(IP)

Proposal Number 610802
Proposal Acronym CloudWave

Title: Agile Service Engineering for the Future Internet
Duration in months: 36
Funding: 6.300.000 €

Abstract:

CloudWave will revolutionise modern cloud infrastructures and tools by enabling agile development and delivery of adaptive cloud services which dynamically adjust to changes in their environment so as to optimise service quality and resource utilization. In doing so, CloudWave will advance the competitiveness of the EU economy by providing a powerful foundation, based on open standards, for creating innovative Future Internet services, realising the full promise of cloud computing.

CloudWave will tangibly deliver (1) an open architecture and standards-based reference implementation of an advanced cloud software stack, with novel capabilities for adaptation across all cloud layers; and (2) tools and methods for agile development of reliable and adaptable cloud services, facilitated by the new stack.

Technologically, CloudWave will advance the state-of-art along three dimensions:

•Execution Analytics: A new framework where specialised algorithms dynamically analyse cloud infrastructure and application behaviour, seamlessly integrate data pertaining to physical and virtual resources and IoT elements, and provide consolidated feedback to drive service evolution and adaptation

•Coordinated Adaptation: A new software technology where cloud services, cloud infrastructure and end-user devices exploit Execution Analytics to collaboratively and automatically undertake complex adaptation actions across the cloud stack, ensuring quality of service and effective utilization of ICT resources

•Feedback-Driven Development: A new agile approach for developing cloud applications, where developers exploit Execution Analytics to incrementally determine and evolve application features, extensions and optimizations, based on observed user needs

Using three industry case studies, we will demonstrate how the CloudWave technology supports emerging Future Internet services and provides high and quantifiable improvements in service delivery quality, productivity, reliability and cost.

Coordinator:
IBM ISRAEL - SCIENCE AND TECHNOLOGY LTD

Partners:
ATOS SPAIN SA
Cloudmore OÜ
IBM ISRAEL - SCIENCE AND TECHNOLOGY LTD
INTEL PERFORMANCE LEARNING SOLUTIONS LIMITED
SAP AG
TECHNION - ISRAEL INSTITUTE OF TECHNOLOGY
TELECOM ITALIA S.p.A
UNIVERSITA DEGLI STUDI DI MESSINA
UNIVERSITAET DUISBURG-ESSEN
UNIVERSITAET ZUERICH
VALADIS SYSTEM SA

(IP)

Proposal Number 611068
Proposal Acronym CoherentPaaS

Title: A Coherent and Rich PaaS with a Common Programming Model
Duration in months: 36
Funding: 4.950.000 €

Abstract:

In the cloud computing landscape it is clearly understood that there is "no one size fits all". This has resulted in a blooming of different cloud data management infrastructures specialized for different kinds of data and being able to perform orders of magnitude better than traditional approaches.

Unfortunately, this trend has resulted in a wide diversification of APIs, the loss of a common programming paradigm and the lack of coherence across different cloud data managers. The loss of a common programming paradigm is quite noticeable. The ACID coherence provided in traditional environments has been totally lost in the cloud landscape. Even within a single cloud data store providing transactional semantics (such as key-value data stores or sharded SQL databases) no transactional semantics is provided at the data store level. Instead, sharding (splitting the data into many different pieces) is applied to be able to scale linearly. This results in transactions that cannot span across shards resulting in an effective loss of transactional semantics.

CoherentPaaS addresses precisely these two major issues in the cloud PaaS landscape. CoherentPaaS will develop a PaaS with a rich diversity of cloud data management technologies including no SQL data stores such as key-value data stores and graph databases, SQL data stores such as in-memory and column-oriented databases, hybrid systems such as SQL engines on top on key-value data stores, and complex event processing data management systems. All these systems will be programmed by using a common query language under a uniform paradigm. Additionally, a scalable transactional management system will provide holistic coherence across data stores. With the outcome of CoherentPaaS cloud application developers will be able to develop applications using multiple cloud data management technologies with a single query language and holistic coherence, therefore, being able to develop applications faster and with higher quality.

