

# FP7 Call 1 Achievements in Software & Services

SOFTWARE & SERVICE ARCHITECTURES,  
INFRASTRUCTURES AND ENGINEERING

... FP7-ICT-2007-1 – Objective 1.2  
September, 2011



European Commission  
Information Society and Media

#### LEGAL NOTICE

By the Commission of the European Communities, Information Society & Media Directorate-General, Software & Service Architectures, Infrastructures and Engineering Unit.

Neither the European Commission nor any person acting on its behalf is responsible for the use which might be made of the information contained in the present publication.

The European Commission is not responsible for the external web sites referred to in the present publication.

The views expressed in this publication are those of the authors and do not necessarily reflect the official European Commission's view on the subject.

More information on the European Union is available on the Internet (<http://europa.eu>).

Cataloguing data can be found at the end of this publication.

Luxembourg: Publications Office of the European Union, 2011

ISBN 978-92-79-20614-6

doi:10.2759/49652

© European Union, 2011

Reproduction is authorised provided the source is acknowledged.

*Printed in Belgium*

PRINTED ON RECYCLED PAPER

# Table of contents



Introduction ..... 5

ADMIRE ..... 14

ALIVE..... 16

COMPAS..... 18

DEPLOY..... 20

DIVA ..... 22

FAST..... 24

FLOSSINCLUDE ..... 26

IRMOS ..... 28

MANCOOSI ..... 30

M:CIUDAD..... 32

MOST..... 34

NESSI 2010..... 36

NEXOF-RA ..... 38

OMP ..... 40

OPEN ..... 42

PERSIST..... 44

PROTEST..... 46

Q-IMPRESS..... 48

RESERVOIR ..... 50

ROMULUS ..... 52

S-CUBE .....	54
SERVFACE .....	58
SERVICE WEB 3.0 .....	60
SHAPE .....	62
SLA@SOI .....	64
SMARTLM .....	66
SOA4All .....	68
STREAM .....	70

# Introduction

## Objective 1.2: “Service and Software Architectures, Infrastructures and Engineering”

Over the next 15 years, advanced network and service infrastructures will open the door to a variety of new networked applications, costing less to build and run. To reap full benefit from these opportunities, Europe’s current internet, mobile, fixed and broadcasting networks must progress, together with the platforms for the deployment of services and the ability to engineer dependable software and services.

### Outcome of FP7 Call 1 for proposals

The overall goal of Challenge 1 within the FP7 Information and Communication Technology Research Programme was to enable the emergence of network, software and service technologies that open up new application scenarios and innovative business models, thus creating novel business opportunities and growth. The aim was and still is, though adjusted and revised, to provide technology for scalable networks of interconnected devices, machines, individuals and organisations, served by new, configurable, and dynamic software services.

This will eventually empower users to manage networked media content and will allow organisations to adapt to rapidly changing, networked business ecosystems.

At the time of the call within Challenge 1, Objective 1.2: “Service and Software Architectures, Infrastructures and Engineering” was concerned with research on virtualisation of resources, service platforms and service/software engineering.

Three years ago, in June 2008, we published the brochure illustrating the portfolio of projects selected after the first call of the Framework Programme 7, [ftp://ftp.cordis.europa.eu/pub/fp7/ict/docs/ssai/ssai-fp7-project-portfolio-final\\_en.pdf](ftp://ftp.cordis.europa.eu/pub/fp7/ict/docs/ssai/ssai-fp7-project-portfolio-final_en.pdf).

Now, three years later, most of these projects have terminated their activity and this brochure collects their achievements and successes.

A lot has happened during these three years. The Future of Internet landscape has consolidated and new realities have come to play a significant role. The Future of Internet Assembly is a regular recognised event to discuss topics relevant to future developments, both in the Internet of Service area and in other areas such as for example, the Internet of Things. The Future Internet PPP has been launched and the first projects started this year. Cloud Computing has established itself as a reality which any new development or implementation in this area has to consider and evaluate.

In this evolving reality our projects have found their way. Sometimes having to readapt their objectives to the changing world around them, but most of the time being at the forefront of this movement, turning ideas into new tools and thoughts into implementations.

## Working Groups in the Software and Services area

To help and support projects to better achieve their goals and to enhance the impact potential, since the end of FP6 collaboration among projects has been stimulated.

Collaboration has been encouraged in order to ensure co-ordination and coherence of project activities in the area, to increase the impact in common tasks that require critical mass and to optimize the use of resources. The overall aim was to build a stronger community and to achieve better and more visible results. Collaboration work was and still is structured around *Collaboration Working Groups*. Working Groups topics and leadership have changed over the years following the interest and the needs of the projects. The groups themselves have been jointly decided by the projects and participation to each or some of them is decided on a voluntary base and variable geometry. The HOLA! Support Action currently hosts the CWG homes at <http://www.holaproject.eu/>.

Several CWG have been the source of cooperative activities: workshops have been organised jointly, summer schools, books published, etc. Others have been more a source of information exchange among projects. In some cases, these exchanges have resulted in a closer cooperation among projects and reuse of results.

Among these Collaboration Working Groups, whose exhaustive list is available on the HOLA! web site, a few have achieved the results presented in the following sections.

These working groups results complement the information on individual project results that is provided in the rest of this publication.

### Collaboration Working Group: Quality of Service & Service Level Agreements

The advancements in service-based environments have made feasible the delivery of Internet-scale services by enabling a rich set of services to efficiently operate, manage and reconfigure computing, storage and network resources. Such environments allow various capabilities to be offered as on-demand services (i.e. X-as-a-Service), driving the emergence of corresponding providers (e.g. infrastructure, platform, software, etc). Furthermore, service-centric computing paved its way to break down the monolithic application silos and provided opportunities to any organization to re-think business processes and traditional service models. In this context and given the different stakeholders as well as their incentives for participation in value-chains, research efforts were focused on Service Level Agreements (SLAs) and the corresponding frameworks that manage these agreements (i.e. publishing, negotiation, signing, billing, etc). Another challenge that

#### Identity Card

##### CWG title

Quality of Service & Service Level Agreements

##### CWG current objectives

Share mutual understanding of QoS & SLAs across various projects and develop a mediation framework that allows for translation between different perspectives. Examine the role of QoS & SLAs in the Future Internet and drive joint initiatives towards broader communities like NESSI or FIA. Identify the dependencies / barriers etc between QoS provision and SLAs

##### Contact person

Dimosthenis Kyriazis  
National Technical University of Athens  
[dkyr@telecom.ntua.gr](mailto:dkyr@telecom.ntua.gr)

service-based environments faced refers to the nature of the applications being deployed and executed, which turned from batch processing to interactive multimedia ones. The latter raised the need for efficient adaptation of resource provisioning to the dynamic demands of the applications, the majority of which tends to be real-time and interactive. To facilitate this requirement, research efforts aimed at providing *Quality of Service* (QoS) guarantees at various levels of service-based environments (e.g. modeling, monitoring and triggering corrective actions, employing resilience mechanisms, etc). In the framework of different European Projects, significant outcomes have been published leading both to academic recognition in the worldwide research community and to industrial exploitation. Representative examples are the GT4 SLA framework by BEinGRID, the Quality Reference Model by S-Cube, the Claudia framework for SLA management by RESERVOIR and the ISONI framework that guarantees QoS on the infrastructure level by IRMOS.

## Topics Addressed

The aim of the Collaboration Working Group (CWG) is to provide a forum for collaboration amongst the researchers in the areas of QoS and SLAs, by ensuring that they share a common understanding of the research topics across various projects (e.g. terminology, functional and non-functional requirements, supporting technologies, relevant standards, etc). To this direction, CWG members from more than twenty (20) different projects brought into the group the outcomes and research interests from their projects in order to identify and highlight topics on which collaboration is possible with other projects. In many cases, fruitful discussions provided useful feedback and directions regarding ongoing work, while in other cases research outcomes were used as a baseline to build on top of them.

In the area of Quality of Service, CWG members have addressed topics related to *modeling and analysis of quality attributes* as well as *prediction and control mechanisms* for supporting management decisions according to these quality attributes. Besides, within the framework of different projects, efforts focused on *trade-off analysis between different service level objectives* (e.g. performance, reliability, maintainability, providers' reputation, etc) in order to perform service selection taking into account the trade-offs that may result to different end-to-end levels of quality (AssessGrid, IRMOS, Q-ImPRESS, RESERVOIR, S-CUBE, SORMA).

In the area of Service Level Agreements, a number of interesting topics have been addressed, such as *specifications and multi-level decomposition of SLAs* based on a hierarchical form to allow different providers in the value chain to provide, exchange and share services (IRMOS, RESERVOIR), *SLA lifecycle management* in different layers (i.e. business and technical) considering their relations and *fully automated SLA negotiation* between different actors in the value-chain (IRMOS, SLA@SOI, SmartLM). There has to be noted that the aforementioned automated negotiation was developed as a concept following discussions taking place within the CWG with regard to what has already been developed by other projects (i.e. ARGUGRID, AssessGrid, BEinGRID). Members have also tackled the topics of *dynamic service and resource provisioning* based either on SLA protection against QoS violations or on business rules (BEinGRID, RESERVOIR) as well as how different *mapping schemes* can be deployed in service based platforms: between high-level application terms and fine-grained resource-level attributes and between business-level objectives to resource management policies and vice versa (BEinGRID, IRMOS). Finally, *trust models in SLAs*, efficient *accounting and billing* mechanisms, and *monitoring approaches triggering events* for runtime adaptability of resource provisioning have also been addressed (ARGUGRID, AssessGrid, BEinGRID, IRMOS, RESERVOIR).

## Next Steps

The QoS&SLAs CWG will continue serving as a forum to establish collaboration and find synergies amongst researchers – topics highlighted by projects started in 2010 refer to metamodels for QoS (e.g. SRT-15 and FITTEST), lifecycle of SLAs including legal terms (e.g. Cloud4SOA and OPTIMIS), SLA management for data-intensive services (e.g. VISION Cloud), term languages and ontologies (e.g. CONTRAIL and CloudTM).

Nevertheless and given that the realization of the Future Internet vision and in general of any future architecture will primarily depend on resource provision and the environments able to support the different offerings and roles of users in value-chains, within the framework of the QoS&SLAs CWG, members will address topics related to SLAs and business models due to the emergence of new service platforms (e.g. social marketing environments) and revise quality concepts to manage the uncertainty introduced by various services and data (e.g. smart phones, social networks, sensors, user generated content, etc) in today's rich digital environment.

## Collaboration Working Group: Service Front-End

FP7 Call 1 projects have been prolific in innovation and results for the organisations involved on it. Three have been the main areas in which these projects have focused their efforts although all of them have participated in these three areas in one way or another.

One of these groups is *Dynamic Interfaces*. In this research line, Call 1 projects have provided results in delivering seamless and transparent support to users in carrying out their tasks when changing services and/or devices, even in multi-user applications, while offering personalised interaction. Very interesting results were achieved in automatic discovering services, operations and data types to orchestrate new complex services for final users, based on their preferences and contexts.

