated heterogeneous private clouds.....	72
SUPERCLOUD - User-centric management of security and dependability in clouds of clouds.....	73
SWITCH - Software Workbench for Interactive, Time Critical and Highly self-adaptive Cloud	
applications	74

A portfolio of offers for trusted and secure services

The CloudWATCH project is delighted to present a portfolio of 67 stimulating service offers for trusted and secure services that will emerge from the E2 unit of DG CONNECT Software & Services & Cloud Computing. The service cards aim to provide a focus on the **cloud ecosystem emerging from European research and innovation projects**. It takes a pragmatic approach to market uptake and sustainable competitiveness for wider uptake and commercial exploitation.

Concertation meeting in numbers




European research priorities have increasingly adapted to the fast evolution of cloud and related services. Developments taking place today in FP7 ICT (e.g. Call 8 & Call 10) and now demonstrated with successful Call 1 projects from H2020, range from big data and content, to the Internet of Things (IoT), cloud computing and the Internet of Services. These research and innovation activities will take us to a new era of digital societies and economies based on an agile, adaptive and dynamic collaboration between organisations, communities and individuals, with ubiquitous and instant access to information.

Europe is a world leader in research and innovation on cloud software and services and it is imperative that it retains this position. This very much depends on Europe innovating at the technological level and demonstrating tangible business value. European research projects are thinking strategically which means looking at technology and market uptake as part of the same equation. These service offers demonstrate this through challenging, interoperable cloud services playing which play a very important role in extending the market and in bringing business benefits to both the supply and demand sides.


This portfolio has been compiled by the CloudWATCH project. We'd like to thank all projects for providing their service offers.

We hope you enjoy them

On behalf of the CloudWATCH Consortium



Cloud Computing, Internet of Services and Advanced Software Engineering Market ready in 2015- 2016





ARTIST – Advanced software-based seRvice provisioning and migraTion of legacy Software

Project start date 01/10/2013

Project finish date 30/09/2015

What user need or pain point is being addressed?

ARTIST targets mainly software owners who need to modernize their applications and business model towards the cloud. ARTIST focuses on the modernization of applications based on three pillars that in most occasions cannot be tackled independently. The business model modernization and the selected deployment model involves some architectural constraints in the application. Furthermore, for companies offering their applications as a service, a transformation in the organizational processes is also needed in order to support the delivery of these services.

ARTIST-based services can be offered by ISVs (Independent Software Vendors) and other owners of bespoke software systems that require modernisation to the cloud. In particular, clients requiring the modernisation of several related software programs will be targeted. In some cases these end clients can perform the modernisation themselves using the tools, providing they have the necessary IT skills base. The effort for the transformation of the product is reduced thanks to the guided process and semi-automatic tools provided by ARTIST.

For software owners of non-cloud software applications, ARTIST can help companies in transforming non-cloud software applications not merely to run on cloud, but to take full advantage of cloud features by shifting to a software-as-a-service (SaaS) business model.

For cloud providers, the ARTIST solution can be a complement to their cloud infrastructure attracting offering potential customers with non-cloud complaint software a suite of tools. For example, IT consultants who adopt the tools and incorporate them into their own offerings will most likely be oriented around application portfolio management. This will be the case of the larger industrial partners in the consortium. Also, the migration encompasses fewer risks due to a previous feasibility analysis.

For software integrators, ARTIST provides business opportunities in supporting third companies in the migration of their applications.

How will the solution/service benefit the end-user?

Many software companies feel that cloud may bring plenty of new business opportunities to them but do not know how to approach it, or even where to start to migrate their applications to run on cloud. They may see migrating to the cloud as a way of reducing operational costs for code maintenance and ad-hoc customisations for obsolete technologies that are perhaps no longer supported due to a lack of skilled staff. With a fast evolving market, it's important that they can adapt easily to new challenges. Thanks to ARTIST, they will be able to migrate to the cloud and create future value through new market generation for their products. They'll also improve customer retention, application performance and increase revenue growth.

ARTIST provides a tool-supported methodology to guide the user in performing the modernization. Unlike most alternatives available today, the methodology is Open Source, is source and target agnostic, and emphasizes the reuse of work from one project to the next. This leads to greater cost reductions and shorter cycles. Additionally, ARTIST lets the user select the best cloud provider for migrating the application. It is focused on Cloud-compliant architectural issues at both application and infrastructure levels and allows code optimization by following a methodology-driven approach. Finally, it includes a feasibility analysis before any investment is actually made, which reduces the risks associated with migration.

Website: <http://www.artist-project.eu/>

BETaaS - Build the Environment for the Things as a Service

Project start date 01/10/2012

Project finish date 31/03/2015



What user need or pain point is being addressed?

The BETaaS project is addressing developers communities in the field of IoT and M2M, Hardware vendors, Development Companies, M2M & IoT standardization groups and frameworks, Service Providers (SaaS, PaaS, BaaS, Mobile, Telcos and Utility Companies). The objective of the BETaaS project is to ensure interoperability with M2M frameworks with QoS, Trust and Security under control, helping developers and vendors create new world class applications and experiences over a world of Things.

How will the solution/service benefit the end-user?

The BETaaS project enables application developers to easily use things-based environments as services. It also allows hardware vendors to make their products available for use in the IoT context by a wider range of users and to create/extend innovative products or experiences for their product lines creating differentiated enriched experiences (app-things). BETaaS also allows big players to build business models around the concept of IoT enabling services needed by IoT using compatible Things in their offerings. The BETaaS platform also allows the big players to incorporate such services in a scalable fashion, being able to offer built-in Security and QoS capabilities and the creation of applications that interact with different vendors' technologies and products (hardware, sensors, devices, etc.) without worrying about security and compatibility issues.

Website: <http://www.betaas.eu/>



BIGFOOT – BigData Analytics of Digital Footprints

Project start date 01/10/2012

Project finish date 30/09/2015

What user need or pain point is being addressed?

BigFoot addresses some of the most vital issues related to the deployment of large-scale data analytic technologies and innovative analytics-as-a-service business models :

- » **Data interaction is hard.** Current approaches lack an integrated interface to inspect and query (processed) data. Moreover, not much work has been done on the literature to optimize the efficiency, and not only the performance, of interactive queries that operate on batch processed data.
- » **Parallel algorithm design is hard.** While the design of parallel algorithms is already a difficult topic per se, current systems make the implementation of even simple jobs a tedious and exhausting experience.
- » **Deployment tools are poor.** Management tools are still in their infancy and target solely bare metal clusters.
- » **Impact of virtualization.** The effects of compute and network virtualization on the performance of data-intensive services has been largely overlooked in the literature and available solutions.
- » **Lack of optimizations.** Current systems entrust users with the task of optimizing their queries and algorithms. Moreover, dataflow and storage mechanisms are data-processing oblivious, which leave room for several optimizations that have not been addressed by current solutions.

How will the solution/service benefit the end-user?

A platform-as-a-service solution for processing and interacting with large volumes of data. Bigfoot builds upon and contributes to the apache hadoop ecosystem and the apache openstack project.

Key differentiating benefits provided by BigFoot include :

Analytics-as-a-Service:

- » Self-tuned deployments in private (and public) clouds
- » Hardware and data consolidation through virtualization
- » Performance enhancements to mitigate bottlenecks
- » Multi-site add-ons for geo-replication

Resource allocation mechanisms:

- » New scheduling components to deal with heterogeneous work-loads
- » New work-sharing optimizations for both batch and interactive engines

In-situ querying of RAW data:

- » Distributed query mechanism to operate on heterogeneous RAW data
- » On-the-fly indexing for modern storage devices

High-level languages:

- » Scalable Machine Learning library
- » Time Series Library

Target groups benefiting most of BigFoot solutions (see our website for details):

- » In a generic context - Academic Researchers, Engineers and Data Scientists, Big Data companies
- » In the cyber-security context - Security software companies, CERT teams, Security researchers
- » In the Smart Grid context - Residential electric consumers, Managers of utility companies, Energy data scientists, Smart city services operators

Website: <http://bigfootproject.eu/>

Broker@Cloud - Enabling continuous quality assurance and optimization in future enterprise cloud service brokers

Project start date 01/11/2012

Project finish date 31/10/2015



What user need or pain point is being addressed?

As the number of externally-sourced services in an enterprise cloud environment increases, it becomes increasingly more difficult to keep track of when and how services evolve over time. Moreover, it becomes increasingly more difficult to appreciate the impact that a change can have with respect to a service's compliance to different policies and regulations, its conformance to normative technical specifications and contracts (Service Level Agreements), and generally, with respect to the fulfilment of all different kinds of functional and non-functional requirements surrounding a particular service's usage. The availability of increasing numbers of cloud services that offer similar functionality under comparable terms of provision, despite the obvious benefits, is nevertheless adding to the overall complexity and contributes towards even higher management costs, as enterprises have to invest a lot of effort in identifying suitable alternatives, either proactively, or on a need-to basis.

To address the aforementioned mentioned pain points, the Broker@Cloud framework mechanisms will help cloud service brokers and service consumers to evaluate and analyze the execution of cloud services offered by the service providers. Broker@Cloud will also enable optimized service selection, failure prevention and recovery capabilities to consumers and service providers during service operation.

How will the solution/service benefit the end-user?

Broker@Cloud provides a unique holistic solution for quality assurance, optimization, failure prevention and recovery brokerage mechanisms together with a baseline framework for building new cloud service brokerage mechanisms. Through the Broker@Cloud framework mechanisms, all services offered on the integrated cloud service brokerage platform are continuously monitored against SLA agreements. In case of a violation, the framework's mechanisms will reason about appropriate adaptation actions such as substituting a service. In addition, the failure prevention mechanism can indicate that there are risks for other consumers of the service that has just failed, and produce warnings or proactively generate an adaptation plan. Through complex event processing, the brokerage framework mechanisms can also detect that the response time of a service is dropping rapidly, and predict a forthcoming failure. The prevention mechanisms of Broker@Cloud can then plan and enact the necessary adaptations. Moreover, Broker@Cloud provides a user-friendly way for the cloud consumers to declare their preferences based on a number of precise and imprecise criteria and receive recommendations through the comparison of the available cloud services. Finally, based on the Broker@Cloud framework, novel advanced cloud service brokerage mechanisms could be built, interconnected with each other and integrated into cloud service brokerage platforms.

Website: <http://www.broker-cloud.eu/>



CELAR – Automatic, multi-grained elasticity-provisioning for the Cloud

Project start date 01/10/2012

Project finish date 30/09/2015

What user need or pain point is being addressed?

CELAR targets both users that use cloud infrastructures to deploy and operate their applications and cloud IaaS operators that want to offer extra features to their customers. Cloud users find it hard to detect, deploy, utilize the correct amount of needed resources in order to offer services to their end users, whereas IaaS operators would like to offer more advanced services than simple resource provisioning. Over-provisioning is costly, whereas under-provisioning is not efficient, especially in a dynamic setting where the workload changes constantly during time. Internet serving applications built on top of IaaS clouds such as sites, services, etc, typically encounter these problems. CELAR targets the problem by offering a fully automated fine grained elasticity mechanism that can be utilized by cloud applications to adjust their resources according to user-defined policies when resources need changing. The project's results are open-sourced and offered through a simple installation package that can be deployed on top of any IaaS cloud.

How will the solution/service benefit the end-user?

The CELAR software can be utilized by end users to deploy, manage, monitor and elastically scale their applications on top of any openstack compliant cloud stack, such as Amazon, Rackspace, etc. Users will be able to define complex elasticity rules (i.e., to dictate when and how the system will scale its resources) utilizing CELAR's powerful decision making module by taking into account application specific KPIs or generic monitoring metrics. Apart from legacy coarse-grained elasticity actions, such as virtual machine additions or removals that are already supported by other cloud vendors, CELAR offers fine-grained elasticity actions (i.e., specific resource scaling such as CPUs, RAM, hard disc space, etc), in order to take more informed decisions and meet user defined SLAs. All these features can be deployed in any openstack compliant IaaS cloud, thus avoiding vendor lock-in in a public cloud setting, or enabling in-house elasticity management in a private cloud setting.

Website: <http://www.celarcloud.eu/>

Project start date 01/10/2012

Project finish date 31/12/2015

What user need or pain point is being addressed?

Cloud providers theoretically offer their customers unlimited resources for their applications on an on-demand basis. However, scalability is not only determined by the available resources, but also by how the control and data flow of the application or service is designed and implemented. Implementations that do not consider their effects can either lead to low performance (under-provisioning, resulting in high response times or low throughput) or high costs (over-provisioning, caused by low utilisation of resources).

How will the solution/service benefit the end-user?

CloudScale provides an engineering approach for building scalable cloud applications and services. Our objectives are to:

1. Make cloud systems scalable by design so that they can exploit the elasticity of the cloud, as well as maintaining and also improving scalability during system evolution. At the same time, a minimum amount of computational resources shall be used.
2. Enable analysis of scalability of basic and composed services in the cloud.
3. Ensure industrial relevance and uptake of the CloudScale results so that scalability becomes less of a problem for cloud systems.

CloudScale enables the modelling of design alternatives and the analysis of their effect on scalability and cost. Best practices for scalability further guide the design process.

The engineering approach for scalable applications and services will enable small and medium enterprises as well as large players to fully benefit from the cloud paradigm by building scalable and cost-efficient applications and services based on state-of-the-art cloud technology. Furthermore, the engineering approach reduces risks as well as costs for companies newly entering the cloud market.

Website: <http://www.cloudscale-project.eu/>



CloudSpaces - Open Service Platform for the next generation of Personal Clouds

Project start date 01/09/2012

Project finish date 30/09/2015

What user need or pain point is being addressed?

CloudSpaces aims to protect personal data of users and organisations stored and synchronized in Personal Clouds. Nowadays, Personal Clouds like Dropbox, Google Drive or Box among others have complete access to sensitive data.

The project is targeting mainly organizations that need to protect sensitive data like public institutions, Governments, Universities, Hospitals or even companies. Another important stakeholder are European infrastructure providers willing to offer secure Personal Clouds to users and institutions.

The major problem is that existing solutions do not ensure the privacy of the information stored in the Personal Cloud provider. Furthermore, there are no solutions offering privacy-aware data sharing in a granular way to third-parties.

Another key challenge is to offer open source and scalable solutions that offer multi-device and multi-OS synchronization and sharing of information to a large number of users in a organization.

Here are some examples:

Cybersecurity: Can your Personal Cloud offer a zero-knowledge model to users ? Can your Personal Cloud enable fine-grained sharing with other users, organizations or third-party applications while protecting user's data ? Does your Personal Cloud provide indicators about risks involved sharing sensitive data ?

How will the solution/service benefit the end-user?

CloudSpaces has created StackSync: an open source Personal Cloud for organizations.

StackSync is a secure Personal Cloud that ensures zero-knowledge for the Cloud provider and enable users to retake control of their information. It offers advanced encryption services like client-side encryption and secure sharing using Attribute Based Encryption. The StackSync privacy-aware sharing component gives the end user indicators about the risks that are posed by sharing data, and also proposes and implements policies that can be adopted to mitigate those risks.

There is no other open Personal Cloud in the market offering the security and privacy guarantees to user's data. Since it is open source, it opens up many opportunities to Infrastructure providers and organizations. Furthermore, it includes open APIs that permit to share information with other Personal Clouds and thus avoiding vendor lock-in.