Coordinator:
UNIVERSIDAD POLITECNICA DE MADRID

Partners:
FOUNDATION FOR RESEARCH AND TECHNOLOGY HELLAS
INESC PORTO - INSTITUTO DE ENGENHARIA DE SISTEMAS E COMPUTADORES DO PORTO
INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE
INSTITUTE OF COMMUNICATION AND COMPUTER SYSTEMS
MonetDB
NEUROCOM SA
PORTUGAL TELECOM INOVACAO SA
Quartet Financial Systems Ltd
SPARSITY TECHNOLOGIES
UNIVERSIDAD POLITECNICA DE MADRID

(STREP)

Proposal Number 610874
Proposal Acronym ASCETiC

Title: Adapting Service lifeCycle towards EfficienT Clouds
Duration in months: 36
Funding: 3.190.000 €

Abstract:

Adapting Service lifeCycle towards EfficienT (ASCETiC) is focused on providing novel methods and tools to support software developers aiming to optimise energy efficiency and minimise the carbon footprint resulting from designing, developing, deploying, and running software in Clouds. At the same time, quality of service, experience and perception will still be taken
Into account, so energy efficiency will complement them and boost cloud efficiency at several dimensions.
This project focuses on Cloud services made of several shared software components, which are likely to be used and reused many times in many different applications. The major contribution to the carbon footprint of Cloud services is the energy consumed in its operation, thus the primary aim of ASCETiC is to relate software design and energy use, which will depend on the deployment conditions and the correct operation of the service by means of an adaptive environment.
The project has the following main objectives:
a) Development of models for green and efficient software design, supporting sustainability and high quality of service levels at all stages of software development and execution;
b) Development and evaluation of a framework with identified energy efficiency parameters and metrics for Cloud services;
c) Development of methods for measuring, analysing, and evaluating energy use in software development and execution, complementing quality measures;
d) Energy and quality efficiency integration into service construction, deployment, and operation leading to an Energy Efficiency Embedded Service Lifecycle.
ASCETiC will result in an implementation of an open-source Cloud stack providing energy efficiency at software, platform, and infrastructure layers. We will provide incremental versions of ASCETiC with respect to energy efficiency, approaching the problem in three levels: static, intra-layer and inter-layer adaptation. We will demonstrate the ASCETiC solution in two commercial use cases.

Coordinator:
ATOS SPAIN SA

Partners:
ATHENS TECHNOLOGY CENTER SA
ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS - RESEARCH CENTER
ATOS SPAIN SA
BARCELONA SUPERCOMPUTING CENTER - CENTRO NACIONAL DE SUPERCOMPUTACION
CENTRE D'EXCELLENCE EN TECHNOLOGIES DE L'INFORMATION ET DE LA COMMUNICATION
GREEN PREFAB ITALIA SRL
HEWLETT PACKARD ITALIANA SRL
TECHNISCHE UNIVERSITAT BERLIN
UNIVERSITY OF LEEDS

(STREP)

Proposal Number 610711
Proposal Acronym CACTOS

Title: Context-Aware Cloud Topology Optimisation and Simulation
Duration in months: 36
Funding: 3.450.000 €

Abstract:

Last years have seen a dramatic increase in cloud infrastructure scale, heterogeneity and complexity. Demand for better energy efficiency have lead to a variety of different technological options to build servers from different CPU architectures such as ARM or x86 as well as specialised options for highly parallel tasks (Manycore boards such as XeonPhi or GPGPUs). Additionally the service complexity has evolved from models very similar to traditional server hosting to more interactive services (e.g. remote rendering or gaming) as well as moving towards more complex services on top of Hardware and basic platform services. Similarly data centres have made significant investments in energy efficient buildings, server racks and facility management technology and understand themselves as Smart Consumers in evolving SmartGrid environments.
In order to cope with the challenge to optimise the mapping of services to a variety of different resources from hardware over software (e.g. licenses) demands for new intelligent and cross-domain integration of actual and historical data. Topology aware mapping and placement of them across geographically distributed centres is necessary. The propose will address the follow challenges:
* Modelling of heterogeneous workloads, infrastructure landscapes as well as facility management information and energy supplier information
* Collection and Analysis of historical use data and derivation of intelligent management strategies integrating research results from the Cloud and Data Centre Management field as well as from Mathematics
* Realisation of management methods including dynamic workload placement, scheduling and migration by continuous optimization across multiple dimensions
* Validation by simulations conducting costs and risk analysis, performance validation of the developed strategies
* Validation by deployment in two distinct scenarios for a business analytics an technical computing use case