*Context Management* can be named as another line where very positive results have been got. The focus of this research line is to define new services based on the user preferences. One of the main results in this group is the concept of personalised smart spaces where end users can define their preferences and can modify their behaviour according to the existence of other smart spaces of other users.

And finally, when talking about defining user-centric services, why not providing tools for end users to define and generate their own services for their business and daily activities? *Prosumers, user empowerment and mashup*, is another research line where Call 1 projects have been very active. Several proposals for service generation environments have been defined, covering both traditional (wired) and mobile environments. These environments provide end users the possibility of creating new services by composing other simple services and pieces of code, commonly called *building blocks*.

All these activities and results achieved by the FP7 Call 1 projects form the base for the new challenges and research activities to be faced in this Working Group.

### Identity Card

#### CWG title

Service Front-Ends

#### CWG current objectives

Share mutual understanding of relevant topics for the Service Front End across various projects. To address in particular research topics such as context modeling and management, or evolution of web technologies that enable users, organized in communities, to mash-up, configure connect, and share services in a knowledge-aware manner.

#### Contact person

Kevin Doolin

Waterford Institute of Technology, Ireland.

kdoolin@tssg.org

The most important current challenge for this CWG is standardisation. There are very few standards regarding Service Front-Ends. Standardisation has demonstrated its value for evolution at research time and it is time to embody these achievements in European and World-wide standards.

Another challenge is automation. There are techniques for discovering services, and generating new complex services based on more simple ones. However, new research activities have to be performed in order to decide the most suitable alternatives to automate service creation, to decide how these criteria will be analysed, how the final user will be involved in this selection process and how these systems will solve the possible mistakes when not providing the service requested by the user.

## Collaboration Working Group: Internet of Services – Trust and Security Developments

A cross-project collaboration group on Trust and Security was established as early as FP5. This is an important area for Software and Services projects, because it is not (and should not become) the focus for research in these projects, yet the results of all IoS projects must be trustworthy and secure if they are to have wide impact on the Internet and European society. The group provides a forum for projects to pool their requirements in this area, and identify solutions from within the projects or from elsewhere in the Programme. Call 1 projects have made a substantial contribution to the group, during which time advances were made in several areas.

### Security requirements analysis

One of the biggest challenges for advanced, service-oriented infrastructure and applications is the difficulty of getting service operators and users (who are usually not security experts) to describe their security requirements and concerns. The DEPLOY project focused on making formal methods more accessible to a wider

#### Identity Card

##### CWG title

Trust and Security

##### CWG current objectives

Trust and Security covers a very broad area including technical issues such as the need and mechanisms for authentication, business issues such as the grounding of trust in business processes or in Virtual Organisations, and operational issues such as the need for user training and software maintenance. The aim is to identify and discuss common interests across multiple projects that can lead to benefits for these projects.

##### Contact person

Mike Surridge

IT Innovation, University of Southampton  
ms@it-innovation.soton.ac.uk

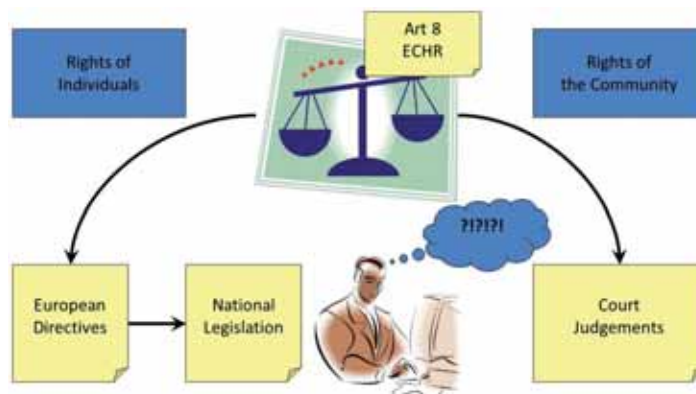


body of professional users, allowing them to express requirements (including but not limited to security requirement) very precisely. The COMPAS project also used formal methods to express requirements and verify compliance, showing that business regulatory compliance (e.g. with SOX, Basel II or ISO 27001) could be translated to formal technical security requirements. The IRMOS project took a less formal approach to capture and analyse security requirements in real-time

multimedia applications. To raise end user awareness of security issues, a survey was first carried out, identifying major concerns for the confidentiality and integrity of multimedia content and streams. Then the users were asked to provide detailed security input using OCTAVE Allegro worksheets, a method chosen because although less precise than formal methods, it is accessible to end users who may lack specialised knowledge in ICT security. IRMOS then used off-the-shelf security mechanisms to address the identified concerns, collaborating with the earlier FP6 Edutain@Grid project to integrate these mechanisms to address real-time multimedia streaming services.

## Privacy requirements

Privacy is one of the most important drivers of security requirements for service-oriented applications. This is a huge challenge for research and development because case law and legal precedents can introduce new requirements at any time. The Trust and Security group has sought to address this by inviting experts in privacy law to present regular updates to projects. In



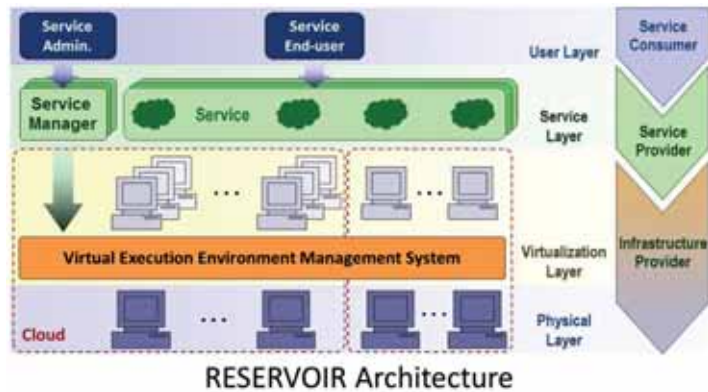
this period, new requirements emerged from the European Court of Human Rights judgement in July 2008 in the case of I versus Finland. This placed the obligation on data controllers to prove that personal data was not improperly accessed. Service oriented applications involving personal data must therefore be auditable, which implies cross-domain applications must have cross-domain (federated) security. This is highly relevant for IoS technologies such as Clouds.

## Cloud security technologies

Call 1 included some of the first FP7 research projects focusing on the development and use of Cloud Computing technology. The RESERVOIR project developed a federated Cloud infrastructure, allowing service providers (i.e. Cloud users) to describe and deploy services to any Cloud hoster. The hoster can then run the service using its own resources, or those from other 'federated' RESERVOIR hosters. This makes Cloud Computing more scalable and flexible, and allows services to migrate to stay close to users if they change or move around. RESERVOIR identified a wide range of security concerns in such a Cloud, starting with the need for service providers (i.e. Cloud users) to describe their security (including privacy) requirements so they can be respected when deploying or migrating services. It is also necessary to isolate services so rogue or compromised service providers can't attack each other or the Cloud itself, and to secure the protocols used to deploy or migrate services and manage the Cloud. RESERVOIR addressed many of these challenges using established security mechanisms such as X509 certification, and developed new ways for service providers to specify security requirements and match them with Cloud resources. These formed the basis for contributions to best practice guidelines from the Cloud Security Alliance, and to standardisation through ENISA, the Open Grid Forum (OCCI), and the Distributed Management Task Force (OVF). Other contributions to Cloud security from Call 1 included a novel software licence manager from the SmartLM project, which allows users flexible access to licenced software in elastic and virtualised (e.g. IaaS Cloud) environments. This makes it possible for software vendors to 'trust' Cloud infrastructures by giving them a means to enforce the terms of their licences without negating the benefits of Clouds to end users.

## Security architecture, specifications and standards

The introduction of standardised transport level security protocols such as HTTPS has provided a vast improvement in the security of websites and other services. But in advanced service-oriented applications this isn't enough. One also needs to standardise the way service communicate securely with each other across domain boundaries, often on behalf of users who are outside



the security domain of either of the communicating services. This challenge has been much studied by the Grid community in FP5 and FP6 projects, during which time a whole stack of specifications have been developed to address security of services, including Web Service Security stack from OASIS (WS-Security, WS-Trust, WS-Federation, etc). FP6 and FP7 projects have continued to collaborate to define architectures and design patterns for using these and related standards. By the start of FP7, it was possible to identify best practice from this research, and this was done by the Call 1 project NEXOF-RA, which engaged with the community and documented a reference architecture based on the 'greatest hits' from this body of work. This has formed an important starting point for subsequent developments, contributing to the conceptualisation of a Future Internet platform now being implemented in the Future Internet PPP.

## Collaboration Working Group: Collecting Use Cases

Use cases are more and more needed in research as a test bed for evaluating new approaches and techniques. The scope of this working group is to define a methodology to describe use cases in a uniform way and to build a repository where use cases and, possibly, associated technological solutions are described. The development of such a repository will allow, on the long run, to build a set of benchmarks that could be useful to researchers in order to compare their solutions with the existing ones on concrete cases. The work has been initiated starting from the needs of the S-CUBE Network of Excellence and by gathering inputs from various projects, including ALERT, BEInGrid, IRMOS, MOSAIC, NEXOF-RA, PLAY, Reservoir, SLA@SOI, SOA4ALL.

## Advances

The working group has achieved two outcomes so far:

- The definition of an approach to describe use cases. The approach leverages from the approaches existing in the literature and provides some guidelines to software engineers on the steps to follow.
- The development of a wiki-based repository containing descriptions of use cases developed adopting our description approach. (<http://scube-casestudies.ws.dei.polimi.it>)

The project addresses the needs of researchers who are looking for use cases shared with others and that can, therefore be, used as examples on which to apply the approaches they develop in order to perform comparisons and benchmarking.

The working group contributes to interoperability as it proposes a uniform way to describe use cases and to make them comparable and classifiable.

The users of the working group results will be researchers and developers interested in offering their use cases to others or in exploiting our use cases to evaluate their solutions. In this way there is the possibility of building proper benchmarks for experimenting with new solutions to IT problems.

Examples of use can be found in <http://scube-casestudies.ws.dei.polimi.it> and concern the usage of our approach to describe cases from various business domains (automotive, health-care, wine, traffic management, crisis management).

## List of all Collaboration Working groups on Software & Services

- **Collecting Use Cases:** This group has developed a methodology to describe use cases and a repository to store use cases so that they can be reused among the projects.
- **Dissemination:** This group has developed a web portal that supports the collaboration among the projects and they organise a yearly summerschool in Greece.
- **Formal Methods for SOA (Service Oriented Architectures) and Future Internet:** This group investigates how established software engineering techniques can be improved for use in future situations, where software systems are dynamic and have no well-defined boundaries. Their main objective is knowledge exchange and organisation of common workshops.
- **QoS and SLAs (Quality of Service and Service Level Agreements):** Share mutual understanding of QoS & SLAs across various projects (terminology, functional and non-functional requirements, relevant standards & technologies) and drive joint initiatives towards broader communities (e.g. NEXOF-RA, Future Internet Assembly, standardisation).
- **Semantics:** Link ontologies that have been developed by the different projects that participate, in order to obtain ontologies with a wide scope and stimulate reuse.
- **Service Architectures:** This group is jointly developing a body of knowledge on service oriented architectures.
- **Service Engineering:** this group has written a book, where one common software problem was solved in different innovative ways by the projects that wrote the chapters.
- **Service Front Ends:** This group has organised the 1st international workshop on Service Front End technologies in the future Internet of Services conference. Their objectives are to develop:
  - A common vision on the technologies and architecture associated to Service FrontEnds in the future Internet of Services
  - Open specifications and, potentially, open source reference implementations of components in the envisioned architectureWhen agreed, results of these joint effort will be submitted for adoption by the NEXOFRA initiative promoted by NESSI.