Website: <http://cloudspaces.eu/>

Project start date 01/11/2012

Project finish date 31/10/2015



What user need or pain point is being addressed?

COMPOSE addresses the complex challenges of developing an IoT application. COMPOSE provides a complete ecosystem for the developer, including the different stages of application development lifecycle, from design time through deployment and run-time within a secure cloud environment. The developer only needs to understand the domain of his intended application, rather than worrying about the entire complex HW / SW stack required to properly run the IoT application. The ultimate goal is to simplify IoT application creation thus contributing to unleashing the huge full potential of this market. COMPOSE focuses in particular on opening the door to small and medium players in this field.

COMPOSE presents a user-friendly portal for developers through which all interactions with the platform are made. The platform offers connectivity to IoT devices accompanied by advanced data management capabilities, including real-time data processing capabilities. Thanks to COMPOSE, the developer no longer has to start from scratch. Rather he or she is presented with discovery capabilities to locate relevant building blocks which exist within the platform and that can be readily used to create a new application. Finally, with the click of a button the designed application is automatically deployed to a secure cloud environment.

How will the solution/service benefit the end-user?

So who's going to gain the most from the COMPOSE platform? Well firstly it's the IoT application developer. Once he starts using the COMPOSE platform it will be much quicker and easier for him to realize his dream application. What's more, he'll not even need a deep knowledge of technological aspects which are maybe out of his scope of knowledge and interest. Nor will he have to worry about things like making sure that the device's data is available; that actuators can be called; that the application is deployed in a cloud environment; that it's secure and that data access policies are followed, and so on. As a direct consequence our IoT application developer will be less stressed and will definitely drink less coffee.

Instead of worrying and wasting time on all these aspects our developer can simply concentrate on the business logic he'd like to apply to create his novel application. Seeing the ease and speed which COMPOSE gives, he'll probably create other such applications while he's at it. As an indirect consequence additional stakeholders, such as city councils, will be able to rapidly get or create new applications that will make the European population happier. And this is what we're here for.

Website: <http://www.compose-project.eu/>



Project start date 01/10/2012

Project finish date 30/09/2015

What user need or pain point is being addressed?

Only the biggest information technology players, e.g., Google or Amazon, have access to the necessary infrastructure for storing and processing data available on the Web. Small and medium companies (SMEs) have no other choice than relying on such companies with dedicated data centers to provide them the necessary resources to store and mine public data. The LEADS project aims at developing a set of solutions to change this state of the Cloud business. To that end, LEADS provides a novel Data-as-a-Service framework that makes storing, processing, and querying public data available to almost every size of business or organization in an elastic, low-cost and energy-efficient manner.

How will the solution/service benefit the end-user?

LEADS fosters the development of data-driven companies. To that end, the LEADS platform provides a long chain of promising value-added services.

1. The framework allows SMEs to aggregate their storage and processing capabilities.
2. SMEs can collaboratively crawl the web, resulting in an index similar to those of commercial and non-commercial search engines.
3. On top of this index, either on their own infrastructure or using the aggregated resources available in LEADS, SMEs can process public data and extract meaningful information.
4. Large-scale analyses of collected data can run online and concurrently to the gathering in a stream-oriented manner, as well as offline using classical Big Data oriented processing platform (e.g., map-reduce).
5. The LEADS platform stores both the results of the collaborative effort to crawl public data, and the outputs of private computations.
6. SMEs may re-sell their private results to others using the LEADS Data-as-a-Service framework as a marketplace.
7. Companies and organizations may also use this marketplace to sell and buy business relevant data from and to other entities.

Website: <http://www.leads-project.eu/>

MARKOS - The MARKET for Open Source - An Intelligent Virtual Open Source Marketplace

Project start date 01/10/2012

Project finish date 30/03/2015



What user need or pain point is being addressed?

In Europe, the Open Source Software (OSS) industry frees software users and developers from technological barriers, allowing them to explore, compare, integrate and customise OSS solutions efficiently and systematically, while avoiding license conflicts and infringements.

The MARKOS Prototype is a real OSS facilitating solution. It provides an integrated view of OSS projects available on public forges together with semantic query features, with a focus on the functional, structural and licensing aspects of software code. It supports users in:

- » Searching for open-source projects and software components,
- » Checking the compatibility of software licences,
- » Obtaining information about software dependencies.

Thus, MARKOS frees software users and developers from certain risks and difficulties they face every day:

- » Having to check all the license types underlying an OSS, since improperly managed Open Source code could result in copyright infringement with legal and financial consequences;
- » Trying to finding the right OSS component required, and also finding correlations among different OSS in a variety of forges and other online sources;
- » Spending large amounts of time understanding the OSS project structure and dependencies, in order to further use it as components for new Open Source or commercial solutions.

How will the solution/service benefit the end-user?

The MARKOS Prototype (<http://demo.markosproject.eu/>) offers innovative services, which make it stand out from others already available on the market.

MARKOS users can now:

- » Index and analyse 1,540 Java, JavaScript and HTML OSS projects (a total of 3,885 releases) available on forges such as Apache, SourceForge and GitHub.
- » Easily search for specific components based on their type, programming language, keywords and on different kind of license compatibility constraints. They can also use smart syntax for more advanced users and queries.
- » Browse and navigate source code and dependencies across projects using an IDE-like web user interface.
- » Coarse automatic license analyses, to quickly work out whether a project risks infringing the license of copied code or used libraries.
- » Analyse in detail license compatibility issues. This provides developers and lawyers with tools to create arguments from semantic models of copyright law, as well as visualizing these arguments using argument maps.
- » Get Notifications. Firstly, downstream which allows a user to subscribe to specific events about a particular project. Secondly, upstream which lets users submit patches to a software release, with the project manager being notified about the availability of the new patches
- » Get RDF descriptions through the use of linked data API and to perform complex SPARQL queries for searching code information.

Website: <http://www.markosproject.eu/>



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

Project start date 01/10/2012

Project finish date 30/09/2015

What user need or pain point is being addressed?

Let's take an example. MODAFIN is a software development company specialised in IT applications for financial services. Its main product line is a proprietary solution for stock market operations, cash administration, and lending management. The most profitable activities are software customisation and life-cycle management. MODAFIN currently fulfils the quality requirements with a capacity management consultancy team following the application life-cycle at the customer's site. However, the MODAFIN customers require an increased assured quality, such as high-availability for real-time calculations during market hours, scalability and low operational costs for batch analytic workloads running after-hours. During the day, calculation engines are expected to ramp-up computing power when the stock market gets volatile. At night, customers want to run their batch analytic workloads at the cheapest operational costs of on-spot instances. Moreover, some customer applications are collecting and processing stock market data using Cloud datastore services. Customers also want to cut costs spent in consultancy services for life-cycle management. MODAFIN needs to address these new requirements, through integration of the existing legacy codes with customer Cloud services, replicating on several Clouds to provide quality assurance, have a flexible architecture which is adaptable to new Cloud offers, and support a life-cycle management deployed on multi-Clouds.

How will the solution/service benefit the end-user?

The MODAClouds solution is exactly what MODAFIN needs. The MODAClouds approach allows the design, development, and re-engineering of existing components into software modules that operate directly on multiple Clouds. It allows one to include a description of MODAFIN's legacy software in the design model and automatically generates interfaces and connectors for application integration with customer systems for all of the supported target platforms. A Decision Support System is used to determine which Cloud to adopt for hosting the different components of new solutions, comparing costs, risks, and analysing non-functional characteristics for each alternative provider. This gives users the right information to make informed decisions and trust cloud solutions. Furthermore, a run-time management API (independent of Cloud vendors) is beneficial to natively implement Cloud-to-Cloud migration and Multi-Cloud load-balancing. Workloads can be run on different platforms depending on the customer's preferred datastore and hence enjoy the performance benefits of customer data locality. The run-time operation can easily be integrated inside MODAFIN's software development process through a feedback system that automatically provides recommendation on the best design and deployment patterns for a new custom module. This makes it adaptable to context and requirement changes.

Website: <http://www.modaclouds.eu/>

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



Project start date 01/10/2012

Project finish date 30/06/2015

What user need or pain point is being addressed?

OPENi addresses the issue of user personal data privacy, allowing users to decide which aspects of their personal data they are prepared to share by giving them the option of fine grained authorisation and access control. Through the OPENi platform we are aiming to alter the dynamics of user control over their personal data. We are addressing the lack of interoperability between cloud based services and enabling applications to access and use a broad spectrum of existing cloud based functionality and content, consistently across different devices and platforms we are enabling a more user-centric application experience. The platform will incorporate an open framework that will be capable of interoperating with any cloud-based service, abstracting the integration challenges to a single open standard without losing any service features.

How will the solution/service benefit the end-user?

The operation of the personal information economy is one which is coming under increasing public scrutiny and regulation. As digital services become increasingly personalised and embedded in our everyday lives, the importance of effectively managing our digital identities becomes more apparent. At present the balance of power favours the service providers and application developers; yet systems which offer user control over where, how and with whom their personal information is shared will ultimately tip the scales towards the user. Our platform enables fine grained sharing, federated user identity, and greater control over personal information OPENi is all about empowering the user to manage and control their digital identity.

Website: <http://www.openi-ict.eu/>

What user need or pain point is being addressed?

OSSMETER targets 4 types of users:

1. Developers and Project Managers who are responsible for deciding which of the vast and overlapping open source technologies publicly available should be adopted for internal use or for inclusion in new products or delivering new services;
2. Developers of publicly available open source software technologies;
3. Quality managers within software development organisations;
4. Sponsor organisations and supporters of projects that produce open source software technologies.

OSSMETER addresses the following pain points

1. Huge number of open source software technologies available over the web from general purpose forges (e.g. SourceForge), to focused forges (Eclipse.org) there are often many choices when trying to decide which open source technology should be adopted.
2. Adopting open source software is an investment and commitment with respect to new products or internal tools so understanding which open source alternative has the best quality is key factor in decision making.
3. Quality of code is critical, but equally important for open source technologies is understanding the quality of the surrounding communities where contributions to the evolution of the open source technologies, the level of activity in the forums and message boards, and responsiveness of the development teams to fixing responding issues or responding to queries play an important part in deciding which technologies to adopt.

How will the solution/service benefit the end-user?

OSSMETER provides a platform and over 200 independent and reliable metrics for evaluating open source technologies. It measures both the quality of the code and the quality and level of activity of the surrounding support communities.

OSSMETER enables the comparison of different open source technologies both as they stand today and also their history and evolution. It can be used for evaluating and monitoring publicly available open source technologies or downloaded and used as a powerful and customisable toolset for managing and monitoring the quality of in-house developed software and support services.

OSSMETER provides on-going monitoring of the quality of open source projects and gives immediate warnings if open source technologies already adopted for use in products or services experience changes in code quality, number of issues reported or drop in the level of support occurs.

OSSMETER uses advanced natural language processing to extract the sentiment of those interacting within the community surrounding an open source project including the manner in which issues are addressed and the overall perceptions of those already using an open source technology. In addition it provides special features to aggregate metrics into meaningful summaries that give an immediate view of key open source quality aspects.

Website: <http://www.ossmeter.org/>



What user need or pain point is being addressed?

Developing once and deploying on different Clouds is not the reality today for software developers and system administrators. Methods and tools are critically missing for supporting the development, the deployment and the execution of Cloud applications.

There exist many different Cloud platforms, based on heterogeneous technology, on non-standard architectures and APIs, thus often implying vendor lock-in and resulting in an interdependence between the client application and the Cloud platform.

Businesses also need technology-neutral methods and tools to allow organisations to scale out from their private Cloud to public Clouds whenever they need it, without a technical chasm between. This cannot easily be achieved today.

How will the solution/service benefit the end-user?

PaaSage is an open and integrated platform, to support both deployment and design of Cloud applications, together with an accompanying methodology that allows model-based development, configuration, optimisation, and deployment of applications independently of the existing underlying Cloud infrastructure technology. Specifically PaaSage is developing CAMEL (Cloud Application Modelling and Execution Language, a series of domain specific modelling languages), an Integrated Development Environment, execution-level Cloud-specific mappers and interfaces, a metadata database collecting models and historical information related to the applications components and external services, such as execution histories, linked with a dedicated social network.

PaaSage software is released under the open source MPL2.0 license and an open community of developers and users is built together under the OW2 (www.ow2.org) umbrella.

PaaSage allows for application deployment on multi-cloud platforms (of any kind). It selects the optimal deployment solution, taking into account the characteristics of the available Cloud platforms, the data to be used and the end-user preferences or restrictions such as price, location, availability, etc. PaaSage supports Cloud bursting and dynamic reallocation of application or its artefacts to other Cloud platforms.

Website: <http://www.paasage.eu/>

What user need or pain point is being addressed?

The open source projects Europe platform is targeting users that want to:

1. promote the use of FLOSS in their EU project,
2. Access free software collaboration tools,
3. Want visibility of their projects software results.

Open source projects Europe offers unique features to EC research projects hosting both open and closed parts of projects in its operating structure. This is to cater for the large number of EC ICT projects have parts of their project which are essentially private to the originating project consortium and also parts that are shared, open collaborative areas for their wider community.

How will the solution/service benefit the end-user?

The OSP Europe platform has been enhanced in order to give users metrics data so that users can get insights into the software contributions and effort given to a project over time. It can also highlight important ongoing trends for the number of contributions and developers for your project.

The OSSEVAL tool for FLOSS evaluation has also been added to the platform. This facilitates procurement and evaluation by third parties, and is a unique technical feature of the OSP Europe platform. It also qualifies and selects open source software demand, evaluating open source specific metrics, ranging from the sustainability of its community to the ease to contribute to the code.

SOS Open Source methodology provides tools useful for determining the risk associated with the usage of specific open source software components. The analysis returns a set of open source candidates that are robust (stable, mature and backed by a viable community), supported (either by a community or vendors) and grant evolvability (readable and maintainable code).

Finally a unique value proposition that the OSP Europe platform has is access to business experts and legal experts in the area of open source software. This can provide mentorship to software developers on the pitfalls of using and distributing open source software.

Website: <http://www.ict-prose.eu/>

PROWESS - Property-based testing for web services

Project start date 01/10/2012

Project finish date 30/09/2015



What user need or pain point is being addressed?

Prowess is addressing the need for effective and efficient testing automation techniques and tools focusing, in particular, on web services-based systems and cloud computing. The possibilities associated with cloud computing provide instant scalability, flexibility and availability for testing on demand with no upfront investment. This provides the industry with a perfect opportunity to utilise powerful high volume automated testing solutions.

The global testing cloud marketplace will allow for the joint collaboration of leading test specialists following industry best practice. This enables firms of all sizes to access the latest test approaches and methodologies whilst providing a unified platform for domain experts to represent business processes and user story acceptance criteria in natural language with content sensitive business validation. Within PROWESS we aim to reach out to web developers that do use Erlang for their backend, as well as web developers unaware of Erlang.

How will the solution/service benefit the end-user?

The Quivq tool QuickCheck is a market leader in the area of automated testing. PROWESS will deliver an integrated set of tools and methods that contribute to the overall goal of strengthen the QuickCheck ecosystem to effectively test internet services. Within the project, testing is seen as a service. This service includes creating models, wrapping these models in a simple interface and training people to use this dedicated version of QuickCheck for testing. One goal is to use the results of this project to offer such services for web developers; most probably as a web service in itself.