Coordinator:
UNIVERSITAET ULM

Partners:
Flexiant Limited
FORSCHUNGSZENTRUM INFORMATIK AN DER UNIVERSITAET KARLSRUHE
SAP AG
THE QUEEN'S UNIVERSITY OF BELFAST
UMEA UNIVERSITET
UNIVERSITAET ULM

(STREP)

Proposal Number 610582
Proposal Acronym ENVISAGE

Title: Engineering Virtualized Services
Duration in months: 36
Funding: 3.180.000 €

Abstract:

Today, reliability and control of virtualized resources are barriers to the industrial adoption of cloud computing. These barriers are not satisfactorily addressed in software engineering. To remedy this, the goal of ENVISAGE is a resource-oriented design-by-contract methodology to efficiently deliver virtualized services. The methodology addresses two major challenges inherent to virtualization:
(1) the leveraging of contracts and service-level agreements (SLA) into programming models, and
(2) the leveraging of resource management into the early phases of service design.
The methodology of ENVISAGE will build on a foundational study of virtualization and SLA, which goes beyond today's cloud technologies. ENVISAGE will create a practical open-source framework to develop SLA-aware services, with highly automated analyses using formal methods. SLA-aware services will be able to control their own resource management and renegotiate SLA across the heterogeneous computing landscape. The methodology and framework will allow services to be delivered in a more effective, efficient, and reliable manner than today, accelerating the development cycle and lowering the operational costs for innovative networked services that make use of cloud computing. By mastering virtualization in the engineering of services, ENVISAGE has the potential to overcome these barriers, significantly improve the competitiveness of SMEs, and profoundly influence business ICT strategies in all sectors.
The consortium has the complementary expertise required for these important challenges, with leading researchers in formal methods, contract-based specification, service-oriented computing and resource analysis, and with innovative industries using virtualized services in different, representative business domains. The project's tasks are carefully designed to jointly achieve its goal through a timely succession of results, building on the consortium's broad experience in European projects.

Coordinator:
UNIVERSITETET I OSLO

Partners:
ALMA MATER STUDIORUM-UNIVERSITA DI BOLOGNA
ATBROX AS
ENGINEERING - INGEGNERIA INFORMATICA SPA
Fredhopper B.V.
STICHTING CENTRUM VOOR WISKUNDE EN INFORMATICA
TECHNISCHE UNIVERSITAET DARMSTADT
UNIVERSIDAD COMPLUTENSE DE MADRID
UNIVERSITETET I OSLO

(STREP)

Proposal Number 611337
Proposal Acronym HEADS

Title: Heterogeneous and Distributed Services for the Future Computing Continuum
Duration in months: 36
Funding: 3.460.000 €

Abstract:

The future computing continuum is composed of a highly heterogeneous network of platforms and devices, with very different capabilities. This diversity is the key to build innovative services seamlessly interacting with the physical world. However, it also raises new software engineering challenges. Current approaches to engineer cloud services provide productive abstractions for virtualizing infrastructures but lack specific support for supporting the rest of the continuum and efficiently interacting with the physical world by exploiting the capabilities of mobile devices, smart-objects, sensors and actuators.

The objective of the HEADS project is to provide a cost-effective and integrated programming model which allows full exploitation of the capabilities of the individual open, legacy or proprietary devices and platforms composing the future computing continuum.

The HEADS approach combines and advances model-driven and generative software engineering techniques to (i) efficiently exploit cloud resources and resource constrained devices and networks, (ii) optimize distribution of data processing and (iii) automate deployment and validation of systems executing on the future computing continuum.