- **Standards:** The majority of the Software and Services projects have a strong interest in taking forward its results into standardisation and in producing results which comply with existing and evolving standards. But with so many different and potentially parallel standardisation activities, which are mostly driven by globally operating standardisation bodies, it is vital to co-ordinate the projects' efforts to strengthen collaboration, exploit synergies and increase the impact of European standardisation activities on the world stage, which is done in this group.
  
- **Trust & Security:** It is clear that there are distinctive requirements for dealing with trust in industrial or commercial projects. With this wide remit, it will not be possible to find a common technical approach addressing all project requirements. The emphasis is therefore to identify more focused topics where there are common interests across multiple projects, and discussion can lead to benefits for these projects. The group maintains a list of current topics and circulates summary information about their scope and eventual conclusions to all participants.
  
- **Virtualised Service Platforms (cloud computing):** The main activities of this working group are:
  - Define what is meant by a service, and determine the set of requirements a service must meet in order to be supported. This will take into account envisioned future Internet trends such as sensors, RFIDs, "smart dust" etc.
  - Identify other FP7 stakeholders in virtualized service platform and solicit their involvement.
  - Formulate the non-functional requirements on a virtualized service platform from an application provider point of view.
  - Verify the set of requirements against contributed use-cases.
  - Start working on a common definition of the interfaces required by service/application providers for lifecycle management.
  
- **Best practices with open source:**
  - the use of forges, development and release tools, and development platform
  - best practices to manage a developers' community
  - how to increase the likelihood of living code after the end of the project
  - how to select a FLOSS license, and how to get help with licensing issues.

All these best practices have been documented in the "FP7 guide to Open Source Software available at [http://ec.europa.eu/information\\_society/events/cf/ios10/document.cfm?doc\\_id=14042](http://ec.europa.eu/information_society/events/cf/ios10/document.cfm?doc_id=14042)



# ADMIRE



### At a Glance

**Project title**

Advanced Data Mining and Integration Research for Europe

**Contact person**

Dr Parsons Mark Ian  
The University Of Edinburgh, Epcc  
m.parsons@epcc.ed.ac.uk

**Website**

www.admire-project.eu

**Total cost / EC contribution**

4,241,573 € / 3,001,662 €

**Start date /end date**

March 2008 / May 2011

### Scope

The world is undergoing a digital-data revolution. More and more data is born digital. Almost every business, government and organisational activity is driven by data and produces data. Science, engineering, medicine, design and innovation are powered by data. To thrive in this new environment requires new strategies, new skills and new technology.

ADMIRE has created a single, comprehensive platform for knowledge discovery.

As an example, analytical customer relationship management (ACRM) is an established field in business analytics. Its aim is to make use of operational customer data to enable the measurement, analysis and, ultimately, the optimisation of customer relationships – but it is struggling to manage the increasing tides of data in privately-held customer databases, and increasingly in public data sources such as social media sites. ADMIRE has applied data-intensive methods to ACRM in

the telecoms domain, demonstrating that not only can traditional knowledge discovery and business intelligence processes be supported by ADMIRE’s data-intensive architecture and tools, but also that it accelerates the development of methods and delivers agile data exploration.

### Advances

Since 2008, the ADMIRE Project has combined data-intensive strategies, skills and technology to create a single platform for knowledge discovery, combining data access, integration, pre-processing, data mining, statistical analysis, post-processing, transformation and delivery. ADMIRE has created:

DISPEL, a powerful, Java-like language for describing complex data-intensive workflows.

A Java and web-service-based software platform designed to run distributed data-intensive workflows in a streaming fashion, avoiding data bottlenecks.

Rich semantic descriptions of workflow elements based on a common network of ontologies.

A library of 150+ standard workflow elements for accessing, integrating, transforming and moving data.

Visual programming tools based on the eclipse platform.

A book, to be published late 2011 or early 2012, describing new strategies for data-intensive methods in business, government and science.

ADMIRE software is free, open-source and already in use in scientific applications throughout Europe.

### Positioning in global context

There are many firms and groups engaged in producing better tools and methods for knowledge discovery. We believe the strategies and technologies offered by ADMIRE are unique in the open-source world, and broader in scope and potential reach than many if not all commercial offerings in the data integration and workflow management space.



### Contribution to standardization and interoperability issues

ADMIRE is based entirely on web standards, using Java as a primary implementation language, http-based web services for communication and Eclipse as a basis for its tooling.

### Target users / sectors in business and society

ADMIRE is targeted at three categories of expert:

1. professional users of data who need to improve the path from data to knowledge within their own organisation;
2. professional developers of the tools, methods and software that comprise data-intensive workbenches and that accelerate the path from data to knowledge;
3. professional data-intensive engineers who build and run the systems that provide the computational power and data handling needed to accommodate the growing demand for data-intensive applications.

### Overall Benefits for business and society

ADMIRE aims to assist researchers in companies, governmental organizations and academia who are supporting strategists and decision makers or who are developing new knowledge.

### Examples of use cases

Traditional attempts to understand environmental risk have relied on large-scale simulation, of weather, for example. These simulations are

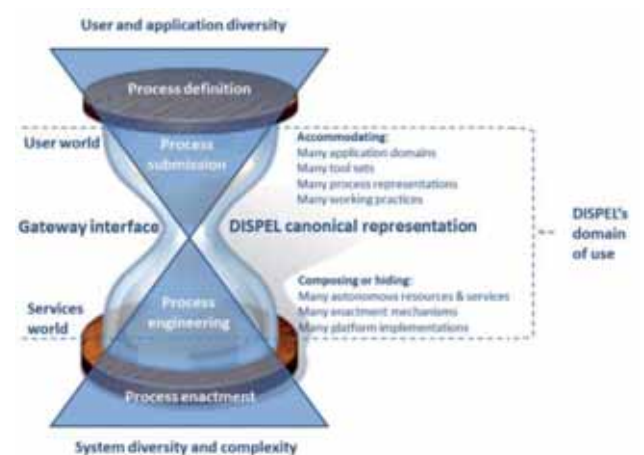
notoriously difficult, and in terms of supporting decision-makers in charge of managing a major river system, alternative approaches based on sifting historical and real-time data can offer more pragmatic solutions. ADMIRE has applied the use of data-intensive computing to the management of a river system and the prediction of river flood levels, providing domain experts at the Slovak Hydrometeorological Institute with access to new data sources and new knowledge-discovery tools. ADMIRE provides a single, comprehensive solution which integrates, assimilates, filters and analyses river data, weather data and other sources of information across multiple sites.

### Achievements

The ADMIRE Book.

The ADMIRE Software Platform – public beta.

The ADMIRE Workbench tools – public beta.





ALIVE



### At a Glance

#### Project title

Coordination, Organisation and Model Driven Approaches for Dynamic, Flexible, Robust Software and Services Engineering

#### Contact person

Javier Vázquez Salceda  
Universitat Politècnica de Catalunya  
jvazquez@lsi.upc.edu

#### Website

[www.ict-alive.sourceforge.net](http://www.ict-alive.sourceforge.net)

#### Total cost / EC contribution

3,777,385 € / 2,809,213 €

#### Start date /end date

February 2008 / October 2010

### Scope

There is a new emerging trend of Internet networked applications based on the notion of software-based services that are combined to create a wide variety of powerful software systems, from business e-Commerce environments to e-Government systems to social networking applications. But combining these services in non-trivial cases is a real challenge. There is a need for technological solutions that not only give support in the design of such systems but also provide mechanisms to dynamically combine and adjust the services to an ever-changing digital environment.

The ALIVE project is based around the idea that many of the strategies used today to organise the vastly complex interdependencies found in human social, economic behaviour will be essential to structuring future service based software systems. The project has combined cutting edge Coordination and Organisation mechanisms

(providing flexible, high-level means to model the structure of interactions between services in the environment) and Model Driven Design (providing for automated transformations from design models into multiple implementations) to create a framework for software and services engineering for “live” open service-based systems.

### Advances

The project has extended current trends in service-oriented engineering by adding three conceptual layers, described as follows:

- The **Service Layer** augments and extends existing service models in order to make components aware of their social context and of the rules of engagement with other services by making extensive use of semantic web technologies.
- The **Coordination Layer** provides the means to specify, at a high level, the patterns of interaction between services, using a variety of powerful coordination techniques from recent (European) research in the area.

The Organisational Layer provides context for the other levels – specifying the organizational rules that govern interaction and using recent developments in organizational dynamics to allow the structural adaptation of distributed systems over time.

The Model-Driven Engineering paradigm chosen in ALIVE greatly helps developers as it provides for automated transformations from models of these three levels to potentially multiple target platforms

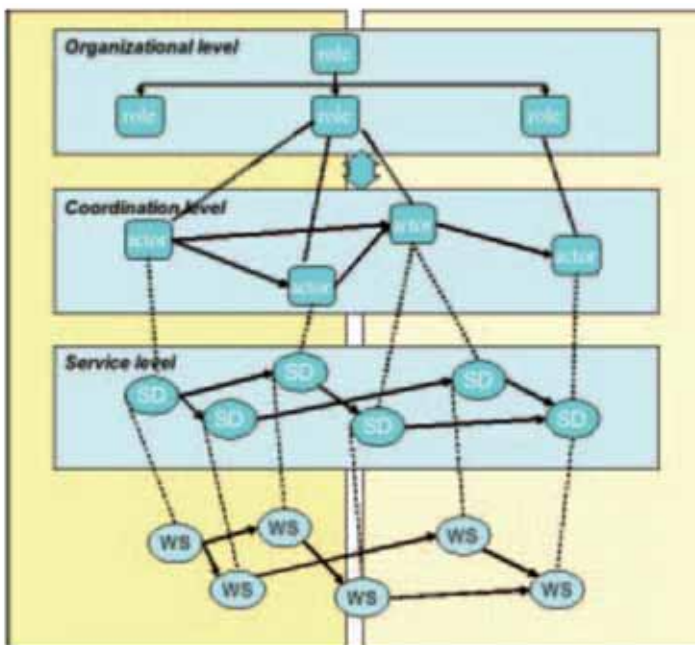
### Positioning in global context

The ALIVE approach and framework makes several notable contributions to the development of dynamic and robust service-oriented systems. In addition, the ALIVE approach is supported with a framework and tool-set that enables the smooth design of service oriented applications based on organisational and coordination concepts, their

integration with standard technologies and external tools as well as their (semi)automated creation from model abstractions. We also observe that systems engineered with our approach have distinctively autonomic and robustness features, in that the system looks after itself. Malfunctions at the service level trigger adjustments/changes at the coordination level (namely, re-planning), and, in more extreme cases, changes at the organisation level. In that way, the autonomic property allows our system to dynamically adapt and recover from errors.