Website: <http://www.prowessproject.eu/>



RISCOSS – Managing Risk and Costs in Open Source Software Adoption

Project start date 01/11/2012

Project finish date 31/10/2015

What user need or pain point is being addressed?

The RISCOSS project aims at proposing a tool supported methodology to evaluate the risks and costs intrinsic to the adoption of Open Source Software (OSS) in companies of different sizes, public administrations and OSS communities.

The main problem in this context is that of identifying and deciding the exposure and impact of the risks on the business goals of the adopting organization. Two exemplar risks are: licensing risks and maintenance risks.

Licensing risk is one of the most critical risks. It refers to the capacity of an adopting organization to use sets of OSS components whose licenses allow releasing the product using an OSS license that does not negatively impact on the business objectives of the organization. This may, for example, limit product sales.

Maintenance risks on the other hand, refer to difficulties organisations may have in being supported by OSS communities in maintaining OSS components used in a given product. In particular, organizational costs related to the maintenance and evolution of a given OSS component included in the product and generating possible misalignments between the internal component and the work of the community.

How will the solution/service benefit the end-user?

The main benefit for an organization adopting OSS is related to the increasing of awareness of possible risks in doing so. RISCOSS, identifies the two dimensions of OSS-related risks. On the one hand, the technical risk directly related to the adoption of OSS, so influencing the way an organization produces software. On the other hand, the business/strategic risks that are related to the overall business dimensions of the organization. These two aspects should allow organizations to better decide how to avoid or mitigate possible risks at different levels of the organization.

For licensing and maintenance, RISCOSS proposes licensing models that let users evaluate the exposure of an organization to this risk and evaluate the possible impacts on the goals and assets of the organization. This allows the organization to decide if a particular set of components is compliant or, as a mitigation strategy, whether the set should be changed with other components which have the same functionalities. In the case of maintenance, RISCOSS evaluates the exposure and impact of the risk on the organization allowing for possible mitigation activities in the case maintenance problems arises from OSS.

Website: <http://www.riscoss.eu/>

U-QASAR – Universal Quality Assurance & Control Services for Internet Applications with Volatile Requirements and Contexts

Project start date 01/10/2012

Project finish date 30/09/2015



What user need or pain point is being addressed?

U-QASAR is addressing the need of obtaining an objective measurement of quality for software (SW) development processes and their resulting products. It is targeting all actors involved in the SW quality process, like CEOs, CTOs, software project managers, quality managers or different categories of software engineers.

Those actors have many different perspectives in quality. Different levels of granularity are required so they can be on the same page. Often users don't know what and where to look at in order to measure the quality of their SW projects and processes.

Data about SW quality of different projects in a company normally come from several distinct and heterogeneous sources (tools, outputs of the different stages of processes, documents) and it is a hard work to compile that huge amount of information and put it together in a document report for managers, decision makers, etc. In many cases the end-users end up having to do it manually.

It is also demanding to follow-up quality during the whole SW lifecycle and to identify areas where actions are needed. For example, too many bugs in the SW product or customer' dissatisfaction etc.

How will the solution/service benefit the end-user?

U-QASAR's Methodology allows companies to get a shared view of SW quality so that everyone involved (CEOs, project managers, quality managers, developers, testers...) can collaboratively agree on key quality goals. The result is a commitment from everyone to reach them.

U-QASAR also allows the easy creation of quality models for the different types of projects and development paradigms (V, waterfall, RUP, Agile...). These can even be done from scratch and can be adapted to the requirements of existing quality models and standards, like CMMi or ISO 9126 (25000). It also provides patterns to facilitate this.

The platform allows the integration of diverse tools and gathers and aggregates data from them which can be presented to different stakeholders. This offers an objective overview of quality with different levels of abstraction. It is also easy to extend the tool for integrating new data sources.


Users, whether they are experts or not, can easily find out about and follow-up the overall status and quality of different project items. Analytics and reports of the whole SW development lifecycle can also be obtained. This lets CTOs make decisions based on the history data stored in the platform.

U-QASAR's business model allows the solution being accessible to a wide range of users such as SMEs, freelancers and large companies.

Website: <http://www.uqasar.eu/>

Software Engineering, Services & Cloud Computing

Market ready in 2015- 2017



ASCETiC - Adapting Service lifeCycle towards Efficient Clouds

Project start date 01/10/2013

Project finish date 30/09/2016



What user need or pain point is being addressed?

The ASCETiC project focuses on issues of energy efficient computing, specifically on design, construction, deployment and operation of Cloud services.

For software developers, it proposes novel methods and development tools to support them in monitoring and optimizing (minimizing) the energy consumption resulting from developing and deploying software in Cloud environments.

For Cloud providers, ASCETiC can be a complement to their offering in order to attract users while improving their corporate image as a means of offering differentiation and adhering to green legislations. At the same time it provides internal benefits decreasing energy consumption and as a result associated costs. As the green-aware market segment grows, this reduction in energy consumption can also help generate sales from niche market segments making it stand out from other competitors .

For end users, ASCETiC helps to improve their Corporate Social Responsibility, as well as guaranteeing the adhesion to green legislations regarding the optimization of energy consumptions and greenhouse gas emissions.

How will the solution/service benefit the end-user?

ASCETiC's primary goal is to characterize the factors which affect energy efficiency in software development, deployment and operation. The 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 measurements of energy requirements into the design and development process for software to be executed on a Cloud platform.

For this reason, ASCETiC is providing the ASCETiC Toolbox, an Open Source Cloud stack that takes energy awareness into account, aiming to reduce carbon footprint and minimize energy consumption.

The ASCETiC Toolbox brings an innovative solution to the market, as it focuses on the whole lifecycle of a service operation, not only on parts of it. It offers pricing models based on energy consumption, self-adaptation capabilities and the possibility to reuse its functionality on a single installation.

Most of the work performed for reducing energy consumption is nowadays concentrated at data centre (building) or infrastructure level. ASCETiC is going one step further on this offering, as it is focused at the software level, and it is a complementary solution for those organizations that have already implemented improvements on the above mentioned levels.

Website: <http://www.ascetic.eu/>



CACTOS – Context-Aware Cloud Topology Optimisation and Simulation

Project start date 01/10/2013

Project finish date 30/09/2016

Context-Aware Cloud Topology
Optimisation and Simulation

What user need or pain point is being addressed?

CACTOS addresses the challenge of cloud data centre operators that hardware infrastructure becomes increasingly heterogeneous to meet energy efficiency constraints. This makes the challenge of optimized VM placement and performing optimization within the Data Centre much more complex.

How will the solution/service benefit the end-user?

The overarching goal of the CACTOS project is to provide autonomic and interactive tools for cloud data centre operators to support and improve the efficiency of their 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.

Website: <http://www.cactosp7.eu/>

CloudCatalyst – Reenergize productivity, efficiency and competitiveness of European economy through Cloud Computing

Project start date 01/10/2013

Project finish date 31/10/2015



What user need or pain point is being addressed?

Adopting the cloud, or in other words migrating a business to the cloud, is not necessarily an easy task. Failure to properly prepare business for cloud migration can end up in costly results such as company-wide disorganization, prolonged downtime, loss of productivity and even loss of data.

CloudCatalyst provides comprehensive cloud computing support tools to businesses that are either looking to migrate their business to the cloud or have recently done so. The project will share expert advice on how to safely and efficiently execute a cloud migration, and more specifically, how to avoid the most common (and avoidable) mistakes that most companies tend to make when migrating to the cloud.

How will the solution/service benefit the end-user?

CloudCatalyst offers the following benefits to the end-users:

1. Startups and SMEs (vendor / cloud supply perspective): Guidance to help companies finding gaps that may be new business opportunities to effectively target the cloud market
2. Startups and SMEs (end user / cloud demand perspective): Guidance to help companies that will be “natural” adopters of cloud computing addressing technical challenges and defining their cloud strategy
3. Incubators and entities providing support to new market entrants: Guidance to help their companies to define their IT strategy.

Website: <http://www.cloudcatalyst.eu/>



Project start date 01/07/2013

Project finish date 30/06/2015

What user need or pain point is being addressed?

CloudingSMEs takes a comprehensive approach to the use of cloud computing by SMEs, which considers both the supply side (i.e. results and endeavors of ICT SMEs with expertise on cloud computing) and the demand side (i.e. the adoption and use of cloud computing technologies by SMEs). SMEs represent 99.8% of the European enterprises. 92 % employ less than 10 people – the so-called micro-enterprises. SMEs play a crucial role for economic and social stability at local and regional level, where SMEs are the “basic fibre” by providing goods and services for the daily life. CloudingSMEs is motivated by the need to lower the barriers associated with the adoption of cloud based solutions by SMEs, while also considering the need to facilitate SME providers of cloud solutions with the aim of sustaining and strengthening their positions in a very competitive cloud computing market. The main motivating factors behind CloudingSMEs are:

- » The unsatisfactory adoption rates of cloud computing technologies within SME communities, despite relevant benefits.
- » The fact that SMEs need practical support in their cloud computing related decisions and initiatives.
- » The potential of cloud computing technologies to boost competitiveness help SMEs grow and the need to unveil and boost this potential.
- » The lack of market solutions focused on micro SME needs.
- » The need for consolidating SME requirements as a means of achieving/boosting interoperability, harmonisation and economies of scale.

How will the solution/service benefit the end-user?

CloudingSMEs is providing support to SMEs that aim to adopt and/or exploit cloud computing, through a practical toolbox (i.e. the CloudingSMEs toolbox) that will facilitate their cloud-related decisions. CloudingSMEs is committed to supporting SMEs in contractual, legal and cost-related issues, including privacy and security issues. CloudingSMEs has already implemented, evaluated and fine-tuned on number of interactive tools, including:

- » Cloud Security Scorecard
- » TCO/ROI Calculator
- » Strategic Considerations Scorecard
- » Contents of Contracts and SLAs
- » Cloud Standards Catalog
- » SLA Guide
- » Cloud Services and Solutions Providers Searchable Catalogue

CloudingSMEs will also establish clear policy recommendations for the benefit of all European SMEs.

Cloud Computing can help SMEs lower their IT costs by outsourcing IT services and obviating the need for in-house infrastructures. Furthermore, the pay-as-you-go nature of the cloud ensures that SMEs only have to pay for what they use. In this way, SMEs can transform the capital expenses (CAPEX) associated with IT infrastructure acquisition to operational expenses (OPEX) associated with usage-based billing.

Software-based services on a global scale greatly reduce requirements for capital expenditure and lower operating costs (including energy costs). For SMEs this can be an element of their competitiveness (since it allows them to operate based on a reduced IT budget and with higher efficiencies).

Website: <http://www.cloudingsmes.eu/>

CloudWATCH - A European cloud observatory supporting cloud policies, standard profiles and services

Project start date 01/09/2013

Project finish date 31/08/2015



What user need or pain point is being addressed?

Small- and medium-sized enterprises (SMEs) are the backbone of the European economy. But when it comes to adopting new services or implementing information security systems, SMEs face significant challenges. They typically have restricted budgets, limited resources and limited expertise in information security. They are increasingly less likely to adopt cloud because of security concerns, complex terminology, fear of vendor lock in and lack of transparency.

There is also a growing demand from customers for freedom of choice and increased control. Security and privacy certifications and attestations have been identified as one of most effective and efficient means to increase the level of trust in cloud services and stimulate their adoption.

So how can we best facilitate SMEs in migrating to the cloud and help them innovate in the global marketplace where new business value creation is becoming increasingly important? How do we bring together small firms and cloud service providers in a way that facilitates information in a neutral and objective way?

More knowledge for making the right strategic decisions

CloudWATCH is creating new services and tools to help navigate SMEs to the cloud confident that they have at their fingertips all the practical, security and legal information they need to make the right decisions [3].

Online tools for SMEs: CloudWATCH has already led the way with the launch of the European Cloud Scout, a novel interactive tool for SME managers in Europe with low-threshold information on legal, organisational and technical requirements for the successful use of cloud services [4]. Several country and language specific versions already available online and more will follow.

CloudWATCH is now poised to launch a highly practical and insightful web application for SMEs. Through a portfolio of cloud service providers and their services which will enable SMEs to get a much clearer picture of cloud services, and match current offers on the market with their specific business needs and processes. The tool will also provide information, tips and check-lists on key pain points for SMEs: legal issues, contracts and SLAs; certification and compliance; interoperability and performance.

Commoditising cloud

Currently, albeit its already massive scale, Cloud adoption can still be seen as first mover market from a risk management and service procurement point of view. To become a true market at scale, Cloud services need to turn into commodity (for large scale B2B economics) as well as a utility (for B2C and end-user consummation) just as electricity is. This is achievable only through agreed information models and access interfaces commonly agreed across market participants and would be entirely impossible without agreed standards in the market. Standards provide a level playing field for market participants to compete on service delivery, service management, customer satisfaction and service pricing for the benefit of the customer. CloudWatch aims at driving the Cloud market towards Cloud services being evolved into and treated as commodities and utilities through promoting uptake of publicly defined open standards on technology, legal and governance processes.

More choice for customers: CloudWATCH is developing a set of common standards profiles based analysis of clustering of 55 cloud projects and scored on the NIST criteria for cloud computing. A profile on a standard clarifies in an unambiguous way how a standard has to be interpreted, explaining how to implement it based on a specific use case. Based on a portfolio of European and international use cases covering technical, policy and legal requirements, such as service level agreement management CloudWATCH will develop and test a set of common standards profiles around the federation of cloud services which can be picked up and used easily and thus driving the Cloud market into a state of commodity and utility.

More trust for customers: CloudWATCH has analysed all currently available cloud security certification schemes and through a set of recommendations is providing guidance for cloud service customers, especially public administrations and small and medium companies, cloud service providers and policy makers in their evaluation of possible options for "certifying" the level of security and privacy of cloud services.

Website: <http://www.cloudwatchhub.eu/>

Project start date 01/11/2013

Project finish date 31/10/2016

What user need or pain point is being addressed?

Cloud computing's initial success and impact in the IT sector is undeniable, but as its service offerings continue to evolve, both cloud infrastructure providers and application developers alike look towards a significant improvement in terms of performance, reliability and agility.

Cloud providers (IaaS) look for ways to better optimize their infrastructure and drive down OpEx. Meanwhile, their customers, cloud application developers (SaaS), seek quicker time-to-market, shortened development/maintenance cycles and increased QoS/QoE for the end-user.

Holistic and cloud-specific engineering methods are needed to address the pain points of both stakeholders. Rising trends such as cloud-based monitoring and DevOps seek to better facilitate interactions between operations with development, but current solutions fall short in integrating such services between the infrastructure and application layers of the cloud.

How will the solution/service benefit the end-user?

CloudWave (www.cloudwave-fp7.eu) enables a next generation of cloud infrastructure operations and agile development by dynamically adapting cloud services to their environment, improving service quality and optimizing resource utilization. Inspired by the emerging concept of DevOps for facilitating application development, CloudWave empowers cloud infrastructure providers (IaaS) and their hosted applications developers (SaaS) to transparently collaborate to obtain high levels of service at lower costs.