The main result of the HEADS project is a new IDE and an operational methodology for specification, validation, deployment and evolution of software-intensive services distributed across the future computing continuum. The HEADS approach exploits existing execution platforms and the IDE provides a plug-in mechanism to easily extend the IDE to support new devices and execution platforms

The HEADS consortium consists of six partners from five EU countries. Two academic partners: SINTEF (Norway), INRIA (France); two technology providers: Software AG (Germany), eZmonitoring (Ireland); and two use-case providers: ATC (Greece) and Tellu (Norway). The total budget of the HEADS project is 4.8M euros (3.4M euros in requested funding).

Coordinator:
STIFTELSEN SINTEF

Partners:
ATHENS TECHNOLOGY CENTER SA
GEOGRAPHICAL INFRASTRUCTURE SOLUTIONS LIMITED
INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE
SOFTWARE AG
STIFTELSEN SINTEF
Tellu AS

(STREP)

Proposal Number 611125
Proposal Acronym MONDO

Title: Scalable Modelling and Model Management on the Cloud
Duration in months: 30
Funding: 2.670.000 €

Abstract:

As Model Driven Engineering (MDE) is increasingly applied to larger and more complex systems, the current generation of modelling and model management technologies are being pushed to their limits in terms of capacity and efficiency, and as such, additional research is imperative in order to enable MDE to remain relevant with industrial practice and continue delivering its widely recognised productivity, quality, and maintainability benefits. The aim of MONDO is to tackle the increasingly important challenge of scalability in MDE in a comprehensive manner.

Achieving scalability in modelling and MDE involves being able to construct large models and domain specific languages in a systematic manner, enabling teams of modellers to construct and refine large models in a collaborative manner, advancing the state-of-the-art in model querying and transformations tools so that they can cope with large models (of the scale of millions of model elements), and providing an infrastructure for efficient storage, indexing and retrieval of large models. To address these challenges, MONDO brings together partners with a long track record in performing internationally-leading research on software modelling and MDE, and delivering research results in the form of robust, widely-used and sustainable open-source software, with industrial partners active in the fields of reverse engineering and systems integration, and a global consortium including more than 300 organisations from all sectors of IT.

Coordinator:
X/OPEN COMPANY LIMITED

Partners:
ASSOCIATION POUR LA RECHERCHE ET LE DEVELOPPEMENT DES METHODES ET PROCESSUS INDUSTRIELS - ARMINES
BUDAPESTI MUSZAKI ES GAZDASAGTUDOMANYI EGYETEM
IKERLAN S.COOP.
SOFTEAM
SOFT-MAINT
UNINOVA - INSTITUTO DE DESENVOLVIMENTO DE NOVAS TECNOLOGIAS
UNIVERSIDAD AUTONOMA DE MADRID
UNIVERSITY OF YORK
X/OPEN COMPANY LIMITED

(STREP)

Proposal Number 609828
Proposal Acronym ORBIT

Title: Business Continuity as a Service
Duration in months: 30
Funding: 2.290.000 €

Abstract:

More and more areas of public life become dependent on availability of IP based services. Banks, logistics, travel, sales and media - to name a few - are severely hit by small and specially larger and longer outages of services. As events in 2012 have shown datacenter outages can result from a number of unforeseeable events. For example a datacenter in Ireland was affected by a lightning strike. In another event a failure of data collection services had widespread effects on cloud resources in Northern-America, taking down a number of high-trafficked websites.

From an EU perspective there are additional aspects to be taken into consideration. For example in case of service discontinuity a provider might replicate data to other data centers, though by doing that breaching compliance rules for data in the EU.

In short: Loss of availability of mission-critical services, such as SAP HANA in-memory relational database or other data repositories with dependency on very high availability, may result in both capital and human loss.

And while addressing HA for statefull and fully-consistent applications, such as SAP HANA, can be done via specialized hardware of heavy application modifications, in this project we aim at achieving this end by applying generic techniques which require neither of the two, ultimate attempting to provide HA as a standard cloud service. We farther investigate and address the full range of outage (or service downtime) causes: from single hardware failures (commonly referred to as faults) effecting a host or a network switch, up to major faults (commonly referred to as disasters) bringing down an entire data center or site.