### Contribution to standardization and interoperability issues

Members of the ALIVE consortium have contributed in early phases of the SOA-ML standard definition process-driven by the SHAPE project.



### Target users / sectors in business and society

Target users are *Software and Application Developers* willing to make more flexible and robust service-oriented applications.

### Overall Benefits for business and society

We expect the ALIVE results to be a relevant contribution to the Service Engineering field as:

- ALIVE promotes reliability and stability for service-oriented applications, by modelling multiple levels of a system, enabling designers to keep slowly changing elements separate from dynamic aspects of the application;

- ALIVE provides a sound paradigm for the deployment of service-based applications across multiple organizations;
- ALIVE models have the potential to structure service-based systems in such a way that they are easier to understand for non-technical individuals – paving the way for broader involvement in software and service design.

### Examples of use cases

*Dynamic generation of personalized information in interactive community displays:* Interactive community displays (ICDs) integrate urban information (in real time) and create public spaces in the Internet for people living in or visiting a city. However much of the digital content offered tends to be static and unpersonalised. By the usage of ALIVE technology we have created ICD's that can adapt to the user needs up to the point to recommend activities or services and get up-to-date information directly from the Internet.

*Policy making support for Crisis management through simulation:* Simulations of crisis scenarios have the potential to improve the organizational structures and the quality of response of the parties involved. As a crisis evolves, organizational structures must be systematically updated to reflect the changes in the numbers of personnel and the seriousness of the crisis. Simulation allows the effectiveness of these organizational changes to be evaluated. In the real world, simulations are very cost- and time-expensive as they typically use active emergency service personnel. Thanks to ALIVE we have created an emergency simulation tool for the Dutch emergency forces. This simulator effectively explores diverse crisis management scenarios when a natural disaster has escalating city- and nation-wide consequences, allowing policy makers to evaluate current and potential emergency policies and procedures.

### Achievements

*ALIVEclipse:* a set of metamodels and tools which implement architectural elements necessary for the three levels of the ALIVE Theoretical Framework. This includes the means to 1) specify and deploy the organisational rules of engagement which apply to a service environment; 2) the patterns of interaction between components within a service environment, and hence enable the connection of individual coordination patterns to the organisational objectives; and 3) enable services to manage their dependencies with other services.

*ALIVE Methodology:* a novel methodology for the design and development of service-based systems. The methodology includes tutorials and design wizards which provide methodological heuristics and guidelines.

# COMPAS



### At a Glance

**Project title**

Compliance-driven Models, Languages, and Architectures for Services

**Contact person**

Prof. Schahram Dustdar,  
Technische Universität Wien,  
Distributed Systems Group  
dustdar@infosys.tuwien.ac.at

**Website**

<http://www.compas-ict.eu>

**Total cost / EC contribution**

5,602,803 € / 3,917,723 €

**Start date /end date**

February 2008 / January 2011

*the underlying technology platform* that realizes business processes with embedded compliance concerns.

### Advances

Compliance solutions of today are built by addressing each compliance issue as a separate project, resulting in a somewhat unstructured, even chaotic collection of compliance-assurance applications. The COMPAS project realized an integrated software framework, comprising various components whose ultimate goal is to support the design, implementation, monitoring and evolution of compliance concerns, in a more systematic way. The software framework specifically enables realization of compliance in a SOA based system; moreover, the framework follows a model-driven approach, which means that new, domain-specific languages and models are the primary artifacts, and that the underlying source-code generation is automated. Domain-specific languages enable IT personnel to concisely capture, share and discuss implementations of compliance concerns with compliance experts. These languages are supported in the backend by models, which the framework leverages when it generates the actual underlying technology platform that implements the compliance concerns.

### Scope

Due to corporate scandals of the last years, business compliance, i.e., the conformance of business procedures with laws, regulations, standards, and best practices has become a major concern for corporations. Given the pervasiveness of IT in supporting everyday business activities, it is a natural step to investigate the possibilities of IT to support business compliance.

The COMPAS project proposed to design and implement novel models, languages, and an architectural framework to ensure compliance of business systems. The project achieved this in the context of Service Oriented Architecture (SOA) design principles, which is proliferating in today's business IT environments. One of the key achievements in the project has been to realize an extensible software framework that integrates the functions of *modelling and visualizing compliance concerns* for SOA, and automatically *generating*

### Positioning in global context

The integrated approach from the COMPAS research and products provide a systematic approach to addressing compliance through a single framework. Additionally, the research products are mostly built on top of open source software products. The project ultimately demonstrates that systematic compliance can be achieved in a cost-effective manner.

### Contribution to standardization and interoperability issues

The COMPAS project contributed two proposals for standardization, the Compliance Request



Language (CRL) proposal, regarding a language for specifying compliance, and BPEL4CFrags, a proposal of a specification format for reusable process fragments that enable implementation of compliance requirements in BPEL processes.

### Target users / sectors in business and society

Among end-users for the COMPAS research results we have; compliance domain experts, management, software developers, software architects, software / business process designers and software test engineers (verification and validation).

### Overall Benefits for business and society

The centralized, systematic solutions developed under the COMPAS project make improvements in the compliance implementation process in terms of: easing *maintenance and evolution* of compliance related code in the system, enabling *reuse of compliance code* that may be applied to different parts of an organization system, and improving on the *understandability and communication* of compliance implementations between the technical and business people in an organization.

### Examples of use cases

In executing the COMPAS project, a case study scenario from the telecoms sector was used to provide guidance and evaluate research results. The WatchMe scenario focuses on advanced telecom services offered by mobile virtual network operators (MVNO). The services offered by WatchMe include TV streaming services, which offer real-time combined audio and video streams to end-users. In order to deliver streams, the MVNO service (WatchMe) has to acquire them from the third party video and audio providers. For this purpose contracts have to be signed between the MVNO and 3<sup>rd</sup>-party providers. The contracts document license plans as well as a specification of quality requirements that have to be met – e.g. only predefined compositions of video and audio streams from different providers are allowed. The MVNO is also required to deliver a minimum Quality of Service (QoS) when streaming to end-users.

Using the COMPAS results, we capture the requirements in the form of a domain-specific language designed to express licensing compliance requirements. From the compliance specification, we are able to generate a monitor that can observe the MVNO services to detect violations in license compliance.



### Achievements

The following software prototype systems were realized in this project:

- An overview architecture defining integration points for compliance-assurance components.
- View-based Modelling Framework – a software framework for modelling compliance requirements, coupled with a number of tailor-made domain-specific languages for specifying compliance.
- MORSE – a model repository for storage of models and making them accessible for compliance checking at runtime.
- Compliance Governance Dashboard, Business Process Intelligence, BPath – Compliance monitoring and management facilities that provide visualisation and drill-down capabilities for compliance issues in a system.
- BPELC4Frag, Fragmento – concepts, language and runtime support for reusable process artefacts for compliance-assurance.
- Compliance Request Language, Reo, QuaLa, Security DSL, License DSL – domain-specific and formal, expressive languages for compliance concerns.



## DEPLOY



### At a Glance

#### Project title

Industrial deployment of advanced system engineering methods for high productivity and dependability

#### Contact person

Alexander Romanovsky  
University Of Newcastle Upon Tyne  
alexander.romanovsky@ncl.ac.uk

#### Website

<http://www.deploy-project.eu/>

#### Total cost / EC contribution

18,700,469 € / 12,403,399 €

#### Start date /end date

February 2008 / January 2012

### Scope

Formal engineering methods enable greater mastery of complexity than do traditional software engineering processes. It is the central role played by mechanically-analysed formal models throughout system development that enables mastery of complexity. As well as leading to big improvements in system dependability, greater mastery of complexity leads to greater productivity by reducing the expensive test-debug-rework cycle and by facilitating increased reuse of software. The overall aim of the DEPLOY Integrated Project is to make major advances in engineering methods for dependable systems through the deployment of formal engineering methods. The work is driven by the tasks of achieving and evaluating the industrial take-up of the DEPLOY methods and tools, initially in the five sectors which are key to European industry and society.

### Advances

DEPLOY is delivering methods and tools that support the rigorous engineering of complex resilient systems from high level requirements down to software implementations via specification, architecture and detailed designs; support the systematic reuse and adaptation of models and software thus addressing industry's requirement for high productivity and requirements evolution; have been field-tested in and adapted for a range of industrial engineering processes; are accompanied by deployment strategies for a range of industrial sectors; and are based on an open platform (Eclipse) and will themselves be open.

### Positioning in global context

DEPLOY is developing an advanced industry-strength toolset supporting rigorous design of complex critical systems. The applicability is demonstrated in four major industrial domains: transportation, business information, automotive and space.

The unique features of this toolset are in its open extendable architecture based on Eclipse, open source distribution, and novel approaches to supporting interactive modelling.

### Target users / sectors in business and society

Potential users are:  
Software and Application Developers  
Systems Integrators  
Technology Providers  
Service Providers

### Overall Benefits for business and society

The main benefits for the business are in the improved productivity and dependability of their products. Formal development provide evidence for ensuring quality of the critical products

### Achievements

DEPLOY is delivering the Rodin development environment (openly downloadable from <http://www.event-b.org/>) supported by a rich choice of documentation and tutorials ([http://wiki.event-b.org/index.php/Main\\_Page](http://wiki.event-b.org/index.php/Main_Page)) and a library of models demonstrating how the DEPLOY methods and tools have been applied

in developing medium- to large-scale industrial applications <http://deploy-eprints.ecs.soton.ac.uk/view/type/>.

The methodological guidelines, cookbooks and modelling patterns are provided in <http://deploy-eprints.ecs.soton.ac.uk/view/type/> to help the developers.



Dynamic Variability in complex, Adaptive Systems

DiVA



### At a Glance

**Project title**

Dynamic Variability in complex, Adaptive systems

**Contact person**

Geir Horn  
Stiftelsen Sintef  
Geir.Horn@sintef.no

**Website**

<http://sites.google.com/site/divawebsite>

**Total cost / EC contribution**

4,966,794 € / 3,510,409 €

**Start date /end date**

February 2008 / March 2011

### Scope

Software today is increasingly used in mobile devices ranging from portable computers, tablets, and smart phones, with many different operating systems and hardware choices within each category. This implies that the software must be able to adapt to different environments, and dynamic changes in context and use. Traditionally the developer would maintain different software releases for different platforms, and use “if” statements to deal with context variation of the form “if the battery level is low, then reduce the screen brightness”. However, in a mobile setting the multiplicity of simultaneously occurring events and context changes give rise to a very high number of conditions to consider. This makes efficient adaptive software costly to develop and difficult to maintain. DiVA changes this by offering a way to create such applications using a model driven development approach throughout, even at runtime, enabling the resulting software to reason about the current

context and reconfigure the application to a version that is momentarily deemed to give the best utility for the user based on high level goals and constraints set by the developer.