Execution Analytics: CloudWave improves existing cloud monitoring solutions with a more holistic and efficient approach towards IaaS and SaaS services. Unified monitoring consolidates infrastructure vs. application data, as well as virtual infrastructure vs. physical hardware. Programmable monitoring enhances filtering and delivery of data for analysis, allowing for better management of cloud-based resources.

Coordinated Adaption: CloudWave enables reconfiguration of the infrastructure and application in real-time to compensate for a variety of performance factors, resulting in an increasingly resilient, automated and optimized cloud deployment.

Feedback-driven Development: CloudWave advances current DevOps solutions with developer-oriented data based on its innovative monitoring. It mixes automation (coordinated adaptation) with customizable feedback for improved agile development, resulting in quicker time-to-market, shortened maintenance cycles and more reliable cloud applications for their end-user customers.

Website: <http://cloudwave-fp7.eu/>

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

Project start date 01/04/2013

Project finish date 30/04/2016



What user need or pain point is being addressed?

Let us outline four scenarios where users can benefit from the ClouT solution.

Scenario one

I am a city municipality who needs to manage weather risks and emergencies and inform my citizens about specific emergency alarms, about environmental risks and on what to do in critical emergency situations. I need to meet their needs by giving them relevant and timely information. The problem is that I need to avoid capital expenses and I need to be able to store historical sensor data and build statistical information in an optimal way.

My needs are met through ClouT because I do not have to buy expensive machines, for one. Moreover, I have access to cloud storage where I can store and receive data, benefitting from reliability, scalability, and elasticity aspects. It represents fast and secure access to my event data.

Scenario two

I am a city municipality which is looking to improve the quality of life of my citizens by offering better event and mobility management. I need to meet their needs by offering real time information and suggestions on the best options for reaching a specific destination. The problem is I need to be able to gather and process data more efficiently, optimize city council service provision and save on costs.

My needs are met through ClouT because I can benefit from the cloud storage and receive a more balanced and faster service, even in regards to scalability. Moreover, I can gather and manage information coming from different sources better. Furthermore, I can build enhanced services, such as integrating new functionalities.

Scenario three

I am a city municipality which needs to motivate my elderly citizens by providing them with interesting suggestions on what to do, such as nice places to visit or interesting events to attend. This in turn will encourage the elderly to go out more often, take longer walks, and ultimately promote a healthy lifestyle. I need to meet their needs through innovative social network application that persuades them to go out, benefit from the suggested activity, share the positive feedback and in turn persuade others to do the same. The problem is I need to use a platform that provides me with the possibility of matching participatory data with open city data, as well as manage and deliver information in a timely and effective manner.


My needs are met through ClouT because I can try my new idea without an important initial investment, and benefit from the rapid development via PaaS. Furthermore, I can gather and manage information coming from different sources better and take advantage of the scalability aspect.

Scenario four

I am a city municipality which is looking to provide city information, such as events and traffic in a timely manner to visiting tourists and citizens. Moreover, I am looking to provide important information about emergency situations. I also need to create a way to encourage tourists to sightsee and motivate them to continue to stay in the area. Additionally, I need to help revitalize the shops and restaurants in the area. I am looking to meet their needs by creating innovative applications aimed at influencing tourist behaviour through incentives such as discounts at local shops. The problem is I need to successfully leverage the IoT sensors and actuators which are locally and remotely located. Furthermore, I need to manage disaster data which needs to be reliable.

My needs are met through ClouT because I can gather and manage information from various sources better and deliver this to my tourists and citizens in real-time. Moreover, I can benefit from the resiliency ClouT offers making my smart city services reliable in the event of a local disaster.

Website: <http://clout-project.eu/>



CoherentPaaS

CoherentPaaS - Coherent and Rich PaaS with a Common Programming Model

Project start date 01/10/2013

Project finish date 30/09/2015

What user need or pain point is being addressed?

I am the head of the IT department of a large enterprise. Just like many other companies, we are using a combination of data management solutions, such as SQL databases, NoSQL databases such as document data stores, graph databases, key-value data stores and complex event processing. However, we are encountering some problems as some business actions require us to update multiple data stores. In addition, whenever there is a failure our federation of data stores loses its consistency in dealing with multiple problems with our customers.

What is more, we now have to perform operations across our heterogeneous data stores. This is a big problem for us as when we have to join data, our application developers have to develop the code. Most of the time they are not able to do this. This implies that we also have to move the data from the data stores to our data warehouse in order to be able to perform these queries. This is also painful because the data warehouse is a SQL OLAP database and does not support the graph and document data models for either of the corresponding query languages.

How will the solution/service benefit the end-user?

CoherentPaaS addresses my two pains points. Now we have transactional semantics to update across data stores getting all-or-nothing semantics for business transactions which guarantees our data consistency in the advent of failures. This means no more problems with customers due to data lost during failures!!!!

CoherentPaaS has also successfully solved our pains with making queries across data stores. Now, we can write subqueries in the native query languages, exploiting the full power and efficiency of the underlying data stores, including their data models and query languages/APIs. When we have queries across data stores we write them declaratively in SQL. This SQL observes the resulting sets of the subqueries as temporary SQL tables. With CoherentPaaS, now we do not have to waste any extra time and effort on the heterogeneity of our data stores and we can fully exploit the data.

CoherentPaaS has helped us to solve two major pains points we encounter due to the introduction of multiple data store technologies in the company and has reduced the huge overhead in terms of personnel effort that we previously had to perform to access the data scattered across data stores.

Website: <http://coherentpaas.eu/>

Project start date 01/01/2014

Project finish date 30/06/2016



What user need or pain point is being addressed?

A majority of European mobility-related public and private sector institutions still struggle to take maximum advantage of the cloud to deliver services and data online in a cost-effective and scalable manner. At the same time, developers and SMEs who seek such services and data to create useful and innovative mobility applications, have difficulties in finding them and increasingly depend on disparate service catalogues/marketplaces to procure them.

The European Cloud Marketplace for Intelligent Mobility (ECIM) set out to help 3 key audiences – service and data providers, developers (service re-users), citizens (service users) – to benefit from the ability of cloud computing to deliver mobility apps and services more efficiently and effectively.

- » Service and data providers: Private or public providers who seek to offer their services or data in order to expand their geographic scope and increase their target market.
- » Developers: Including SMEs, who need to quickly and easily access/purchase and use innovative web services anytime, anywhere.
- » Citizens: Who are forced to use different applications, user accounts and payment processes to cover their every-day mobility needs (e.g. parking, transport planning, etc)

How will the solution/service benefit the end-user?

ECIM offers a state-of-the-art cloud platform that will become the European marketplace of choice for city mobility services that are secure, scalable and easy to use.

It will achieve this goal by creating a flexible, cloud-based solution that makes it easier for public and private service providers to publish their services and data, while enabling developers to combine them and re-use them. In summary, ECIM benefits its end users by providing:

- » A pan-European mobility service and data catalogue/marketplace with a common interface for developers to purchase and access them, irrespective of the type of service and geographic location.
- » Common Mobility API to facilitate the easy creation of new mobility applications, by alleviating the need for developers to learn and use different service specifications.
- » An easy-to-use interface where providers can publish their services and make them available through the Common Mobility API.
- » A common payment system to facilitate the interaction and financial transactions between service providers and service re-users.
- » The potential to combine web services and data from different providers into a single application, with seamless experience for the app end user.

Website: <http://www.ecim-cities.eu/>



Project start date 01/10/2013

Project finish date 30/09/2016

What user need or pain point is being addressed?

Cloud computing allows software developers to control resource parameters such as the choice of processors, the amount of memory, storage capacity, and bandwidth. When software is designed based on specific assumptions about the size of data structures, the amount of RAM, and the number of processors, rescaling requires extensive design changes. If software were designed for scalability, the client of a cloud service can not only deploy and run software, but also fully control trade-offs between the incurred cost and the delivered quality-of-service.

Traditionally, deployment happens late in software development. For static deployments, IBM Systems Sciences Institute estimated that a defect which costs one unit (hour, \$, etc) to fix in design, costs 15 units to fix in testing and 100 units in production, without considering delayed time to market, lost revenue, lost customers, and bad public relations. Considering the complexity of virtualized resource management, more significant differences can be expected. Deployment errors, such as the failure to meet a service level agreement (SLA), must be fixed already in design! To make full usage of the cloud, software developers need a methodology that supports deployment modeling at early design stages and permits the detection of deployment errors early and efficiently.

How will the solution/service benefit the end-user?

The goal of Envisage is to move deployment decisions into the design phase of the software development chain without convoluting the design with deployment details. To fully exploit the opportunities of cloud computing, the software developer working with the tool-supported development environment of Envisage uses advanced model-based analysis techniques to devise resource-aware deployment strategies when developing software services. The resulting software will monitor its own load, congestion and client traffic to scale in a timely and cost-efficient manner.

The software design methodology developed in Envisage is based on “early modeling” of virtualized services destined for cloud deployment. The behaviour of virtualized services is specified as an abstract behavioural model in the executable modeling language ABS. ABS models interact with a Cloud API, a highly configurable abstract interface to the provisioning layer. The models make it possible to develop and compare different client-side resource management schemes and observe their impact on cost and performance at design time. Model-based analysis facilitates the detection of deployment errors early and efficiently, helped by software tools, such as simulators, test generators, and static analyzers. The methodology further includes automated support for runtime monitoring of service-level agreements for services deployed on the cloud.

Website: <http://www.envisage-project.eu/>

HEADS - Heterogeneous and Distributed Services for the Future Computing Continuum

Project start date 01/10/2013

Project finish date 30/09/2016



What user need or pain point is being addressed?

The goal is to empower the software and services industry to better take advantage of the opportunities of the future computing continuum and to effectively provide new innovative services that are seamlessly integrated to the physical world making them more pervasive, more robust, more reactive and closer (physically, socially, emotionally, etc.) to their users. We denote such services HD-services. HD-services (Heterogeneous and Distributed services) characterize the class of services or applications within the Future Internet whose logic and value emerges from a set of communicating software components distributed on a heterogeneous computing continuum from clouds to mobile devices, sensors and/or smart-objects. The main objective of the HEAD project is to enable efficient exploitation of the broad diversity of the future computing continuum for rapid service innovation of advanced services, by providing service developers with new agile tool supported software engineering methods enabling short innovation cycles.

How will the solution/service benefit the end-user?

This HEADS integrated development environment (IDE) enables the efficient documentation and formalization of the knowledge of platform experts (e.g., cloud expert mastering advanced frameworks, or embedded system developers able to optimize memory or power consumption on resource-constrained platforms) into reusable forms. The capabilities and peculiarities of each platform should be well-documented and should be leveraged at their maximum instead of being "harmonized" according to their least common features denominator.

The HEADS IDE provides service developers with new abstractions to cost effectively design, validate, deploy and evolve HD-services for the future computing continuum. These abstractions should provide service developers with an explicit access to the diversity of the platforms of the continuum while relieving them from their technical implementation details. This way, developers can focus on the logic of their services and should be guided by a methodology and an IDE to distribute this logic in a way that takes advantage of the whole continuum.

Website: <http://heads-project.eu/>



HTML5Apps - HTML5 for Apps: Closing the Gaps

Project start date 01/10/2013

Project finish date 30/09/2015

What user need or pain point is being addressed?

HTML5Apps is targeting app developers who want to break free of the limitations of today's "native apps" (centralized and proprietary app stores, lack of cross-device support) by using HTML5, CSS and Javascript to develop their applications.

While already powerful, HTML5 cannot be used to fully replace native apps today. This is because it is lacking a number of important functionalities such as

- » support for "installing" an app
- » support for handling of payments
- » rich APIs to interact with devices (e.g. to access bluetooth or NFC functionality)

How will the solution/service benefit the end-user?

With the standards started with the help of the HTML5Apps project, app developers will increasingly be able to use HTML5 for apps that they could not do before, creating an alternative to app stores and avoiding many of their issues (app store approval process, revenue sharing with the app store provider, proprietary development environment, ...). In particular, HTML5 standards on support for installing apps, handling of payments and access to device features such as Bluetooth or NFC are well advanced and first implementations in HTML5 platforms are available.

Website: <http://html5apps-project.eu/>

MONDO – Scalable Modelling and Model Management on the Cloud

Project start date 01/11/2013

Project finish date 30/04/2016



What user need or pain point is being addressed?

Mondo targets software developers and software development managers who have adopted or are planning to adopt model-based engineering (MBE) methods for software development. In particular, developers of large and complex software applications, or software development teams who need to work together to develop a software application or system using MBE. Mondo also targets modelling tools and model management technology providers.

Mondo addresses a number of pain points for these stakeholders. While MBE has been shown to provide substantial improvements in software development productivity and significantly enhance maintainability, consistency and traceability of software, existing modelling and model management technologies are being stressed to their limits in two dimensions: 1) their capacity to support collaborative development where several software developers or large teams of developers need to design or maintain a model for a software application; and 2) their capacity to efficiently manage and store large models of more than a few hundreds of megabytes in size.

Many organisations have committed substantial investments in moving towards MBE for software development only to discover that when using MBE for their larger software applications existing modelling tools have performance issues and create bottlenecks.

Some industries (e.g. Construction) have adopted standardised models that provide important benefits in automation and traceability, but often when these standard models are deployed the tools and storage technologies become overloaded.

How will the solution/service benefit the end-user?

MONDO will enable common tasks for MBE software development, like model queries and transformations, to be provided as Cloud services. Modelling tools as well as management and storage technologies will be able to scale to address MBE development for large or very complex software applications.

The extensible architecture of the MONOD platform will allow the integration with existing commercial and open source technologies to deliver value-added services such as improved performance for version-control management systems, model validation and verification facilities, and more sophisticated access control mechanisms for supporting large software development teams.

The openness and standards compliance of the MONDO platform will enable SMEs and individual developers to select and combine the Cloud-based services provided by the platform, and build upon them to develop innovative domain-specific services and solutions, without needing an upfront investment in expensive proprietary tools. SME modelling tool providers will be able to concentrate on the novel features of their tools that can give them a competitive advantage.

MONDO's robust and scalable platform for collaborative and controlled model development and management is an essential for MDE techniques to continue to deliver to software development organisations well-recognised benefits as complexity, diversity and size of software systems grow.

Website: <http://www.mondo-project.org/>



ORBIT - Business Continuity as a Service

Project start date 01/10/2013

Project finish date 31/03/2016

What user need or pain point is being addressed?

More and more areas of public life are becoming dependent on availability of Internet based services. Banks, logistics, travel, sales and media, to name a few, are severely disrupted when hit by service outages. Outages have serious implications on the continued operation of businesses - causing direct loss of revenue, legal liabilities as well as long term damage to reputation and brand names. A recent survey estimated a staggering €20 billion of annual revenue is lost because of IT downtime.

ORBIT provides high performance fault tolerance solutions for zero downtime in virtualized environments, enabling businesses to offer guaranteed services to their customers.

How will the solution/service benefit the end-user?

ORBIT solutions will minimize the potential damage from service outages for end users by providing a series of services for a variety of needs targeting individuals, businesses, governments and the public sector.

ORBIT will enhance the availability of large-scale facilities to help deal with the explosion of digital data results and the resulting need for larger computing infrastructures, to store and process increasing amounts of data. Many of these infrastructures are distributed, based on a collective ensemble of machines. As the size of the distributed infrastructure increases, so too does its susceptibility to failures.