Coordinator:
INSTITUTE OF COMMUNICATION AND COMPUTER SYSTEMS

Partners:
DEUTSCHE WELLE
IBM ISRAEL - SCIENCE AND TECHNOLOGY LTD
INSTITUTE OF COMMUNICATION AND COMPUTER SYSTEMS
SAP AG
UMEA UNIVERSITET

(STREP)

Proposal Number 610764
Proposal Acronym PANACEA

Title: Proactive Autonomic Management of Cloud Resources
Duration in months: 30
Funding: 2.130.000 €

Abstract:

The main objective of the project "PANACEA" is to provide Proactive Autonomic Management of Cloud Resources as a remedy to the exponentially growing complexity.
If you look at the system resources (Internet) at the bottom of the stack, that system resource can be servers, storage, data centres, and network resources, the concept is then to build a level of virtualization of those resources so that any given event is not tied to one box necessarily or to one storage disk. Once you get that kind of leverage, you can build the set of functions that relate to autonomic self-* properties: configuring, healing, optimizing and protecting. The design that you have to have holistically has to deal with the fact that components are going to fail. The aim of a Cloud Computing platform is to support redundant, self-recovering, highly scalable programming models that allow workloads to recover from many inevitable hardware/software failures and monitoring resource use in real time for providing physical and virtual servers, on which the applications can run.

It will propose innovative solutions for autonomic management of cloud resources, which will be based on a set of advanced Machine Learning Techniques and virtualization. A Machine Learning (ML) framework will be created for a proactive autonomic management of cloud resources. It will allow predicting the failure time of software, or user applications running on Virtual Machines (VM) and the violation of expected response time of services.

To deal with the vast number of possible resources to monitor, our main approach will consider the use of mobile agents, which will move on the cloud, interacting with other agents, reading computing and network sensors, and making autonomous decisions on what to measure, when to report and to whom. Distributed Machine Learning, based on Reinforcement Learning and Neural Networks, will be used to enforce "self-organizing paths".

Coordinator:
IRIANC

Partners:
ATOS SPAIN SA
CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE
IBM ISRAEL - SCIENCE AND TECHNOLOGY LTD
IMPERIAL COLLEGE OF SCIENCE, TECHNOLOGY AND MEDICINE
IRIANC
QOS DESIGN
UNIVERSIDAD COMPLUTENSE DE MADRID

(STREP)

Proposal Number 610717
Proposal Acronym S-CASE

Title: Scaffolding Scalable Software Services
Duration in months: 36
Funding: 2.520.000 €

Abstract:

The goal of S-CASE is to accelerate the software development lifecycle for cloud services by introducing a new agile prototyping paradigm that automates the process of mapping user requirements to concrete software specifications and generates operational code (RESTful services). S-CASE aspires to support software developers in identifying software requirements and business processes in various formats, including textual (requirements/use case documents), formal (UML diagrams), and visual (images of UML diagrams or storyboards) content. In order to realise this vision, S-CASE employs appropriate multimodal information processing techniques, such as natural language and image processing. S-CASE also aims to provide the appropriate mechanisms for synthesising composite executable workflows of resources (software solutions, services, and devices), both proprietary and open source. To this end, semantic matchmaking and service orchestration techniques will be applied for the dynamic discovery of resources and their composition into operational workflows that match the software developer envisaged requirements in the best possible way. All R&D outcomes will be integrated into the S-CASE realm, a cloud-based platform composed of tools for developers and service providers, open-source software repositories and a discoverable service ecosystem that provides the appropriate resources in the form of RESTful services. Three pilot applications will be developed and deployed in real operational environments by two industrial partners and one SME, in order to validate and evaluate the S-CASE paradigm. The aim of the pilots is to showcase how S-CASE can successfully meet different needs and technology requirements in the three domains of energy, social networks and cloud computing infrastructure. In all pilots, an evaluation procedure will be applied for assessing potential gains with respect to the acceleration of the software development process.