### Advances

DiVA offers the *software developer* a complete tool supported methodology for the model driven development of adaptive applications from the requirement engineering to the runtime platform independent configuration selection with platform specific deployment. The integrated solution, the *DiVA Studio*, is available as open source in the form of a set of Eclipse plug-ins for design time modelling and components to support runtime reasoning and application deployment on component platforms like OSGi<sup>1</sup> or Fractal<sup>2</sup>. Experimental results have shown that developers using the DiVA approach become more efficient and capable of handling adaptive applications that was previously unattainable. Furthermore, the focus on models allows transparent and rapid deployment of the application on multiple execution platforms.

### Positioning in global context

DiVA builds on state of the art research in areas like natural language processing for the requirements engineering; aspect oriented development for the adaptation models; constraint based reasoning for the configuration selection; and runtime weaving of the variability aspects with the application base model. However, the integration in DiVA is unique, and there is no similar holistic approach to adaptive software development.

### Contribution to standardization and interoperability issues

The DiVA Adaptation and Architecture Modelling framework will be taken forward in the ongoing

<sup>1</sup> <http://www.osgi.org/Main/HomePage>

<sup>2</sup> <http://fractal.ow2.org/>



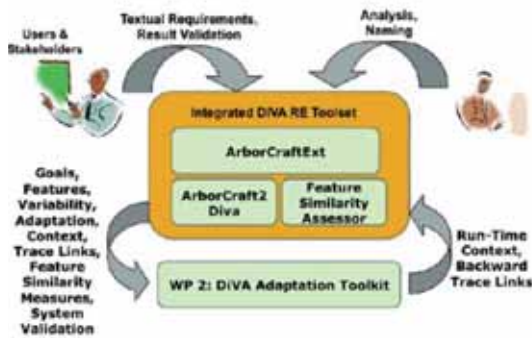
standardization effort within the OMG<sup>3</sup> on the Common Variability Language (CVL)<sup>4</sup>. DiVA has initiated and leads an Eclipse technology project to standardise the EMF<sup>5</sup> feature model API<sup>6</sup>.

### Target users / sectors in business and society

The target users of the DiVA methodology are software developers with a need to develop and maintain adaptive applications, possibly over multiple deployment platforms.

### Overall Benefits for business and society

Software is omnipresent and crucial factor for businesses and society. The same is mobility. DiVA has the potential to increase the productivity in the development of context aware and adaptive applications, and to reduce the cost of maintaining such applications over time.



### Examples of use cases

The DiVA methodology and tools have been tested with two real industrial applications:

1. A Crisis Management System for a civil airport. This is a decision-making tool, made to assist managers in taking decisions in such crisis situations. The system monitors the airport infrastructures and environment and automates a number of tasks to help humans to have the right level of information at the right time and let them focus on important decisions.

2. A Customer Relationship Management (CRM) application: A centrally hosted SaaS<sup>7</sup> CRM service platform offering user access in various ways, e.g., with a client application, a standard web browser running on a desktop PC or a laptop, via mobile access from a PDA or a mobile phone or directly from other applications using its web service interface.

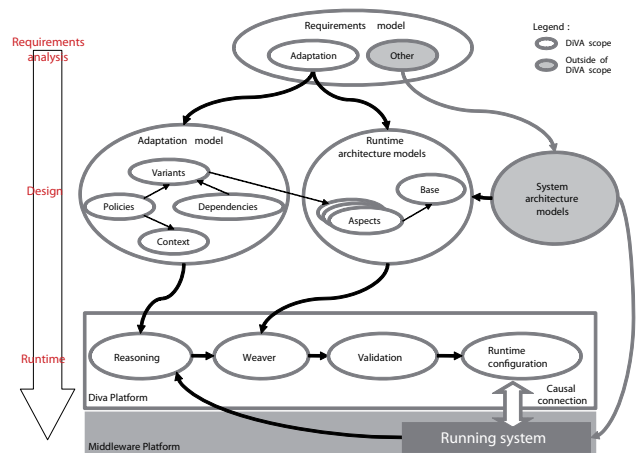
### Achievements

The main result of this project, the set of tools composing the DiVA Studio, as well as the case studies that make use of the technology developed this project are available through the project website.

The DiVA Studio provides an integrated tool-chain and an execution platform supporting the DiVA methodology. This tool supported methodology provides an advanced model based approach to design, implement and execute adaptive systems. The studio offer support for the different lifecycle phases of an adaptive system (from requirement elicitation to runtime management).

The integrated tool suite of the DiVA Studio is integrated with the Eclipse Integrated Development Environment. Most components of the DiVA Studio are publicly available under open-source licenses.

On the website includes full tutorial for installation and use of the DiVA Studio, as well as the possibility to download a full demonstrator based on a simplification of a Dynamic Customer Relationship Management (D-CRM) system. The demo allows the user to simulate the environment and visualize the decisions taken by the reasoning engine. Depending on the context, the system will dynamically adapt, so that the user will be notified by pop-up windows or voice messages.



<sup>3</sup> Object Management Group, <http://www.omg.org/>

<sup>4</sup> RFP with AB changes, document ad/2009-12-03: <http://www.omg.org/cgi-bin/doc?ad/2009-12-3>

<sup>5</sup> Eclipse Modeling Framework, <http://www.eclipse.org/modeling/emf/>

<sup>6</sup> Application Programming Interface

<sup>7</sup> Software as a Service



FAST



### At a Glance

**Project title**

Fast And Advanced Storyboard Tools

**Contact person**

Dr. Javier Caminero  
Telefonica investigacion y desarrollo sa  
fjcg@tid.es

**Website**

<http://fast.morfeo-project.eu>

**Total cost / EC contribution**

5,532,606 € / 3,564,151 €

**Start date /end date**

March 2008 / February 2011

European work in the Future Internet is related to the usage of available services by users. The research in front-end technologies like mash-ups and gadgets can provide valuable insight, here. Since mash-ups are expanding quite rapidly and with it the mash-up platforms. But still there is a lack of tools to create the visual interfaces of the services which will be deployed to these platforms. This is the place where FAST steps in.

We expect our tools and research to be carried on in other contexts now. E.g., to mention a number of follow-up projects that FAST consortium members are involved in: in the 4Caast Project, FAST-created gadgets can either play the role of an application provider, or the FAST platform can be used to assist the mashing-up of different applications. Similarly, the Cloud4SOA project can apply the technology and lessons learned for semantic interoperability of application components in FAST. On the other hand, projects such as Puzzled by Policy will be able to apply the FAST platform (or a derivative) to empower citizens to make better use of public sector data and services.

### Scope

The FAST project goal was to put the **tools to create small but complex Web applications** – call them enterprise mashups or situational widgets or gadgets – directly **into the hands of end users**, without having to rely on programmers and developers to do the job for them.

### Advances

The research, the design and the development to produce these tools have been done within the project, and significant efforts were spent to test and evaluate the tools, focused a lot on usability, and made sure they really work for the people who should use them.

### Positioning in global context

The European Commission, founder of FAST project, has pledged to have an important stake in the creation of the Future Internet. Indeed, Europe has the capacities to play a guiding role in creating new trends. One of the main fronts of the

### Contribution to standardization and interoperability issues

Despite the standardization efforts made by the consortium, no FAST result has been incorporated or influenced any standardization process. The most relevant standardization body for enterprise mashups, the Open Mashup Alliance, was contacted in order to influence the Enterprise Mashup Markup Language (EMML) with the FAST research results.

### Target users / sectors in business and society

An important inspiration for FAST came from the problem of business process management and business integration. For this reason, one of the main target audiences for FAST is businesses, both small and medium enterprises and



large corporations. The software tools developed within the project empower business consultants to speed up processes, organize management procedures or integrate with business partners and sub-contractors. Several scenarios of this type have been thoroughly and successfully evaluated during the course of the project. Similarly, the project vision and tools have been demonstrated to business audiences at several occasions, gathering a lot of interest.

### Overall Benefits for business and society

The use and application of FAST is in no way limited to businesses. Similarly, public sector organizations such as governments can use the platform to easily provide better services to their citizens. E.g., a project such as the recently started Puzzled by Policy (ICT 256261) can employ FAST to reach its goal of tackling the problem of detachment and disengagement between citizens and policy making.

At the same time, also private individuals can make use of FAST to create and use situational applications and therefore improve their daily knowledge work and activity on the Web. Also this possibility has shown great promise in the evaluations carried out, particularly in the third year of the project.

### Examples of use cases


Using visual programming method, developers begin with the front-end gadgets, meaning the user view and functionality; after finishing these components, the developer will proceed to visually establish the connection to back-end web services, this may eventually mean going through process execution flows. All of these programming steps are guided by the FAST tools delivering the best possible help to the developer.

The front-end gadgets created with the FAST tools are afterwards combined into a mashup. This will enable gadget users to have direct access to all required information using only this gadget inside the mashup-platform. The users work load can be simplified this way, since he no longer has to browse recurrent web resources to perform his daily tasks. The FAST complex gadget will automatically guide the user through his daily work steps. This will be seamless for the user and is carried out by a complex process flow which is part of the gadget itself.

### Achievements

During the life of the project, many advances beyond the state of the art in end user development and gadget creation have been carried out. As a summary it can be highlighted the following:

- **The gadget development method.** A user-centred method guided by end-user requirements and focused on end-user development of complex gadgets. It is inspired by agile patterns, but focused on easing the user experience in the process of design gadgets.



- **An innovative complex gadget architecture.** It is oriented to create compositional applications, using semantically enriched building blocks which allow publish and discovery using social catalogues. One of its main achievements is to offer automatic creation of coherent component control- and data-flows among building blocks using their semantic information.
  - **The Gadget Visual Storyboard.** The end-user tool to easily create complex gadgets from its lower level building blocks until its deployment into a wide range of destination platforms with the aid of the storage and deployment service.
  - **A service wrapping and semantic mediation operator designer.** A semi-automated tool to create service wrappers and semantic matching and mediation operators.
- All these results are accessible through the project website.

# FLOSSINCLUDE

## FLOSS [include]



### At a Glance

#### Project title

Free/Libre/Open Source Software:  
International Cooperation  
development Roadmap

#### Contact person

Meiszner Andreas  
Universiteit Maastricht  
andreasmeiszner@gmail.com

#### Website

www.flossinclude.eu

#### Total cost / EC contribution

1,007,254 € / 703,465 €

#### Start date /end date

February 2008 / July 2010

### Scope

Free/Libre/Open Source Software (FLOSS) holds great potential for developing countries. FLOSS uses open standards avoiding lock-ins and allowing for flexible solutions. Programs can easily be translated into local languages. The international developers' community acts as a free-of-charge training environment. Yet FLOSS use and development in the global South has so far stayed below its potential. The FLOSSInclude project therefore aimed to help developing countries to use FLOSS to their advantage. The project started in February 2008 and ran for 2 ½ years studying needs to increase the deployment, development and societal impact of FLOSS in Africa, Asia and Latin America. Upon identified needs, the project partners developed and tested innovative solutions to satisfy specific local ICT needs within Africa, Asia and Latin America. The results of those works have been taken forward in a Roadmap outlin-

ing future possibilities for research cooperation between the EU and developing countries. The project consisted of five core tasks, from analysis to Roadmap:

1. analysis of available data to identify key problem areas and areas of blocked potential for FLOSS in the target regions
2. dissemination and networking, to identify and federate local and regional initiatives
3. requirements analysis, to show with concrete cases the specific technical, business and socio-political needs for the growth of FLOSS use, deployment and development in target regions
4. pilots, to ensure that FLOSS solutions, tools and services can be cost-effective and practical
5. research and cooperation Roadmap, supported by regional initiatives, concrete cases for clearly identified requirements, with solution areas proposed that have been validated through pilots. Validated pilots and a coordinated Roadmap for future EU development research cooperation aim to ensure that the impact of FLOSSInclude will be sustained beyond the duration of the project.