ORBIT will enable services to become global and available 24/7, without the risk of downtime.

ORBIT will provide cost-sensitive and eco-efficient solutions, given that high level business continuity is resource-consuming both in terms of human capital and in financially, through for example dedicated hardware. Redundant hardware also increases the already-high carbon footprint of data centres.

Website: <http://www.orbitproject.eu/>

Panacea – Proactive Autonomic Management of Cloud Resources

Project start date 01/10/2012

Project finish date 31/03/2016



What user need or pain point is being addressed?

The main objective of the project “PANACEA” is to provide Proactive Autonomic Management of Cloud Resources, based on Machine Learning, as a remedy to the exponentially growing Cloud complexity. PANACEA will allow users several advanced possibilities, based on the Machine Learning (ML) framework, and the autonomic principles.

How will the solution/service benefit the end-user?

Users will benefit from PANACEA in the following ways:

- » Proactive autonomic management of cloud resources.
- » Proactive software migration within the cloud(s).
- » Creating mission-oriented distributed clouds with autonomic self properties.
- » Efficient use of cloud resources.
- » Monitoring, controlling and pro-actively managing applications’ executions (VM migrations, proactive rejuvenation, predicting the threshold violation of response time and the time to crash).

Website: <http://projects.laas.fr/panacea-cloud/>



What user need or pain point is being addressed?

Me and my friend are co-founders in a small mobile dev shop.

Our tech expertise is on the development of mobile apps, but for almost all of our projects, we need some back-end services running on the cloud in order to upload or download user data and contact 3rd party web services, like some weather services we needed in our latest app.

Not being web service tech savvy, we either have to outsource the back-end development, leading to increased costs, or we have to do it ourselves which slows down time-to-market. In addition, there is the potential risk of underperforming, especially in terms of reliability and security. And you know how competitive and unforgivable the mobile market can be...

We would almost certainly adopt an open source solution that would help us alleviate this pain, especially one with an active community, some successful case studies and a well documented website with lots of tutorials and examples. This is exactly what S-CASE is all about!!!

How will the solution/service benefit the end-user?

S-CASE has helped us build robust and scalable back-end web services for our mobile apps.

Being users of S-CASE for almost a year now, we have seen several benefits. We avoid writing boilerplate code for our RESTful web services which reduces lead times. Secondly, S-CASE incorporates good design practices and patterns in our codebase. Moreover, through S-CASE, we keep requirements and the respective source code always aligned.

External web services were never easier to use. We can now add to the codebase, almost automatically, the web service we want by just describing the functional and non-functional requirements it must satisfy. S-CASE recommends a list of services to select from.

In addition, even though our apps are customized per client, the back-end services are quite similar in terms of functionality. Through S-CASE we are able to identify similar functionality and re-use requirements, code and web services for every new app we build.

Building back-end services has never been easier! In fact, seeing how S-CASE has facilitated our workflow, we also pitched it to our friends working on pure Web API development.

Website: <http://www.scasefp7.eu/>

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

Project start date 01/10/2013

Project finish date 30/03/2016



What user need or pain point is being addressed?

SeaClouds is a novel open source framework that performs Seamless Adaptive Multi-cloud Management of Service-based applications.

This cloud's standard-based framework consist on an Application Management System over Clouds (AMSoC) which makes more efficient the design, development, planning and management of complex business applications across multiple and heterogeneous clouds, something unfeasible hitherto.

SeaClouds is a valuable tool for:

- » Application Designers who designs a service composition and interacts with the SeaClouds' planner component to obtain a Deployment Plan.
- » Cloud providers, who may not interact directly with the SeaClouds Platform, but the services offered are exploited by the platform to run service compositions.
- » Application administrators who will take advantage of SeaClouds to control the correct execution of the service composition deployed on multiple clouds.

How will the solution/service benefit the end-user?

SeaClouds solves a number of current problems and barriers existing on the cloud:

- » **Support for application deployment and migration to different providers.** SeaClouds provides support for deploying and migrating applications composed of several services. It takes care of the synchronization of services and their reconfiguration, without requiring the user to manually intervene.
- » **Management and monitoring of underlying providers.** Properties over applications and services deployed on multiple clouds can be ensured and managed in a standardized way by using unified metrics and automated auditing.
- » **Increased availability and security.** The use of formal models to support the management of service-based applications over multi-cloud environments gives more flexibility to reconfigure the distribution when SLA violation occurs.
- » **Performance and cost optimization.** The framework gives users freedom to distribute application requirements over different cloud offerings. Organisations can take advantage of useful and powerful services provided by each platform, avoiding its weaknesses.
- » **Low impact on the code and user-friendly interface.** Firstly, the development process is simplified by using SeaClouds tools and framework that require minor code changes. Secondly, the management of already deployed complex cloud applications is simplified thanks to the SeaClouds dashboard.

Website: <http://seacLOUDS-project.eu/>

What user need or pain point is being addressed?

The goal of SyncFree is to enable large-scale distributed applications without global synchronisation, by building upon the recent concept of Conflict-free Replicated Data Types (CRDTs). CRDTs allow unsynchronised concurrent updates, yet ensure data consistency. This revolutionary approach maximises responsiveness and availability, enabling data to be located near its users, in decentralised clouds. SyncFree aims to enable extreme scale replication, by applying eventual consistency techniques to CDRTs and beyond CDRTs.

Global-scale applications, such as virtual wallets, advertising platforms, social networks, online games, or collaboration networks all require consistency across distributed data items. As networked users, objects, devices, and sensors proliferate, the consistency issue is increasingly acute for the software industry. Classical alternatives are unsatisfactory as they either as they rely on synchronisation to ensure strong consistency, or forfeit synchronisation and consistency altogether with ad-hoc eventual consistency. The former approach does not scale beyond a single data centre and is expensive. The latter is extremely difficult to understand, and remains error-prone, even for highly-skilled programmers.

How will the solution/service benefit the end-user?

CDRTs avoid both global synchronisation and the complexities of ad-hoc eventual consistency by leveraging some simple formal properties of CRDTs. CRDTs are designed so that unsynchronised concurrent updates do not conflict and have well-defined semantics. By combining CRDT objects from a standard library of proven datatypes (counters, sets, graphs, sequences, etc.), large-scale distributed programming is simpler and less error-prone. CRDTs are a practical and cost-effective approach. The SyncFree project develops both theoretical and practical understanding of large-scale synchronisation-free programming based on CRDTs. Project results are new industrial applications, new application architectures, large-scale evaluation of both, programming models and algorithms for large-scale applications, and advanced scientific understanding.

The SyncFree project advances both the theory and practice of large-scale application architectures, and especially of CRDTs and related mechanisms. As the SyncFree industrial partners already have large user bases and feel the need for increased scalability in their applications, the project aims to include an extreme-scale crowd-sourced experiment, pushing the scalability needs of real world applications. An open-source library of CRDTs, 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. Beyond CRDTs the project explores global invariants in an extreme-scale environment to develop programming tools and patterns for extreme scale replication, and to experiment in vivo with extreme scale real applications.

Website: <https://syncfree.lip6.fr/>

Advanced Cloud Infrastructures and Services

Market ready in 2016 -2018



ALIGNED – Aligned, Quality-centric Software and Data Engineering



Project start date 01/02/2015

Project finish date 31/01/2018

What user need or pain point is being addressed?

I am the CIO charged with delivering big data applications. However, my software engineering teams and data engineering teams lack integrated processes and tools for developing systems where the data and software are evolving independently.

How can I efficiently assure application quality through independent software and data quality processes?

How can I increase my organisation's development agility and productivity?

How can I integrate inconsistent, incomplete public data sources with specialised domain expertise to efficiently produce web-scale curated datasets and applications?

How will the solution/service benefit the end-user?

The ALIGNED software and data engineering methodology will provide a basis for unified management views of co-evolving software and data systems.

Tool interoperability will be increased by publishing and consuming enterprise Linked Data that describes the lifecycle of the software and data engineering process.

Maintaining and integrating software and web data sources will be easier and the resultant quality and agility of these systems will be increased.

Website: <http://aligned-project.eu/>

AppHub – The European Open Source Market Place

Project start date 01/01/2015

Project finish date 31/12/2016



What user need or pain point is being addressed?

AppHub helps the market to seamlessly identify, position and implement the software outcomes of your projects.

AppHub provides the software produced by your project as cloud-ready packages that can be executed by a broad range of cloud service providers.

AppHub fosters adoption of projects results by making them trustworthy, easy to find, and easy to download and run.

AppHub, the open source software shopping mall, will also help accelerate open source adoption.

The partners that run and promote AppHub combine unparalleled expertise in open source community management, EU research projects and a breakthrough technology in software asset management.

Website: <http://www.apphub.eu.com/>

ARCADIA – A novel reconfigurable by design highly distributed applications development paradigm over programmable infrastructure



ARCADIA

A novel reconfigurable by design highly distributed applications development paradigm over programmable infrastructure

Project start date 01/01/2015

Project finish date 31/12/2017

What user need or pain point is being addressed?

Given the inability of Highly-Distributed-Application-Developers to foresee the changes as well as the heterogeneity on the underlying infrastructure, it is considerable crucial the design and development of novel software paradigms that facilitate application developers to take advantage of the emerging programmability of the underlying infrastructure and therefore develop Reconfigurable-by-Design applications. In parallel, it is crucial to design solutions that are scalable, support high performance, are resilient-to-failure and take into account the conditions of their runtime environment. Towards this direction, the ARCADIA project aims to design and validate a Novel Reconfigurable-By-Design Highly Distributed Applications Development Paradigm over Programmable Infrastructure. The proposed framework will rely on the development of an extensible Context Model which will be used by developers directly at the source-code level, assisted and validated by IDE-plugins in order to re-assure that the generated executable files contain meaningful semantics. According to ARCADIA's vision, the generated executables should be on-boarded by a Smart Controller which will undertake the tasks of translating annotations to optimal infrastructural configuration. Such a controller will enforce an optimal configuration to the registered programmable resources and will pro-actively adjust the configuration plan based on the Infrastructural State and the Application State.

How will the solution/service benefit the end-user?

ARCADIA's reference implementation and developed toolkits are going to facilitate application developers to design and develop infrastructural-agnostic applications and lead to the evolvement of novel and innovative paradigms for the deployment of advanced applications, boosting in this way the competitiveness of the software development industry. Based on the use of a set of annotations in the source code level, application deployment and management will be supported through the design and implementation of Smart Controllers. The Smart Controllers will orchestrate the dynamic allocation and management of resources (e.g. physical devices, virtual images, linux containers), as well as the configuration of the execution environment in real time. The approach followed ensures the optimal use of the available resources based on the existing policies, as well as the optimal runtime configuration in application and infrastructural level. This is based on orchestration provided by distributed collaboration among the Smart Controllers. Thus, optimisation of the application's execution according to multiple objectives (e.g. energy efficiency, QoS, security), as set by the end users, will be supported.

Website: <http://www.arcadia-framework.eu/>

BEACON - Enabling Federated Cloud Networking

Project start date 01/02/2015

Project finish date 31/07/2018



What user need or pain point is being addressed?

The main users of the BEACON project are organisations that manage IT infrastructures and wish to introduce network virtualisation technologies into their IT infrastructure. The BEACON project aims to integrate network virtualisation and service function chaining into Cloud middleware. The project will integrate the OpenDaylight open source project with the OpenNebula and OpenStack cloud middleware. OpenDaylight is an open platform for network programmability to enable software defined networking (SDN) and create a solid foundation for network function virtualisation for networks at any size and scale.

How will the solution/service benefit the end-user?

BEACON will allow users to provision customisable virtualised networks for their cloud applications. Today an IP network is difficult to customise to the specific needs of an application. With virtualised networks, application specific requirements can be passed to the virtualised network to select the network policies that meets the needs of the application. Applications will be able to select which network services, such as firewalls or deep packet inspection, they wish to use. These network functions will run on commodity servers at the edge of the network.

Website: <http://www.beacon-project.eu/>

What user need or pain point is being addressed?

Business growth as illustrated , for example, by smarter scenarios where new services facilitate daily human activities and give support for the growth of new markets and employment opportunities.

Innovative applications will be delivered thanks to new software engineering approach based on choreography of web services.

The problem is that existing choreography-based service composition approaches are rather static and are poorly suited to the needs of the Future Internet that is a highly dynamic networking environment, further bringing together highly heterogeneous services ranging from Thing- to Business-based services that span different security domains. As a result, the technology is not mature enough for market take-up.

How will the solution/service benefit the end-user?

CHOReVOLUTION adds the automated synthesis of dynamic and secured choreographies to existing choreography technologies. This makes these technologies able to support stringent application requirements in terms of dynamism and cross-organization security.

CHOReVOLUTION develops the technologies required to implement dynamic and secured choreographies in the Future Internet via the dynamic and distributed coordination of services. The distributed coordination of services within choreographies will enable innovative Future Internet applications that can support the dynamic adaptation and evolution of services due to the ever-changing environment.

Website: <http://www.chorevolution.eu/>

Project start date 01/09/2015

Project finish date 31/08/2017

What user need or pain point is being addressed?

Cloud computing and big data are important enablers for productivity and better services. The challenge for Europe is to ensure it retains its leadership position in research and innovation on cloud software and services. This leadership not only depends on Europe's ability to innovate at the technological level, but also to demonstrate tangible business value.

European Research and Innovation (R&I) projects need to think strategically, looking at technology and pricing as part of the same equation. While challenging, interoperable cloud services play a very important role in extending the market and in bringing business benefits to both the supply and demand sides.

Results from projects need to have an impact on the market if they are to become truly sustainable. However, this can be challenging. This is where CloudWATCH2 comes into play.

How will the solution/service benefit the end-user?

CloudWATCH2 will provide a set of services to help European R&I initiatives capture the value proposition and business case so that they can take their outputs to market.

- » **Pricing transparency:** A cloud market structure encouraging transparent pricing, showing how the market could become more efficient, more competitive and bring savings to consumers.
- » **Improved risk assessment:** A set of risk profiles and practical guides that fit around the different budgets and resources of private and public organisations.
- » **Security and legal guides:** A set of security measures that customers are recommended to take matched with expert legal guidance to lower barriers and ensure a trusted European cloud market.
- » **An evolved portfolio of standards for interoperability and security:** This is based on an analysis of implementations in European Research and Innovation projects, with emphasis on the business opportunities enabled by interoperability.
- » **A mapping of technologies, development status and practical support activities:** This provides a platform for projects to cluster and collaborate on common challenges, carry out interoperability testing, validate and endorse levels of interoperability.

CloudWATCH2 will contribute to forming a single market for cloud computing based on best practices and a common understanding of these best practices. This can help guide projects into becoming more competitive in the marketplace and sustainable for the long term.

Website: <http://www.cloudwatchhub.eu/>



CLARUS – A framework for user centred privacy and security in the cloud

Project start date 01/01/2015

Project finish date 31/12/2018

What user need or pain point is being addressed?

2015: I am the CTO of a Healthcare organization and I would like to benefit from cloud services, including software as a service, but my business asset includes sensitive data. The current cloud ecosystem is based on the assumption that appropriate security and privacy-preserving techniques are implemented by cloud providers.

I simply cannot take the risk of security breaches and my concern has increased due to the recent privacy disclosure and uncontrolled data outsourcing.