Coordinator:
ARISTOTELIO PANEPISTIMIO THESSALONIKIS

Partners:
akquinet tech@spree GmbH
ARISTOTELIO PANEPISTIMIO THESSALONIKIS
CENTRE FOR RESEARCH AND TECHNOLOGY HELLAS
ENGINEERING - INGEGNERIA INFORMATICA SPA
ERICSSON NIKOLA TESLA D.D.
Flexiant Limited
THE UNIVERSITY OF EDINBURGH
UNIVERSITA DEGLI STUDI DELL'INSUBRIA
Δ. ΔΑΣΚΑΛΟΠΟΥΛΟΣ Α.Ε

(STREP)

Proposal Number 610531
Proposal Acronym SeaClouds

Title: Seamless adaptive multi-cloud management of service-based applications
Duration in months: 30
Funding: 2.190.000 €

Abstract:

Cloud computing reduces time-to-market and provides on-demand scalability at a low cost. Many private and public Clouds have emerged during the last years, offering a range of different technologies each suited for particular types of applications.
SeaClouds tackles the problem of deploying and managing, in an efficient and adaptive way, complex multi-services applications over technologically dissimilar Clouds environments. SeaClouds will support the best distribution of the modules according to the deployment requirements and the strong and weak points of each offering. This allows organisations to embrace Cloud solutions and, at the same time, avoid risks of unreliability and lock-in.
SeaClouds provides an answer to questions such as:
• How can a complex application be deployed, managed and monitored over multiple Clouds?
• How can a complex application be reconfigured if run-time problems are detected?
SeaClouds' approach is based on the concept of orchestration. Orchestration manages the different services complex applications are composed of, without the need to modify the code of the services. Orchestration can be formally designed to fulfil properties such as Quality of Service of the whole application. Reconfiguring an application is also performed by changing the orchestration of the services, when the monitoring detects that properties such as QoS are not respected.
The objectives of SeaClouds are therefore:
• orchestration, adaptation, and verification of services distributed over different Cloud providers
• monitoring and runtime reconfiguration operations of services distributed over different Cloud providers
• unified application management of services distributed over different Cloud providers
• promote and align with major standards for cloud interoperability, particularly OASIS' CAMP and TOSCA

Coordinator:
ATOS SPAIN SA

Partners:
ATOS SPAIN SA
Cloudsoft Corporation Limited
CYNTELIX CORPORATION BV
Nurogames GmbH
UNIVERSIDAD DE MALAGA
UNIVERSITA DI PISA

(STREP)

Proposal Number 609551
Proposal Acronym SyncFree

Title: Large-scale computation without synchronisation
Duration in months: 36
Funding: 2.670.000 €

Abstract:

The goal of SyncFree is to enable large-scale distributed applications without global synchronisation, by exploiting 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; it enables locating data near its users, in decentralised clouds.

Global-scale applications, such as virtual wallets, advertising platforms, social networks, online games, or collaboration networks, require consistency across distributed data items. As networked users, objects, devices, and sensors proliferate, the consistency issue is increasingly acute for the software industry. Current alternatives are both unsatisfactory: either to rely on synchronisation to ensure strong consistency, or to 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.

SyncFree avoids both global synchronisation and the complexities of ad-hoc eventual consistency by leveraging the 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 will develop both theoretical and practical understanding of large-scale synchronisation-free programming based on CRDTs. Project results will be new industrial applications, new application architectures, large-scale evaluation of both, programming models and algorithms for large-scale applications, and advanced scientific understanding.

Coordinator:
INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Partners:
BASHO TECHNOLOGIES LIMITED
FACULDADE DE CIENCIAS E TECNOLOGIADA UNIVERSIDADE NOVA DE LISBOA
INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE
KOC UNIVERSITY
Rovio Entertainment Oy
TECHNISCHE UNIVERSITAET KAISERSLAUTERN
Trifork A/S
UNIVERSITE CATHOLIQUE DE LOUVAIN