### Advances

The benefits offered by FLOSS (Free/Libre/Open Source Software) appear to be particularly useful for developing countries around the world. However, in many cases it remains below its potential due to various factors; which are hurdles to the widespread adoption of FLOSS. During the work process of the FLOSSInclude project one objective had been to identify the main factors and enablers that would help increase the benefits of FLOSS in the target regions (Africa, Asia and Latin America) and how they can be influenced to use more FLOSS tools.

The FLOSSInclude project has been providing some evidence that in order to foster the adoption of FOSS within the target regions a number of very diverse areas needs to be taken into account, such as skill development, local-

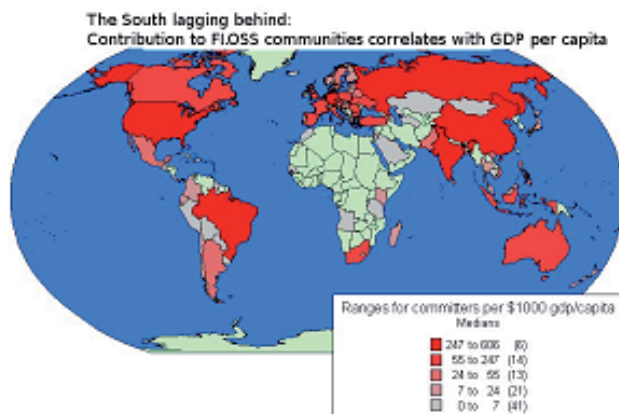
sation, information, documentation, promotion and dissemination. Hence, to comply with all these requirements successfully demands a lot from the various actors in the field, such as professionalised organisations, international support, research, funding, etc. The work submitted during the FLOSSInclude project by the local partners as well as concrete examples of what measures and activities should be undertaken and a prioritization of those is further provided within the FLOSSInclude Roadmap.

### Contribution to standardization and interoperability issues

The need for a widespread awareness of FLOSS and its benefits to society had been identified as an important issue across all areas. Though FLOSS offers great opportunities for education and businesses a lack of awareness often hinders the introduction of free software solutions in these areas. The lack of people skilled in FLOSS is however another severe limitation for the greater spread of FLOSS. In many cases, FLOSS applications that developing countries desperately need exist already but still have to be adapted and localized into the native languages of the users. Regarding technology issues the main challenge in many target regions appear to be the lack of affordable broadband connectivity. Finally, a key factor to making possible the spread of FLOSS is through the support from governments. However, often there seems to be a very low awareness of FLOSS amongst government officials and in many cases, policy environments are hostile to the adoption of FLOSS.

### Achievements

One of the core outcomes of the project has been the FLOSS Research and Cooperation Roadmap 2010-2020 that aims to serve as a guideline for future research cooperation between Europe and developing countries in order to fill the gaps identified by the FLOSSINCLUDE project. Like the FLOSSINCLUDE pilot reports, this Roadmap covers the technical, social, economic, cultural and political dimensions of FLOSS. The basic question that this document tries to answer is: What sort of FLOSS-related R&D in developing countries appears to be useful in the near, mid-term and long-term future and how could it be supported by the EU through international cooperation? The Roadmap highlights approaches likely to be successful as demonstrated by the FLOSSINCLUDE pilots. A remarkable characteristic of this Roadmap that particularly derives from the experience, attitudes and capacities of the FLOSSINCLUDE partners in the project's target areas, is that it is not guided by the expectation that needs only exist in the developing countries and means only exist in Europe. The basic lesson that all participants in the FLOSSINCLUDE project have learned is that both regions dispose of needs and capacities that can be tapped in a mutually beneficial way. The Roadmap is therefore not limited to European funding, and it does not solely address policy makers as enablers of useful future research and collaboration. Rather than that, a wide array of possible activities and actors, namely the FLOSS community itself, businesses, and educational institutions, is taken into consideration in order to ensure that the proposed future research and collaboration activities will have an impact and will be useful to the target regions independently.





## At a Glance

### Project title

Interactive Real-time Multimedia Applications on Service Oriented Infrastructures

### Contact person

Mr. Smithson Stuart  
Xyratex Technology Limited  
stuart\_smithson@xyratex.com

### Website

www.irmosproject.eu

### Total cost / EC contribution

12,759,911 € / 7,902,430 €

### Start date /end date

February 2008 / January 2011

- Applications resource needs are not modeled, predicted and estimated to capture in a coherent and consistent way their real-time attributes

### The IRMOS Solution:

- The IRMOS Team has developed a set of tools and a platform that enables the real-time execution of interactive multimedia applications.
- This includes resilience mechanisms to allow the platform to be not only QoS-aware but also QoS-guaranteed!

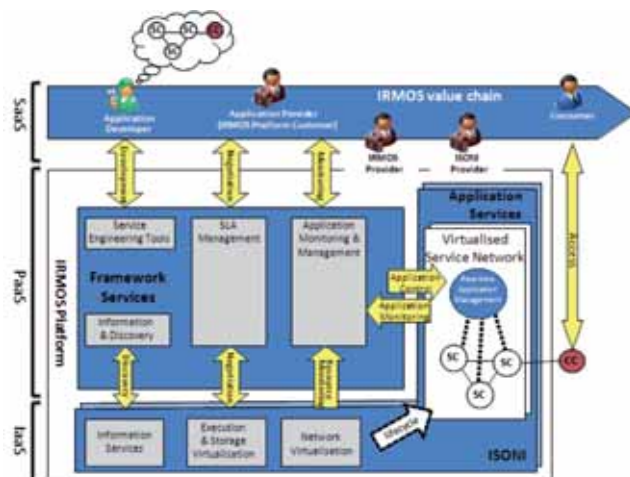
Future Internet applications raise the need for environments that can facilitate **real-time** and interactivity without major modifications in the application domain. IRMOS provides Software as a Service solutions to enable the adaptation of applications for cloud-based platforms. IRMOS provides a Methodology for creating or adapting interactive multimedia applications for virtualized service infrastructures along with an Application Wrapper for service-oriented applications.

## Scope

Today's Service Oriented Infrastructures (SOIs) lack Real-Time (RT) capabilities. This will limit the scope of Cloud Computing Solutions to emerging opportunities, particularly with Multimedia Applications. IRMOS has generated a framework which overcomes these problems and offers solutions for guaranteed Quality of Service in a Real-Time Environment.

### The Problem Areas:

- Service Oriented Infrastructures (SOIs) are not dynamically configurable and adaptable to Real-Time (RT) requirements
- Timing and interaction issues are not thoroughly studied and formally expressed in SOIs
- Web Services lack Real-Time capabilities



## Advances

This project has carried on research in several fields and key innovations have been realized in the various areas. Among them the development of virtualization techniques and intelligent networking that enables the connection of components through an intelligent network to improve efficiency through the autonomous deployment of services. Another important one is the negotiation and monitoring of automated Service Level Agreements within a QoS aware environment. Features have been developed within a Service Orientated Infrastructure that supports Real-Time Attributes. And last but not least, there is the implementation of software tools and modeling environments to enable real-time interactive applications to be written to target the IRMOS framework.

Globally the project has contributed to international standards.

## Overall Benefits for business and society

Besides the specific benefits generated by the target applications within the project, the technology generated by IRMOS may be utilized for alternative purposes or redeveloped for use in new markets. The realised IRMOS approach consists of a set of cloud solutions which will allow the adoption of interactive real-time applications, and especially multimedia applications, enabling their rich set of attributes (from time-constrained operation to dynamic service control and adaptation) and their efficient integration into cloud infrastructures.

The IRMOS solution has been realised in form of a set of software bundles targeting Software-as-a-Service, Platform as a Service and Infrastructure as a Service markets, which can each be exploited independently.

## Examples of use cases

IRMOS has demonstrated its functionality in three different application areas:

- Collaborative Digital Film Postproduction.
- Virtual and Augmented Reality.
- Interactive collaborative e-learning.

Details about these uses cases can be found on the project web-site.

## Achievements

The IRMOS Team has developed a set of tools and a platform that enables the real-time execution of interactive multimedia applications whose Resil-

ience mechanisms allow the platform to be not only QoS-aware but also QoS-guaranteed! These outcomes are available now.

Software as a service provides an application wrapper for service orientated applications and provides a blueprint for creating interactive applications for virtualized real-time service infrastructures.

Although targeted towards real time Multimedia applications the project outcomes may be utilized over a broad range of software application areas, particularly those requiring a guaranteed Quality of Service performance. This has been demonstrated in all the target applications, particularly the video post production market where remote interaction by multiple users on a single target in real time is commercially a significant step forward in reducing the cost of bringing a video to market using global resources.



Full descriptions and software downloads are available from the website <http://irmosproject.eu/>

In particular the IRMOS Project Repository in SourceForge is hosting the OSS components of IRMOS. They are classified within four major bundles containing more than fifty software components.

They are also accessible through the IRMOS web-site, where you can also find contact information related to non-open source components.

The IRMOS group in LinkedIn is open to everybody who is interested to join it, both within and outside the IRMOS consortium. The group is accessible through <http://www.linkedin.com/groups?mostPopular=&gid=2099046>. The group will continue being used by its members as a discussion forum.

# MANCOOSI



### At a Glance

**Project title**

Managing the complexity of the Open Source Infrastructure

**Contact person**

Roberto Di Cosmo  
Universite Paris Diderot – Paris 7  
roberto@dicosmo.org

**Website**

www.mancoosi.org

**Total cost / EC contribution**

4,427,855 € / 3,323,373 €

**Start date /end date**

February 2008 / May 2011

### Scope

Did you ever install an upgrade on your PC, just to find out afterwards that something very important was not working anymore? Or did you ever get stuck in the middle of a software update, with some arcane error about missing packages or software components ? In today’s rapidly changing world, software updates are becoming ubiquitous but in practice they fail all too often to complete successfully.

The problem is that there is no general way to know if a complex software package, with all its dependencies on obscure libraries, configuration files, hardware components or security setup, will work properly on a specific computer. At least, there is no way to know it BEFORE actually installing it, and before finding out the hard way that it has destroyed something else.

MANCOOSI worked hard to find reliable solutions to this problem, focusing in particular on upgrades of systems built using Free and Open Source Soft-

ware (FOSS), which is developed in a cooperation of volunteer programmers, public institutions, and software industry. This mode of development guarantees transparency of the inner working of the software, and encourages users to become contributors themselves. FOSS usually evolves in a rapid succession of new versions. Software developed and published in this way is bundled into packages by software distributions for easy deployment and installation, as for instance distributions of the GNU/Linux operating system and its huge collection of application software.