I think I am one of many potential users wishing to embrace the benefits of trusted cloud services while retaining full control over my sensitive data. I am looking for such services, but I am not sure they are available on the market and I cannot assess the guaranteed security service levels.

How will the solution/service benefit the end-user?

2017: CLARUS has improved privacy, security, and trust. The proxy solution they offer is exactly what I was looking for. It is easy to integrate it in my business and to use it for protecting my sensitive data.

It enables more transparent standardized auditable and controllable cloud services by implementing security and privacy-enabling mechanisms to ensure that the patient records are properly protected before outsourced to the cloud service provider.

This is a great benefit for my organization, now the medical records are safe and I am not worried anymore of potential attackers. Most importantly I can analyze them using the cloud computing power without being concerned of breaching their privacy.

CLARUS has been a model change tackling serious issues giving me as data owner control over my data, I now use cloud services based on standards that can be certified as being compliant with security and privacy.

Our stakeholders and patients are happy because they benefit from a more efficient service without being worried about highly sensitive data, which is very important in the healthcare sector today.

Website: <http://www.clarussecure.eu/>

CLOUDLIGHTNING – Self-Organising, Self-Managing Heterogeneous Cloud

Project start date 01/02/2015

Project finish date 31/01/2017



What user need or pain point is being addressed?

CloudLightning aims to address the problems surrounding cloud infrastructure management, service specification, access to, and provisioning of, cloud resources faced by both consumers and providers of cloud services. CloudLightning's initial target application domains are the oil and gas industry, genomics, and sectors that use ray tracing.

Currently, from a cloud service provider perspective, cloud service delivery models typically rely on over-provisioning to guarantee quality of service and accommodate unpredicted peaks in demand. The typical cloud server operates at approximately 20% computing capacity. The output capacity could be increased to 80% using heterogeneous processing resources with no change to the server's energy consumption. This represents a model of inefficiency, coalescing in ongoing sub-optimal energy consumption and related costs for the service provider and ultimately the consumer. Additionally, service providers must grant users access to resources to allow them to configure the service to meet their requirements, thus requiring security and access protocols to manage user access and permissions.

From a cloud consumer perspective, when interfacing with their chosen cloud service provider, they are tasked with the provisioning, configuring and optimisation of resources to meet their requirements. This requires a commensurate level of IT expertise and knowledge of the provider's infrastructure.


How will the solution/service benefit the end-user?

Consumers now have improved accessibility to heterogeneous cloud resources. The self-organising, self-managing principles used by CloudLightning provide consumers with a user-friendly service level interface to explicitly declare their requirements for service delivery. They can do this by specifying their exact workflow, dataset and time deadline. Through the assembly of dynamic resource coalitions, the self-organising and self-managing system automatically and intelligently locates the required resources and presents the most appropriate configuration of that service to the user for selection.

Cloud service providers now enjoy power-efficient, scalable management of their cloud infrastructures and better overall utilisation of service. CloudLightning's new standard in-service description language is not only easy to implement, but means that the service configuration presented by the self-organising, self-managing system will always be within the constraints of the service provider's resources, thereby providing optimal resource allocation and energy consumption. The coalition formation and service deployment processes negate the need to grant users access to manually configure their service description, thus providing an additional layer of security to the provider.

CloudLightning benefits service providers with the freedom to use its resources optimally and frees the consumer of the need and expertise required to configure cloud service delivery.

Website: <https://cloudlightning.wordpress.com/>



CloudSocket – Business and IT-Cloud Alignment using a Smart Socket

Project start date 01/01/2015

Project finish date 31/12/2017

What user need or pain point is being addressed?

Startups, small and medium enterprises are typically very much focused on their core business. Hence, there are several business processes like customer relationship and campaigning, administrative issues on registration, part of IT services as well as part of after sales support that are necessary for the business success, but can only insufficiently be handled by those organisations.

Business Processes in the Cloud enables brokers that may act as a public or private organisation to offer the cloud-based execution of those business processes the startups, founders and SMEs wants to delegate to reasonable price.

CloudSocket targets brokers like the business incubator or the process broker by providing tools, framework and knowledge to setup smart business and cloud alignment in form of a service.

The Business Incubator focuses on supporting the “Coaching and Finance” effort Start-ups with designing, analysing and simulating individual business plans, business processes, and also concerns a high degree of adaptability of Cloud Services for Start-ups, e.g. Customer Relationship Management, Order Management, Human Resources Management.

The Business Process Broker use case identifies typical business episodes of potential SMEs in different application domains as eHealth, Manufacturing, Photonics, Government, Security, e-Commerce, Retail, etc. but share a common set of business processes.

How will the solution/service benefit the end-user?

CloudSocket can be formulated as “The Smart Cloud Business Process Broker”, which comprises the discovery, orchestration, deployment and execution of services on the cloud. A learning cycle improves cloud individualization over time. This means that the level of integration has to be lifted from the technical to the business level. Currently the model-based approach is prominent for this integration. A detailed analysis of the business requirements based on the phases plan, model, manage and measure is common. Business models such as business processes, business rules and workflows are aligned with IT models such as IT infrastructure and architecture.

The proposed concept extends the integration problem from the IT level to the business level.

This is realised by:

1. supporting the extraction of business and IT-experts’ knowledge in a human and machine interpretable,
2. providing smart and intelligent tools to align business requirements to Cloud offerings,
3. supporting the model-driven and knowledge supported resource allocation of Cloud-based workflows,
4. enabling intelligent execution of models within a smart and adaptive BPaaS middleware
5. feeding back the process monitoring results from Cloud level to business level thanks to semantic enrichment and conceptual analytics.

Website: <https://www.cloudsocket.eu/>

CloudTeams - Collaborative Software Development Framework based on Trusted, Secure Cloud-based Pool of Users



Project start date 01/03/2015

Project finish date 28/02/2017

What user need or pain point is being addressed?

CloudTeams will be a cloud-based platform which will transform software development for cloud services into a much easier, faster and targeted process. We'll engage with communities of users who will participate in the product life cycle to help software teams develop better solutions for customer problems.

Main user-oriented challenges addressed by CloudTeams are twofold. Firstly, the fragmented European market that makes adaptation of innovations difficult and expenses for market research extremely high. Secondly, the lack of tools and incentives for users to collaborate with talented software teams around Europe in order to develop better solutions.

CloudTeams will use endpoints to existing services and tools that are popular in software development, by mashing them up with common practices. It will support developers with a collaborative platform where interaction with customers will feel natural and will validate the final outcome.

CloudTeams is the intersection of three important fields: crowd-sourcing platforms, collaborative software development tools and trusted cloud services delivery. This innovative combination of different tools and practices under a unique concept will be enhanced with a rewarding system to enable end-users to jump into the platform and support solutions they are interested in.

How will the solution/service benefit the end-user?

CloudTeams brings together software developers and prospective users, in an unprecedented collaborative platform that speeds up software development and team collaboration, by managing and linking information to extract knowledge and predict behaviour of users, to deliver better and valuable solutions. CloudTeams will develop a platform that can support the software development process through team collaboration, customer adaptation and constant testing with trusted users, in a transparent, privacy respectful way. By 2020, CloudTeams wants to be the prominent tool for trustful software development in Europe.

The CloudTeam Mission is to put customers' needs at the centre of software development, by providing complete software solutions and an active community of prospective users, with respect to their privacy. CloudTeams will develop a win-win situation for both developers and users. We will facilitate relationships of trust where software will alleviate pain and be beneficial to the users, while at the same time providing viable business models to the developers.

Website: <http://www.cloudteams.eu/>



Project start date 01/01/2015

Project finish date 31/12/2016

What user need or pain point is being addressed?

Complex applications are often distributed between cloud infrastructures to provide resilience against cloud provider outages, higher levels of elasticity, and better response times for their users by placing services near the clients. Moreover, they are often designed to scale automatically in response to demand and to permit live upgrades of the underlying software.

Application service providers and developers experience a number of deficiencies in the current cloud ecosystem:

- » Heavy relay of distributed applications on network services.
- » Lack of a unified identity management scheme between cloud providers
- » Control of network resources between data centres through the cloud APIs is currently impossible.

The CYCLONE project primarily targets application service providers who develop complex computing platforms and deploy them on cloud infrastructures in order to facilitate the deployment, management, and use of complex, multi-cloud applications, as well as enhance the end-to-end security of those applications.

How will the solution/service benefit the end-user?

CYCLONE's primary goal is to provide the software for a production-quality platform that facilitates the deployment and management of complex, cloud-based applications.

The project has identified two flagship applications. An academic cloud platform and associated services for bioinformatics research, and a commercial deployment for smart grids in the energy sector. These use cases are representative of complex applications to be consumed by end-users which demands security, reliability and flexibility while deploying them. Thus, application developers who are the main target of CYCLONE, will be able to easily deploy their applications over distributed federated scenarios with better networking resources control experience, enhanced security and unified identity management.

This is what CYCLONE delivers:

- » CYCLONE allows users to aggregate cloud resources from both private and public providers to build a cloud platform that is tailored to their application's needs
- » CYCLONE enables dynamic allocation of high-bandwidth channels inside and between data centers
- » CYCLONE allow pluggable monitoring services
- » CYCLONE implements end-to-end cloud security

To realize these objectives CYCLONE integrates and improves mature, open-source components, such as StratusLab, OpennNaaS, SlipStream and TCTP

Website: <http://www.cyclone-project.eu/>

DICE - Developing Data-Intensive Cloud Applications with Iterative Quality Enhancements



Project start date 01/02/2015

Project finish date 01/01/2018

What user need or pain point is being addressed?

DataInc is a small software vendor selling cloud-based Big Data services. DataInc has to deliver an initial version of a new Big Data application within 3 months, capable of processing sensor data acquired from a geographical area, with the goal of increasing coverage of areas, sensors and compute capacity on a monthly basis. The contractors require the application to be highly-available, scalable and cost-efficient. Yet, software developers are puzzled on how to implement a complex Big Data application of this kind in just 3 months: they could rush development but how could they satisfy all the quality requirements? What architecture should they adopt, keeping in mind the future evolution of the service? How should they accelerate quality testing for this initial release?

How will the solution/service benefit the end-user?

Software engineers at DataInc discover the DICE Horizon 2020 project. DICE provides an Eclipse-based integrated development environment (IDE) to design the new Big Data application. The IDE features the DICE profile, which extends UML to describe Big Data applications. The DICE quality analysis tools, integrated with the IDE, compare possible architectures by predicting their expected reliability, efficiency and safety characteristics. DICE also generates a cloud deployment plan for the application and offers testing tools for initial quality assessment. Once the application becomes operational, the application quality characteristics are “learned-as-you-go” from monitoring data and are fed back to the developer for evaluation. The iterative quality enhancement toolchain supports the developer in this task by identifying design anti-patterns and performance outliers in the monitoring data.

Using DICE, DataInc can rapidly implement the first version of the Big Data application and continuously evolve its quality characteristics over time.

Website: <http://www.dice-h2020.eu/>



ENTICE – dEcentralized repositories for traNsparent and efficienT vlrtual maChine opErations

Project start date 01/02/2015

Project finish date 31/01/2018

What user need or pain point is being addressed?

In this project, we research and create a novel VM repository and operational environment for federated Cloud infrastructures aiming to:

simplify the creation of lightweight and highly optimised VM images tuned for functional descriptions of applications;

automatically decompose and distribute VM images based on multi-objective optimisation and a knowledge base and reasoning infrastructure to meet application runtime requirements;

elastic auto-scale applications on Cloud resources based on their fluctuating load with optimised VM interoperability across Cloud infrastructures and without provider lock-in.

We gathered an interesting selection of complementary use cases from two SMEs and one industrial partner on energy control and management, earth observation and Cloud orchestration. For example, the WeSave use case provided by a Cloud provider called Wellness Telecom, is an application for energy control and optimisation of buildings that must be elastic to collect streams of data from various locations. WeSave is composed of various software components that require highly dynamic deployment and migration of VMs from one geographical location to another, as well as frequent starting and shutting down the VMs within seconds to provide satisfactory QoS to its end-users, a technology which is currently unavailable.

How will the solution/service benefit the end-user?

The VM images will be evaluated for size, functionality, and delivery time. The optimised images should be more than 60% than regular user-created VM images (comprising application users, OS experts and Cloud system administrators) while keeping their original functionality. We also expect to reduce the delivery time of the images from minutes to around 10 seconds for most cases. Even for more complex applications, where the original VM image delivery time is over 2 minutes, we will automatically create images with and reduce delivery time by more than 30%. For a given large set of VM images (at least 100) in a repository with over 100 GB of cumulative VM image size, we expect to reduce the storage requirements by more than 80%.

We also expect to reduce the time for cross data center deployment (i.e. when a VM is requested in a data center where the image is not available) by over 20% compared to manual or semi-automatic techniques with no storage optimisation involved. The optimisation process should improve the performance of the pilot use cases by at least 30%. Moreover, the multi-objective approach should preserve performance while decreasing the costs and storage requirements by at least 25% compared to solutions without repository optimisation. The new VM management methods will improve the QoS elasticity of the use cases from their currently inelastic status to elastic. This will mean that the percentage of QoS change will be at least equal to, or exceed the percentage of change in resource provisioning.

Finally, the knowledge model will address interoperability and integration issues of the use cases, and will achieve an over 25% of productivity increase in their VM image preparation and deployment time.

Website: <http://www.entice-project.eu/>

ESCUDO-CLOUD - Enforceable Security in the Cloud to Uphold Data Ownership

Project start date 01/01/2015

Project finish date 31/12/2017



What user need or pain point is being addressed?

Today, users placing data in the cloud need to put complete trust that the Cloud Service Providers (CSPs) will correctly manage such data. As a matter of fact, all CSPs can apply security measures in the services they offer, but these measures either give full trust to the CSP and allow it to have full access to the data, or greatly limit the functionality that the CSP is able to offer on the outsourced data.

ESCUDO-CLOUD will provide protection guarantees giving the data owners full control over their data in the cloud, while at the same time giving the cloud functionality over them. The goal of ESCUDO-CLOUD is to empower data owners as first class citizens of the cloud. This goal will be achieved by providing enforceable security, that is, techniques wrapping the data to provide a layer of protection to the eyes of the storing/processing CSP itself, setting the trust boundary at the client side, which means assuming correct and trusted behaviour only by the client. Data owners will then remain in control over their data when relying on CSPs for data storage, processing, and management.

How will the solution/service benefit the end-user?

ESCUDO-CLOUD will be beneficial to both data owners and Cloud Service Providers (CSPs). Data owners will be enabled to outsource their data while maintaining control over them, with the ability to regulate access to them and share them with other users in a selective way and with assurance that their data will remain protected from the CSPs. Data owners will then be able to rely on CSPs and use their services for a wider range of applications. CSPs significantly benefit, in addition to the increased market penetration that robust data ownership would provide, from reduced regulatory risks, audit costs, and general security threats that they would have to face in the absence of such protection. Freeing CSPs from the worries of protecting data, allows them to even handle the data outside their own control. For instance, it would enable a CSP itself to rely on other services for outsourcing storage and computation, behaving as a broker providing a virtualised cloud service, without worrying about the possible improper exposure of user information, which is guaranteed to be self-protected. This would benefit both larger as well as smaller players in the market, as well as individual users.