These collections of available packages are constantly updated for the benefit of the user who needs the latest versions of software products, for example to fix security issues. Consistent distributions of these software packages are published by editors, which may be released in updated versions several times a year, but over time each user’s machine inevitably ends up with a unique mixture of software components coming from different versions of her preferred distribution, and also from external sources. Maintaining the set of components installed on a machine in a consistent state is a daunting task. The quick evolution of the software base makes upgrading a recurrent process that needs to be both safe and easy for home users who do not want to spend time on system administration, and as smooth as possible for power users who cannot afford system downtime due to failed upgrades.

This is a difficult demand: a free software distribution is made up of dozens of thousands of software packages, interconnected by complex dependency and conflict relationships which change with the evolution of the software packages. As a consequence, even if a user wants to upgrade a single package it may be necessary to remove or install many other packages on the system. The process of finding an upgrade path of a complex software installation is an algorithmically complex task, and finding a path taking into account the preferences or resource constraints of the user is even more difficult.



## Advances

Mancoosi has set up a virtuous circle among users, developers and researchers, to ensure that the problems coming from software upgrades are continuously studied, and the tools improved in the long term.

**Better installation plans** On the one hand, a simple, yet powerful preferences language has been designed to allow any user to specify the kind of updates that is sought: if you want minimal changes to your system, or on the contrary, if you want the latest and best components, taking the risk of a more significant change to your system, now you can say it, and the advanced solver engines developed thanks to Mancoosi will take your criteria into account when planning the update.

Installation tools for GNU/Linux distributions have been adapted to allow the users to anonymously report installation problems to the central server of their distribution editor; a selection of these problems is made available, using a common format, CUDF, to researchers worldwide, that participate in an annual solver competition.

All the algorithms participating in the competition can then be plugged as solver engines in the following generation of user installation tools.

**Simulating upgrades, and undoing system modifications.** On the one hand, new theory and tools have been developed to model more accurately the installation process, in order to be able to know, without installing a piece of software, whether its installation will cause problems on a given configuration; this is not easy, and in complex cases can be a very difficult problem to solve. This is why it is

also necessary to build tools that allow to undo the system modifications made by an upgrade, when the result is not satisfactory.

Mancoosi has built models of GNU/Linux distributions that can be used to simulate faithfully package upgrades before installing them, and tools using these models to undo the upgrade. A set of portals allow the users to participate in the process by submitting rules to detect failures and by helping modeling each software components, thus improving the models and the tools.

The results of the project are being actively integrated in the software update management tools of mainstream GNU/Linux distributions, like Debian, Mandriva and Caixa Magica.

## Achievements

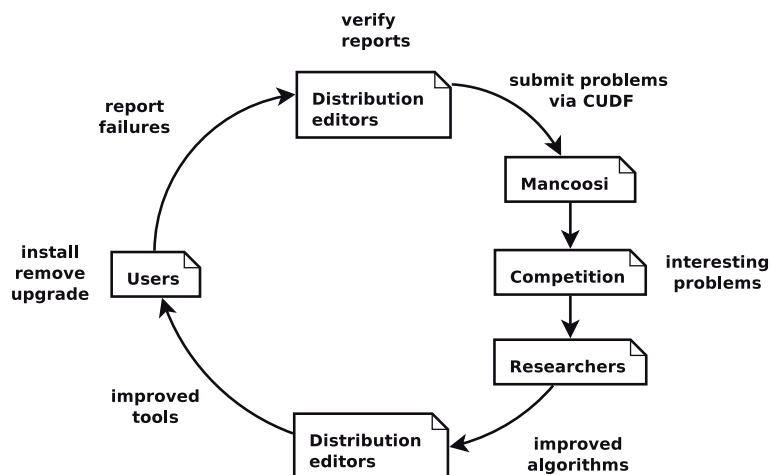
Mancoosi has produced many software components, either directly developed by the consortium, or indirectly fostered through the virtuous cycles established during the project. We mention here only the names of the software packages which are today mature and ready for production usage. The source code, distributed as Free and Open Source Software, can be found online at <http://www.mancoosi.org/software/>

**libCUDF:** reference implementation of the CUDF format specification;

**Dose3:** libraries for managing distribution packages and their dependencies;

**mccs, Packup and apt-pbo:** CUDF solvers;

**EVOSS:** system implementing a model-based approach to support the upgrade of FOSS systems.





# M:CIUDAD

mciudad



## At a Glance

### Project title

A Metropolis of Ubiquitous Services

### Contact person

Maribel Narganes  
Fundacion Tecnalía Research & Innovation  
Telecom  
maribel.narganes@tecnalia.com

### Website

[www.mciudad-fp7.org/](http://www.mciudad-fp7.org/)

### Total cost / EC contribution

3,970,270 € / 2,640,866 €

### Start date /end date

December 2007 / November 2010

## Scope

Nomadic users need mobile services that provide relevant and valuable information: the right content at the right place. Traditional web-based approaches are too static, complex and rigid for mobile users who have to navigate deep into the site to find out what is relevant. Another important challenge in the mobile environment is the fact that the information value dramatically decreases with variables like time, location or user preferences. The scale and dynamism of this information is key for nomadic usage and innovation is needed in order to combine personalization and relevance.

On the other hand, nowadays, mobile terminals have significantly improved its capacities in terms of processing power, interaction methods (e.g. touch screens), capabilities (e.g. location, cameras, sound recorders, NFC readers, etc.) and network connections (close to the always-on paradigm). This makes possible that, the increasing trend in the “fixed” Internet of the “do it yourself” concept, could be applied and continue also in the mobile environment; mobile users could use a device

that always wears not only for accessing to services and presenting data but for providing valuable content to other users.

The creation and access to content in a distributed way may offer the user the expected experience in a mobile environment, and therefore m:Ciudad proposes to turn users into service providers and mobile devices into a service platform, in order to benefit from the vast number of potential mobile sources of the right, key information.

## Advances

The m:Ciudad project introduces a new paradigm of communication between mobile users, setting up the basis for engineering a metropolis of ubiquitous mobile services, relying on a new service infrastructure. While most of new successful Internet services offer User Generated Content, m:Ciudad proposes a native and enhanced infrastructure for User Generated and Provided Micro-services in the mobile environment. m:Ciudad provides a service framework architecture, a set of mobile tools and a platform to allow users:

- To create focused, knowledge-based mobile micro-services;
- To become a service provider using his mobile device;
- To search for distributed services and to reach volatile generated content;
- To exploit the business opportunities that this new scenario brings about.

Micro-services are:

- Small and sharply focused mobile applications which allow user to share personal content (opinions, recommendations, context-based pictures and so on) with fellow users;
- Created and consumed by end-users using their mobile devices, encouraging spontaneous creation of information delivery services;
- Fully running on end user mobile devices;
- Potentially sharable and downloadable.

The main innovations of m:Ciudad are:

- A flexible service description framework for mobile micro-services.

- A novel creation and editing method to create new micro-services on the mobile device, encouraging user-generated-services.
- A new point of view for information provisioning and service creation based on the prosumer concept.
- An optimised service search environment based on semantic data.
- An embedded execution platform for mobile devices, to execute micro-services.

### Positioning in global context

Whereas traditional mobile services are created by professionals focusing on general purpose user needs, m:Ciudad is the first project that faces the concept of User Generated Services, which means end-users are able to create and provide service tuned to their actual needs, providing tailor-made services that will satisfy a Long Tail demand.

### Contribution to standardization and interoperability issues

Although m:Ciudad project results have not made direct contributions to standard bodies, important contributions have been generated in:

- micro-service language description and deployment procedures, which may be an input to W3C in order to complement standardisation on widgets and micro applications;
- micro-service execution environment development and relevant knowledge distribution based on presence and availability, which may be relevant for different working groups in OMA.

Both the reference architecture and the implemented tools and components have been developed following standards (well-defined APIs, web-based languages and protocols) and promoting a non-proprietary solution.

### Target users / sectors in business and society

Potential users may be:

- End-users, who are final users of tools and platform capacities;
- Software and Application Developers, who may design and develop platform upgrades;
- Technology providers, who may introduce relevant innovation on knowledge-based components;
- Service Providers, who may exploit their offer as service capabilities.

### Overall Benefits for business and society

Encouraging users to become creators and providers of mobile services will promote that new and flexible business models will emerge. Therefore, a different approach from the traditional “walled garden“ will be required. m:Ciudad will provide an open framework with flexible and profitable business models that will benefit m:Ciudad mobile service providers. This will create a “fun for all” scenario where service creators will be willing to contribute not only to provide useful services for the community but also to get real profit.

Also, this opens the opportunity for third parties for offering complementary services over the mobile platform, as

many others did in the fixed, allowing for a true syndicated business ecosystem.

### Examples of use cases

For the final evaluation of the resulting m:Ciudad framework with real end-users, a set of case studies (story lines) were defined and executed in several laboratory trials. The goal was to give users the opportunity to create and use micro-services by themselves, including the creation, sharing, executing, discovering, installing and deletion of micro-services.

Two scenarios were selected to assist the evaluations:

Mobile Blog was a single-provider micro-service with private service provision and public service consumption;

Friends Locator was a multi-provider micro-service with private service provision and private service consumption.



### Achievements

Project main results include:

- The m:Ciudad Reference Architecture which aims to serve as a guideline to build an ecosystem of user provided services.
- The m:Ciudad Implemented Architecture which has been fully developed in the project and includes the associated set of tools for covering the lifecycle of micro-services.
- The Service Creation Kit (SCK) which is the mobile application used to create and publish new micro-services.
- The Embedded Execution Platform (EEP) which executes micro-services, both for consumption and provisioning.
- The Searching and Recommendation Engine, which supports semantic search mechanisms and provides efficient search and recommendation methods.
- The Knowledge Warehouse, where knowledge-based description of micro-services and associated metadata are stored.
- The Service Warehouse, where micro-services and associated resources are stored and managed.
- A formal Micro-service Description Language, which is a specific service description language focusing on user-driven mobile service creation, powerful service discovery and platform-independent service execution and rendering.
- A white paper on “Mobile U+ Service Technical Guidelines”



# MOST



## At a Glance

### Project title

Marrying Ontology and Software Technology

### Contact person

Mr Miksa Krzysztof  
Comarch S.A.  
Krzysztof.Miksa@Comarch.Com

### Website

[www.most-project.eu](http://www.most-project.eu)

### Total cost / EC contribution

5,541,703 € / 3,729,258 €

### Start date /end date

February 2008 / April 2011

## Scope

Global competition challenges the competitiveness of the software industry, requiring increased efficiency and quality of software development, improved management of huge and complex software systems, and customer-tailored software products. Model-driven software development (MDS) provides means for separating business and implementation aspects of software at different levels of abstraction. However, key questions related to ensuring consistency within the development of large and complex software systems cannot be answered by today's technology.

*The goal of MOST is enable guided, consistency-preserving software development by bringing results of logic into mainstream software development.*

Ontologies can be used to capture and reason about the knowledge of domains, to evaluate constraints, to prove consistency, and to guide software engineers. Ontologies can also be used to generate proposals how to complete a specification in a reasonable way.