Website: <http://www.escudocloud.eu/>

HOLA CLOUD – Effective collaboration for European RD and Innovation in software, services and Cloud computing: Knowledge discovery and Roadmapping



Project start date 01/01/2015

Project finish date 31/12/2017

What user need or pain point is being addressed?

HOLA CLOUD targets all the stakeholders of the R&D domain in the field of Software, Services and Cloud computing including researchers, industrial players, related national and EU funded projects but also policy makers.

All these stakeholders have the common need to quickly and easily access updated information, obtain relevant contacts, identify the most recent and effective solutions, and compare their research, industrial and policy strategies in their domain.

Furthermore, to keep on top of competitiveness and growth, it is vital to be continuously updated on forefront technologies, so as to ensure that the internal R&D strategies are aligned with the future directions these technologies are taking.

How will the solution/service benefit the end-user?

HOLA CLOUD will build an advanced on-line portal and knowledge repository that ensures the persistent access to information for all stakeholders. This is a one-stop-shop for identifying projects, key players, suppliers, startups and early-stage technologies in the field of Cloud Computing and Software and Services.

This innovative platform will be supported by :

- » The Cloud Futures Conference series: a central element for and effective collaboration and roadmapping and at the same time addressing critical needs from the community behind European R&D and Innovation in Software, Services and Cloud computing.
- » The Cloud Hackademy: a series of events designed for supporting SMEs in ideating, launching, assessing, promoting or consolidating their innovation opportunities for cloud-based service provisioning
- » The Cloud Hackademy: a series of events designed for supporting SMEs in ideating, launching, assessing, promoting or consolidating their innovation opportunities for cloud-based service provisioning
- » The joint Roadmapping exercise to understand how the competitive position of the European Software, Services and Cloud Computing sector could be improved, and identify the way for achieving this vision, by setting mid/long term R&D priorities towards 2030.

Website: <http://www.cloudfutures.eu/>

HyVar – Scalable Hybrid Variability for Distributed Evolving Software Systems

Project start date 01/02/2015

Project finish date 31/01/2018



What user need or pain point is being addressed?

ICT is becoming increasingly integrated into our everyday environment. It is distributed on all kinds of devices from cars to household appliances. On top of that versions evolve continuously according to customer needs. The value chain of this process consists of:

- » End users who want good quality software, which responds exactly to their needs.
- » Device developers who produce the devices that the end-users will use such as control-units in cars, appliances and portable devices. They need to deliver new devices as fast as possible on the market and support the reconfiguration of these devices with requirements to high distribution, flexible maintenance and quality assurance.
- » Software developers who need to produce high quality software for the devices as fast as possible, responding to customer needs and to the device developers' objectives.

This value chain poses challenges in the domain of the Future Internet, combining advanced software management models with highly individualized, evolving expectations and requirements of the end users. In order to meet these challenges, a new approach has to be developed for faster and customisable software design for the following aspects: management of highly distributed applications; software updates exploiting device connectivity; tracing software versions spread everywhere; and maintaining and even improving the product quality in terms of reliability, resilience and customer perception.

How will the solution/service benefit the end-user?

HyVar proposes to integrate and enhance state-of-the-art techniques for the management of complex software systems from software product lines with cutting edge technology for over-the-air software upgrades and scalable Cloud solutions from European industry. This will support highly individualized and reconfigurable distributed applications.

HyVar's objectives are:

- » To develop a Domain Specific Variability Language (DSVL) and a tool chain to support software variability of highly distributed applications in heterogeneous environments, which allows developers to encompass unanticipated evolution as a standard feature of software systems in production.
- » To develop a cloud infrastructure that exploits the software variability supported by the DSVL and tool chain to track the exact software configurations deployed on remote devices and to enable (i) the collection of data from the devices to monitor their behaviour and perform statistical analyses; and (ii) secure and efficient highly customized software updates.
- » To develop a technology for supporting over-the-air updates of distributed applications in heterogeneous environments and enabling continuous software evolution after deployment on complex remote devices that incorporate a system of systems.
- » To test HyVar's approach as described in the above objectives in an industry-led Demonstrator to assess in quantifiable ways the benefits of the approach.

Website: <http://www.hyvar-project.eu/>



Project start date 01/10/2014

Project finish date 30/09/2017

What user need or pain point is being addressed?

We are a city council caring about the health and safety of our citizens. We have a wealth of information coming from various sensors in our city but we can't make sense of all that information and how to use it. It would be so helpful if we could efficiently use this information and the knowledge that can be derived to assist our citizens, especially ones with health concerns, to avoid overly polluted or dangerous areas, well in advance before the situation becomes hazardous for them. We'd also like to assure them about any security or privacy concerns they may have over this process.

We are an application development SME supporting assisted living for our elderly or disabled customers. How can we make good use of a federated cloud backbone consisting of home, area, city clouds so as to have the assisted living application "follow" our customers wherever they go? There are not only technical but also so many security constraints and trust issues as well to be resolved. How can we make sure our customers are assured their privacy is respected?

In any case, resources and software development don't come for free; how can we make the best of what is available?

How will the solution/service benefit the end-user?

Our city council now can exploit the information coming from sensors all around the city area and show our citizens we care both for their health, their time and we respect the council tax they pay. We can now promptly predict situations that can lead to hazards and lost time for them (CO2 emissions, pollen levels, floods, icy conditions, congestion). At the same time we realized that we could do so much more with our current infrastructure by managing data and services intelligently also for additional urban services. We even reduced our operational costs and we were able to pass this reduction down to our citizens.

Our application developer business is now going much better too. We can provide better support for our customers in their daily life wherever they go. We are also able to handle sensitive personal data much easier, as well as reduce the service response time. In addition, we can introduce new features to our products much faster re-using existing software resources, simply combining them in a different way. This together with a much more efficient use of cloud resources has given us an edge over competing solutions and has allowed us to increase our customer base considerably.

Website: <http://ikaas.com/>

INPUT – In-Network Programmability for next-generation personal cloud service support



Project start date 01/01/2015

Project finish date 31/12/2018

What user need or pain point is being addressed?

The INPUT Project aims to contribute to the evolution of the Internet “brain” beyond current limitations due to obsolete IP network paradigms, by moving cloud services much closer to end-users and smart-devices. This evolution will be accomplished by introducing intelligence and flexibility (“in-network” programmability) into network edge devices and by enabling them to host cloud applications capable of cooperating with and of offloading corresponding applications residing in the users’ smart objects and in datacenters, in order to realize innovative personal cloud services. This architecture will allow more attention to be placed on main technical challenges like energy efficiency, network programmability, the “softwarization” of the Internet of Things, network and datacenter virtualization, personal cloud services, etc.

How will the solution/service benefit the end-user?

The INPUT Project aims to design a novel infrastructure and paradigm to support Future Internet personal cloud services in a more scalable and sustainable way and with innovative added-value capabilities. The INPUT technologies will enable next-generation cloud applications to go beyond classical service models (i.e., IaaS, PaaS, and SaaS), and even to replace physical Smart Devices (SD), usually placed in users’ homes (e.g., network-attached storage servers, set-top-boxes, video recorders, home automation control units, etc.) or deployed around for monitoring purposes (e.g., sensors), with their “virtual images,” providing them to users “as a Service” (SD as a Service – SDaaS). Virtual and physical SDs will be made available to users at any time and at any place by means of virtual cloud-powered Personal Networks, which will constitute an underlying secure and trusted service model (Personal Network as a Service – PNaaS). These Personal Networks will provide users with the perception of always being in their home Local Area Network with their own (virtual and physical) SDs, independently from their location.

Website: <http://www.input-project.eu/>



IOStack – Software Defined Storage for Big Data

Project start date 01/01/2015

Project finish date 31/12/2018

What user need or pain point is being addressed?

The main objective is to create IOStack: a Software-defined Storage toolkit for Big Data on top of the OpenStack platform. IOStack will enable efficient execution of virtualized analytics applications over virtualized storage resources thanks to flexible, automated, and low cost data management models based on software-defined storage (SDS).

The user need is to reduce the costs of Big Data Analytics thanks to Software Defined Storage Automation techniques. It is very complex to deploy and operate big data storage and analytic clusters. We will enable the virtualization of analytic technologies (OpenStack Sahara) and the automation and cost reduction of storage with SDS techniques.

Problems solved: optimization of virtualized data analytics services such as Spark/Hadoop, solving I/O Bottlenecks with SDS techniques, storage cost reduction thanks to data reduction techniques.

An example end-user who could benefit from IOSTACK is be a medium IaaS provider or companies offering/operating advanced storage and data analytics services at a low cost.

How will the solution/service benefit the end-user?

Our end-users are normally IaaS providers or companies operating big data storage and analytics resources. They will benefit mainly in simplicity, administration and cost reduction. Instead of operating dedicated clusters, they can virtualize and automate these services thanks to IOStack platform. The advantages of this include high return on investment, cost reduction, and automation.

The IOStack platform will provide a solution in the OpenStack platform to the automated management and deployment of virtualized storage and analytic services. Nowadays, the complexity of the tools imply that medium IaaS providers and companies must operate dedicated storage and computing platforms. Or then resort to the major American players in the industry like Amazon or Microsoft. Our solution will lower the barrier for Big Data storage and analytics to a variety of small and medium providers.

Website: <http://www.iostack.eu/>

Project start date 01/01/2015

Project finish date 11/03/2017

What user need or pain point is being addressed?

Mikangelo targets providers of computational infrastructures, HPC, and big data services. We help providers to take full advantage of cloud computing on virtual infrastructures. Our users use cloud computing to highly utilise their infrastructure, to deploy applications easily, and to isolate applications between tenants as much as possible. Full machine virtualisation fulfills these requirements best, however it exhibits significant degradation of I/O performance, which leads to attrition of physical resources. In turn, poor I/O performance rules out I/O intensive applications. Furthermore, current virtual infrastructures do not handle bursts, which are common in cloud computing, well. Reasons for poor bursting performance are large operating system images and low I/O performance of VMs.

Two large use cases that will benefit from I/O efficient virtual infrastructures are HPC and big data. In HPC, virtual infrastructures are not used, because of the low I/O performance. Thus, end-users need to develop their applications against prescribed operating environments. In big data, end-users also need to adapt their applications to a prescribed big data platform with low flexibility. This lack in flexibility poses an entry barrier for many end-users to use HPC and big data.

How will the solution/service benefit the end-user?

Infrastructure providers that use Mikangelo's stack will be able to host I/O intensive applications in VMs. Providers will achieve better utilisation of their physical resources by providing new platforms to attract more end-users. End-users will gain flexibility, scalability, and elasticity for their applications in the cloud. Most of those benefits reside in the contexts of HPC and big data.

HPC infrastructure providers will run applications in VMs without loss of resources due to I/O inefficiencies. The VM approach to HPC also facilitates rapid scale out of HPC applications to a cloud, which further increases the utilisation of resources across the data centre. End-users will gain the flexibility to choose the best provider since they will be able to define the operating environment in a VM.

Big data infrastructure providers will be able to offer multiple custom platforms to end-users on demand. This flexibility will in turn attract more end-users to use big data applications on the provider's premises.

Both, HPC and big data, will benefit from Mikangelo's service management system with multiple end-user interfaces and an elaborate monitoring system. This monitoring system will allow end-users to identify bottlenecks in their applications to tweak application performance.

Website: <http://mikelangelo.uni-goettingen.de/>



MUSA – Multi-cloud Secure Applications

Project start date 01/01/2015

Project finish date 31/12/2017

What user need or pain point is being addressed?

The main goal of MUSA is to support the security-intelligent lifecycle management of distributed applications over heterogeneous cloud resources, through a security framework that includes: a) security-by-design mechanisms to allow application self-protection at runtime, and b) methods and tools for the integrated security assurance in both the engineering and operation of multi-cloud applications.

The main targeted users are three:

Application developers need tools to easily design the multi-cloud applications not only according to functional features but taking into account also security features such as data confidentiality, data integrity, data access and data location. They also need security mechanisms implemented in the applications to enforce the security at runtime.

System operators need to exploit cloud service combinations as much as possible and need tools to automatically select the best combinations with respect to the application functional and security needs and deploy the components accordingly.

The service administrators need to monitor the proper operation of the application (fulfilment of SLA), including the security features, in order to react to security incidents as soon as possible and to keep the users properly informed.

The three of them need tools that better integrate one another towards seamless assurance of security in the applications.

How will the solution/service benefit the end-user?

The data security incidents in multi-cloud applications have been reduced through the assurance of a secure behaviour of individual cloud-based components and the overall application, even if the data are processed and/or stored by untrustworthy or opaque cloud providers.

The cloud consumers' trust on clouds has enhanced by providing them with tools for expressing their security needs and keeping them informed on the security and performance faults of the multiple cloud services in use.

Application developers have a tool that allow them model the multi-cloud application based on the functional and security features to be offered in the SLA, as well as embed in the application components mechanisms to enforce the security at runtime.

System operators can automatically discover and select the best cloud service combinations that get most out of cloud by balancing performance and security.

Service administrators can assure the secure behaviour of multi-cloud applications and minimize the security risks while keeping the users informed.

This way MUSA has enabled security-intelligent lifecycle of multi-cloud applications.

Website: <http://www.musa-project.eu/>

PaaSword – A Holistic Data Privacy and Security by Design Platform-as-a Service Framework Introducing Distributed Encrypted Persistence in Cloud-based Applications



Project start date 01/01/2015

Project finish date 31/12/2017

What user need or pain point is being addressed?

Although enterprises recognize the compelling economic and operational benefits of running applications and services in the Cloud, security and data privacy concerns are the main barriers in Cloud adoption. Deploying confidential information and critical IT resources in the Cloud raises concerns about vulnerability to attack, especially because of the anonymous, multi-tenant nature of cloud computing. Current Cloud applications and storage volumes often leave information at risk to theft, unauthorized exposure or malicious manipulation.

The most critical part of a modern Cloud application and services is the data persistency layer and the database itself. PaaSword aims to fortify the trust of individuals and corporate customers in Cloud-enabled services and applications. The focus is on secure storage of both corporate and personal sensitive data on Cloud infrastructures. Because valuable business benefits cannot be unlocked without addressing new data security challenges posed by Cloud Computing.

How will the solution/service benefit the end-user?

To address the challenges, PaaSword will introduce a holistic data privacy and security-by-design framework based on distributed and encrypted data persistence and sophisticated context-aware access control mechanisms in Cloud-based services and applications. With the innovative PaaSword approach for key management, customers will have maximal control over the use of their data in Cloud services. PaaSword will extend the Cloud Security Alliance's Cloud security principles, by capitalizing on recent innovations in virtual database middleware technologies that introduce a scalable secure Cloud database abstraction layer with sophisticated data distribution and encryption methods. Furthermore, the implementation of enterprise security governance in Cloud environments will be supported by a novel approach towards context-aware access control mechanisms that incorporate dynamically changing contextual information into access control policies and context-dependent access rights to data stored in the Cloud. Finally, PaaSword will support developers of Cloud applications through code annotation techniques that allow specifying an appropriate level of protection for the application's data.

Website: <https://sites.google.com/site/paaswordeu/>



Procurement Innovation for Cloud Services in Europe

PICSE – Procurement innovation for cloud services in Europe

Project start date 01/10/2014

Project finish date 31/03/2016

What user need or pain point is being addressed?

The public sector is bound by procurement procedures and is not tailored to the dynamic, on-demand and elastic nature of cloud services. Procurement rules in some Member States can also make it difficult to sell cloud solutions to the public sector. PICSE targets the needs of public research organisations and libraries intending to procure cloud services. Many public organisations don't know how to deal with cloud contracts and SLAs or are concerned about their data security and the risk of lock-in, while others don't know how to allocate costs or are struggling to understand the impact that cloud can have on their IT model. So the real challenge is that while technology services continue to evolve, procurement processes and policies have remained firmly rooted in historical practices that are not effective in a cloud environment.

How will the solution/service benefit the end-user?

PICSE tackles these challenges by creating a more flexible and agile procurement model to purchase cloud services that will allow public research organisations to take advantage of the best the cloud market has to offer. The model will be published in June 2015 together with the PICSE wizard, a web-based guide on Procuring Cloud Services, and a self assessment tool that will allow public research organisations that are already procuring cloud services to better understand the gaps in their procurement process. Both tools will help procurement officers and IT managers in procuring cloud in an easiest way. In addition, PICSE has started in October 2014 a consultation phase with public sector organisations by collecting their experiences in procuring cloud services to understand the barriers they encountered in procuring cloud services, the main challenges and the best practices. This information will feed the PICSE final roadmap for cloud procurement over the next 5 years, to be released in March 2016.

Website: <http://picse.eu/>

RAPID - Heterogeneous Secure Multi-level Remote Acceleration Service for Low-Power Integrated Systems and Devices



Project start date 01/01/2015

Project finish date 31/12/2017

What user need or pain point is being addressed?

RAPID targets a novel heterogeneous CPU-GPU multi-level cloud acceleration focusing on applications running on embedded systems found on low-power devices such as smartphones, notebooks, tablets, wearable devices (smart watches, glasses), robots, cars, etc. Many such low-power devices can't always cope with the increased demand for processing power, memory and storage required by several applications in entertainment, vision, security, robotics, and aerospace such as gaming, antivirus, augmented reality, face and speech recognition, movement detection, biometrics, and CCTV. These applications require tremendous performance and cannot run on most existing battery operated devices. As a result, most such applications are only executed on high-end servers.

How will the solution/service benefit the end-user?

RAPID will allow the end-user to run applications that require tremendous processing power, memory and storage on low-power devices. Wearable devices and smartphones will benefit from RAPID by being able to execute more efficiently several applications such as speech recognition, augmented reality, navigation, video editing, as well as gaming. Moreover, power-operated robotic devices can use RAPID to run applications in a wide variety of fields such as medicine and surgery, personal assistance, security, warehouse and distribution applications, aerospace, defence, and ocean and space exploration.

Finally laptops, notebooks and desktops can also take advantage of cloud resources through RAPID by running even more demanding applications such as 3D video and image processing, advanced design, simulation and analysis tools used by engineers, and High Performance Computing.

Website: <http://rapid-project.eu/>



SERECA - Secure Enclaves for REactive Cloud Applications

Project start date 10/03/2015

Project finish date 10/03/2018

What user need or pain point is being addressed?

The Secure Enclaves for REactive Cloud Applications (SERECA) project targets two types of users/organizations: those wishing to move their mission-critical applications and their data to a Cloud infrastructure but do not trust the security of cloud-hosted applications and those who already use modern web applications hosted by cloud providers but the security solutions available to them do not provide adequate performance for latency sensitive.

An example is the case of a Critical Infrastructure that monitor key parameters of a dam for water supply. They would like to migrate to a cloud platform because this new technology can offer them many well-known benefits. However, they fear the integration between the cloud and the set of Industrial Control Systems (ICS) due to the security concerns.

How will the solution/service benefit the end-user?

At the end of the project we will have developed the SERECA secure cloud platform. We will develop the idea of a secure enclave into which applications can be deployed without having to rely on the questionable security mechanisms provided by cloud operators. SERECA aims to provide technical innovations that simultaneously establish sufficient trust and performance in cloud deployments through the secure connection of application components executing on secure commodity CPUs. The result is the secure distributed enclave, a novel technology that shifts the burden of trust from today's cumbersome and vulnerable multi-million-line software cloud stack to a small execution environment, exploiting the features of a commodity trusted hardware platform. This new approach will provide an attractive and scalable solution for cloud application hosting. We will extend the new innovative approach of secure CPU hardware in commodity processors known as secure enclaves (as ARM TrustZone, Intel SGX) and the vert.x reactive framework in order to make an execution of distributed reactive applications inside those enclaves possible. In conclusion, thanks to our innovative solution a user can execute reactive application and be sure that his data won't be touched by anyone, not even by malicious administrators.

Website: <http://www.serecaproject.eu/>

Project start date 01/01/2015

Project finish date 30/06/2016



What user need or pain point is being addressed?

SLALOM addresses top problems for potential adopters of Cloud services:

1. Lack of knowledge about what are fair and reasonable contractual terms and conditions related to service levels.
2. Lack of knowledge about how service levels need to be specified technically to provide meaningful protection for adopters.
3. Lack of resources – financial and personnel – to be able to research the issues.
4. Lack of clout to get some cloud service providers to offer fair and balanced provisions.

Our preliminary study of stakeholders identify the following initial list:

- » Adopters: public administrations, large companies, and SMEs
- » Cloud Service Providers (covering the full stack from IaaS to PaaS, and SaaS and additional services): General purpose providers (small and large); and Niche providers
- » Legal Firms and Professionals (either providing services or expert groups influencing the practice on Cloud SLAs: Legal firms and Profession influencers)
- » Policy Makers: regulators and policy makers; IT related government structures, Large Scale European organizations, experts and working groups (e.g. ECP and C-SIG; EC appointed experts and working groups); and Country level influencer (ministry, bodies and offices)
- » Standardization bodies: ISO, ETSI, OGF, IRTF, TOSCA, CAMP
- » Scientific Community and Researchers: FP7 and H2020 projects into cloud capability and SLA management

How will the solution/service benefit the end-user?

SLALOM provides templates for legal clauses and technical specifications exploiting advances of ISO standards in three levels: Legal – Model Terms (Master Service Agreement, Service Level Agreement); Technical and Technical Specifications. In short, SLALOM provides ready to use Clouds SLAs templates built on top of ISO standards.

One of SLALOM's objectives is to close the gap between cloud adopters and providers:

- » Cloud Service Providers can base their own SLA contractual clauses and technical specifications on the SLALOM recommendations.
- » Cloud Adopters will identify use of SLALOM to mean trustworthy and fair service level contractual terms and technical specifications.

The project will provide the following main positioning benefits: understandable baselines, practical templates for SLA contractual clauses and technical specifications, safe and fair balanced conditions for both providers and adopters.

Website: <http://slalom-project.eu/>

SLA-Ready - Making Cloud SLAs readily usable in the EU private sector



Project start date 01/01/2015

Project finish date 31/12/2016

What user need or pain point is being addressed?

2015: I am the CTO of a European SME. Using cloud services could be the answer to the rising complexity of software systems but ...

I find some of the service descriptions and contractual terms complex and misleading.

I am not happy about typical "take-it-or-leave-it contracts".

I don't understand what the service actually offers, which makes it hard to see the real benefits and be sure I am getting the best deal possible.

I have some sensitive customer data and need to understand how it will be protected and where it will be stored.

How can be sure I my business assets are secure?

How will the solution/service benefit the end-user?

2016: SLA-Ready social marketplace and tutorials-as-a-service have given me a much clearer path to the cloud.

SLA-Ready has given me a better understanding of service level agreements.

Now I can make an informed decision on what services to use, what to expect and what to trust.

Being able use a standardised Cloud SLA gives me a better understanding of the level of security and data protection offered by the cloud service provider.

It gives me the right answers for smooth and effective cloud adoption, including security levels.

Our investors and stakeholders are happy because we're able to innovate in the global marketplace through secure cloud services.

Website: <http://sla-ready.eu/>

SSICLOPS – Scalable and Secure Infrastructures for Cloud Operations

Project start date 01/02/2015

Project finish date 31/01/2017



What user need or pain point is being addressed?

The SSICLOPS project will focus on techniques for the management of federated private cloud infrastructures, in particular cloud networking techniques. Target customers include operators, service providers and other end user that require computing resources across clouds. Application developers, content or service providers that that require dynamic configuration, automated provisioning and orchestration of cloud resources are also SSICLOPS end users.

How will the solution/service benefit the end-user?

The SSICLOPS solution will provide efficient, scalable, and secure intra/inter-DC and client-facing transport mechanisms. This solution will be suitable for end users that require dynamic and efficient allocation of memory, computing and communication resources in same or across federated clouds.

Website: <https://ssiclops.eu/>

SUNFISH - SecUre iNformation SHaring in federated heterogeneous private clouds



Project start date 01/01/2015

Project finish date 31/12/2017

What user need or pain point is being addressed?

The overall problem that SUNFISH is trying to address is the lack of infrastructure and technology allowing Public Sector Players to federate their private clouds. To date, legislative barriers often make it difficult to use many commercial technological solutions. For example, Public Sector Entities in many EU countries cannot have system administrators resident in non-EU countries.

The Beneficiaries will be: EU Citizens who will feel that their data is more secure; and EU Public Sector Entities that will be able to rely on data being hosted on securely federated clouds. The SUNFISH project aims to reduce the management cost of private clouds owned by Public Administrations. Many Public Administrations in Europe need to manage data securely that are growing day by day. Scaling factors are required that overcome economic and technical capabilities of a single entity, even in a big Country. There is a real need to build up a federated cloud infrastructure specifically dedicated to several Public Administration players.

How will the solution/service benefit the end-user?

SUNFISH will develop a secure system to federate private clouds that will guarantee a high level of security, a continuous monitoring of inter-cloud communications, and the ability to roll out services cheaply, quickly, flexibly and securely, even between different private clouds.

These solutions are particularly pertinent given the current context. All European Public Administrations are considering with great interest the evolution of the ICT scenario towards "cloud" technologies, in order to achieve the best results in terms of costs cutting and efficiency improvement. In 2010, the European Network and Information Security Agency (ENISA) included cloud computing on their current and emerging research trends. The European Commission also made an explicit reference to cloud computing in the 'Digital Agenda for Europe 2020'. Following this initiative, the SUNFISH project aims at consolidating Europe as a 'trusted cloud region' in the post-Prism age, through the implementation of high standards for data protection, cyber security and transparency as regards public services.

The SUNFISH project will develop workload management innovations which will allow this to occur.

The private cloud infrastructures that support such applications are often only fully used only during brief periods of time, such as a few days per month. When these huge and costly infrastructures are not optimally used, they represent only a cost for the Public Administration.

Website: <http://www.sunfishproject.eu/>

SUPERCLOUD - User-centric management of security and dependability in clouds of clouds

Project start date 01/02/2015

Project finish date 31/01/2018



What user need or pain point is being addressed?

Despite many business benefits, distributed clouds raise many security and dependability concerns due to an increase in complexity and a lack of interoperability between heterogeneous, often proprietary infrastructure technologies. Current provider-centric clouds are notably faced with three major security challenges:

1. Security vulnerabilities in infrastructure layers: each layer (e.g., customer VMs, cloud provider services, provider hypervisor) is extremely vulnerable to attacks. For instance, the hypervisor and its over-privileged Dom0 is a target of choice for attackers due to its complexity. Hence the difficulty for an integrated protection.
2. Lack of flexibility and control in security management: the problem comes from heterogeneity of security components and policies between cloud providers. This has a major security impact by introducing more vulnerabilities due to mismatching APIs and workflows.
3. Security administration challenges: manual administration of protection of such an infrastructure is clearly out of reach to its complexity and heterogeneity of its components. Automation of security management is clearly necessary but lacking today.

SUPERCLOUD proposes new security and dependability infrastructure management paradigms that are:

1. user-centric, for self-service clouds-of-clouds where customers define their own protection requirements and avoid lock-ins;
2. self-managed, for self-protecting clouds-of-clouds that reduce administration complexity through security automation.

How will the solution/service benefit the end-user?

SUPERCLOUD will build a self-management infrastructure for security and dependability of heterogeneous resources across federated clouds. Customers will be provided with self-service environments enabling adaptive, customizable security for their cloud applications and services.

SUPERCLOUD will provide innovative cryptographic methods and tools for protecting data across distributed clouds through on-demand data security services, such as access control, blind computation, privacy-preserving indexing, and data availability.

SUPERCLOUD will enable resilient network-as-a-service, leveraging software-defined networking paradigms. It will also provide strong guarantees for end-to-end security and integrated trust management across multiple infrastructure layers and cloud domains.

SUPERCLOUD will reach the following objectives:

- » Self-Service Security: Implementation of a cloud architecture that gives users the flexibility to define their own protection requirements and instantiate policies accordingly.
- » Self-Managed Security: Development of an autonomic security management framework that operates seamlessly over compute, storage and network layers, and across provider domains to ensure compliance with security policies.
- » End-to-End Security: Proposition of trust models and security mechanisms that enable composition of services and trust statements across different administrative provider domains.
- » Resilience: Implementation of a resource management framework that composes provider-agnostic resources in a robust manner using primitives from diverse cloud providers.

Website: <http://www.supercloud-project.eu/>



SWITCH - Software Workbench for Interactive, Time Critical and Highly self-adaptive Cloud applications

Project start date 01/02/2015

Project finish date 31/01/2018

What user need or pain point is being addressed?

Time critical applications often involve distributed components, and intensive data communication. For instance, applications which address disaster warning issues often include remotely deployed sensors, and many live event television broadcast scenarios require direction of multiple outdoor video sources. The development of such applications is usually difficult and costly, because of the high requirements for the runtime environment, and in particular the sophisticated optimisation mechanisms needed for developing and integrating the system components. In the meantime, a Cloud environment provides virtualised, elastic, controllable and quality on demand services for supporting systems like time critical applications. However, the engineering method and software tools for developing, deploying and executing classical time critical applications have not yet included the programmability and controllability provided by the Clouds. Time critical applications cannot yet get the full potential benefits which Cloud technologies could provide.

It has therefore become imperative to develop such software methods and tools. SWITCH will develop the software methods and tools for the entire lifecycle of time critical Cloud applications.


How will the solution/service benefit the end-user?

The SWITCH project aims to achieve a major break-through in the development of time critical Cloud applications. The complexity of time critical Cloud applications makes their development much more difficult than development either of classical time critical applications or of normal Cloud applications. Simply applying existing approaches cannot maximize the usage of Cloud environments for time critical applications, because little consideration of Cloud programmability is currently taken in the development of time critical applications. With the new approach developed by the SWITCH project, the user is able to include controllability of the infrastructure as part of the application specification, and make the system self-adaptable in the Cloud.

The programming and control model, and the software tools developed in the SWITCH project, will make considerable impact on

1. Improving development productivity of time critical Cloud applications,
2. Upgrading industrial technologies of time critical applications to use Cloud infrastructure,
3. Improving deployment efficiency of time critical applications,
4. Reducing operational cost of time critical services, 5) Promoting business competitiveness of Clouds.

Website: <http://www.switchproject.eu/>







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

CloudWATCH is funded by the European Commission 7th Framework Programme
Directorate-General for Communications Networks, Content and Technology Software & Services, Cloud. Contract No. 610994



European
Commission