## Advances

The transition from untyped languages to typed languages in the 60s and 70s advanced both the efficiency and reliability of software development as it enabled the validation of type safety during the compilation of software programs. In the 90s the introduction of explicit constraint languages further extended the expressiveness in defining constraints for software specifications. The MOST project continues this evolution by introducing an Ontology-Driven Software Development (ODSD) methodology. This methodology is based on two basic advances.

First, MOST seamlessly integrates MDS with ontologies and reasoning. This integration enables the application of reasoning technology to deliver consistency-preserving languages. These languages provide means for the specification of consistency rules that are based on formalisations in description logics. As well, they employ reasoning tools to ensure the validity of system specification and introduce rich feedback for explaining and repairing inconsistencies.

Second, ODSD is complemented by an integration of consistency-preserving languages with software development processes in order to leverage guidance for software and tool developers. Based on this formal approach, ontology-based specification of artefacts, roles, tasks, dependencies in software development processes and an explicit monitoring of the development status, reasoning technology can be used to guide software developers and suggest due tasks.

### Positioning in global context

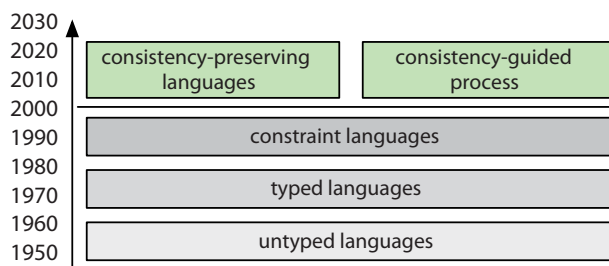
MOST is an early innovation in the field leading to more efficient and higher-quality software development with a better integrated understanding of the resulting software product. While currently the MDSD market is mostly driven and led by the US software industry, Europe has an advantage in ontologies and reasoning techniques. By providing a systematic way of bridging these two technologies, MOST delivers fundamental results that can be leveraged by European software vendors.

### Contribution to standardization and interoperability issues

MOST has been involved in the creation of standards with the two most influential standards bodies in the areas of Semantic Web and Software Engineering: W3C and OMG. The influenced standards were: OWL2, SPARQL1.1, RIF API4KB and ODM.

### Target users / sectors in business and society

These results are considered valuable both for vendors of development tools and software developers. ODSD advances current techniques for metamodelling, i.e., the development of consistency preserving modelling languages and introduces advanced means for monitoring and predefining the software development process. This eases the implementation and customisation of development tools. Software developers can exploit ODSD to challenge problems of software development that still cause tremendous effort for the software developer despite the adoption of MDSD. The results can be used by the software industry in their offerings for any sectors in business.



### Overall Benefits for business and society

Three industrial case studies have been concluded with promising assessment of MOST results. SAP case studies, applying MOST results to ensure consistency of business processes refinements, revealed that the process modeler can be **three times as fast** (productive) as before and make **significantly less errors**. The case study of Comarch, on providing the consistency preservation for the development with Domain Specific Languages, showed over **80% percent improvement** both in the terms of **productivity** and **quality** (time spent on bugfixing).

### Examples of use cases

Let us take an example of a usual situation, when one of the models is invalid. The error can be detected using consistency checking reasoning services. To start the repair actions, the process guidance can then create tasks for appropriate user roles. However, finding a valid solution for the error often requires the deep knowledge about the models and the modeling method. With MOST approach, the assigned user can use the justification reasoning services to detect the root cause(s) of the problem. Then, the proposed suggestions for a valid replacement of the erroneous model elements can be generated by leveraging classification and instance retrieval. After the model is fixed, the guidance engine acknowledges the successful execution of the repair task and may proceed to guide the user to next stages in the development process.

Thus, with MOST technology, model validation can, be made easy by bridging the modeling languages with ontologies and taking profit from the power of semantic reasoning. Additionally, once the languages are integrated, other important services such can be easily made available:

- The errors in the rule base itself can be detected.
- The users can ask for proposals for valid specifications in their model.
- The process guidance can proactively suggest next steps in the modeling process.

### Achievements

MOST durable results are of two kinds: proprietary tools owned by the industrial and scientific partners, and freely available tools. The latter includes:

- TwoUse (integrated modelling with UML and OWL2) <http://code.google.com/p/twouse>
- OntoMoPP (Eclipse tools for ontology engineering. EMFText concrete syntaxes for OWL2, SPARQL, SWRL, OWLCL) [http://emftext.org/index.php/EMFText\\_Concrete\\_Syntax\\_Zoo#Ontology\\_Languages](http://emftext.org/index.php/EMFText_Concrete_Syntax_Zoo#Ontology_Languages).
- DEFT (development environment for self-updating documentation) <http://deftproject.org>
- OWLText (syntax- and semantics-enabled metamodelling) <http://www.emftext.org/index.php/OWLText>
- TrOWL (Tractable reasoning infrastructure for OWL 2 (AGPL licence for open source applications) <http://trowl.eu>

As well, the MOST project delivered the MOST TOPF a product-line approach for the development and customisation of ODSD tool environments. The MOST TOPF was itself implemented using and ontology-based development methodology which demonstrates the benefits of ODSD in product-line engineering.

# NESSI 2010

**NESSI-2010**  
contributing to



## At a Glance

### Project title

Network European Software and Services Initiative

### Contact person

Mr. François- Marsal Bruno  
Thales Services Sas  
bruno.francois-marsal@thalesgroup.com

### Website

www.nessi-europe.com

### Total cost / EC contribution

1,559,996 € / 986,351 €

### Start date /end date

April 2008 / April 2010

## Scope

NESSI 2010 is a support action whose ambition is to sustain the momentum of a community which represents a **1.7 Million strong workforce** and **490 B€ annual revenues**. This community represents leading players including industries, SMEs, Academia and users **sharing the vision** of a common long term strategy on **software and services** to contribute to Europe's competitiveness, job sustainability and quality of life.

Such a community operates together in the context of NESSI – Networked European Software and Services Initiative, the European Technology Platform dedicated to software and services. The overall activities of NESSI cover a wide range of areas from **structuring research, building links to coordinate to national and international programmes, delivering an open service framework** etc.

NESSI 2010 is positioned in a context of *continuing support* – with all the industrial partners of NESSI providing financial resources, all the partners and

members of NESSI providing human resources and a subset of partners delegated by NESSI, namely Atos Origin, Engineering, Thales and TIE, operating selected activities through the NESSI 2010 support action, with a request of funding by the European Commission.

## Advances

NESSI – the Networked European Software and Services Initiative – was launched as a European Technology Platform in September 2005. Coordinated by 22 partners, it unites a community of 300 organisations from industry and academia active in Information and Communication Technologies.

With the support of NESSI 2010, NESSI has defined its strategy to support its community effort over the period 2008-2010. NESSI 2010 results are organized around the support provided to the community including links to national and international initiatives, support of the SME community as well as communicating and disseminating activities and definition of the Strategic Research Agenda which provides the technological vision and implementation path of NESSI.

## Positioning in global context

The goals of NESSI have been clearly defined from its initial vision: to enable a global environment in which all ICT stakeholders, small and large, users and providers, can deploy and locate novel services to the benefit of all. While the underlying infrastructure was identified as the Internet, it is clearly not the existing Internet – which does not offer the minimum level of quality, the required level of trust and security, the capability of continuous and uninterrupted availability. The goals of NESSI are therefore not only to build a service framework that fosters innovation – but also to ensure that the framework can rely on an end-to-end infrastructure that provides all the required features. To this effect, NESSI engaged since 2007 in a common work with the other ETPs eMobility, EPoS, ISI and NEM. This work picked up speed in 2008, and reached a point in time where the

first common vision for the so-called Future Internet was presented in November 2008, and a common event, ServiceWave 2008, was designed to reinforce this cooperation.

In this context, the goals of NESSI 2010 are derived from NESSI's goals as:

- maintaining and expanding the dynamics of the NESSI community
- providing visibility for the SRA – both within and outside the NESSI Community
- ensuring the largest possible community involvement – including that of national and international initiatives, that of the ICT SMEs which constitute a major source of innovation and for which faster deployment and easier access to a larger market base constitute key benefits

### **Contribution to standardization and interoperability issues**

NESSI is committed to interacting with recognised de jure and de facto standards bodies in the S&S field at both the EU and International levels. In 2007 NESSI formed a standardisation committee, COSTA, which has been led by Siemens and has been very active in pulling together NESSI participants interested in this area as well as external parties. COSTA is also promoting NESSI to standards organisations and individuals to ensure that it can be perceived as a focal point for linking all the technologies others are progressing.

### **Target users / sectors in business and society**

NESSI hosts a wide spectrum of different ICT companies that have been classified in 4 segments:

- Industry (for vertical markets)
- Software Vendors
- Telecommunication operators
- IT services

### **Overall Benefits for business and society**

Software is an essential part of an increasing number of products and solutions and thus one of the main drivers of the European economy. Starting from Europe's strengths and weaknesses, opportunities and threats NESSI 2010 supported the analysis of the future trends and the picture of the new software and systems landscape where the development of the future internet and its networked and service oriented systems are one of the most important drivers.

Hence NESSI 2010 actions benefit to research projects, members and EU institutions with a focus on supporting the dynamics of a software and services community and the creation of a strong research coordination.

### **Value proposition for projects**

The value proposition for projects is detailed along five points:

- more visibility for less cost
- community networking
- training support
- sustainability of their results
- channelling of their results towards research or exploitation opportunities

### **Value proposition for members**

- community networking
- training support

### **Value proposition for European institutions and public authorities**

This value proposition focuses on the role of NESSI as a source to present consolidated positions from industry and academia to public authorities in charge of research programmes, economic and strategic agenda such as the Digital Europe or the definition of the Future Internet etc.

### **Achievements**

#### **ServiceWave conference series**

The ServiceWave conference series ([www.service-wave.eu](http://www.service-wave.eu)) is a world-wide event dedicated to the emergence of services in the Future Internet. Service-Wave is a joint industry-academia conference combining a strong scientific track and an innovative industrial program. It is a forum for researchers, industrial practitioners, educators and national/international institutions representatives to expose the latest groundbreaking research works, recent innovations and trends in service oriented computing. Service-Wave fosters the creation of cross-community scientific excellence by gathering academic and industrial experts from various disciplines such as distributed systems, computer networks, business process management, wireless & mobile communication networks, grid & cloud computing, networking, service science and software engineering. ServiceWave is thus the forum to discuss and drive the convergence of the internet of content, internet of services and internet of things.

- ServiceWave 2008 – Universidad Politécnica de Madrid, Madrid, Spain
- ServiceWave 2009 – KTH, Stockholm, Sweden, in cooperation with ICSOC'09

#### **NESSI project summit**

The NESSI Projects Summit was a 2 days event organized in Valencia in April 2010 during the Spanish presidency's full week on "European Innovation and RTD Transforming Sectors" for FP7.

This first edition was organised by NESSI in collaboration with INES – the Spanish National Initiative on Software and Services – as well as with all the NESSI Strategic (NSP) and Compliant Projects (NCP) and network of excellence S-Cube.

With over 250 people registered for the Plenary Session, the NESSI Projects Summit was a unique opportunity for community networking, with NSPs and NCPs presenting and demonstrating their results.



### At a Glance

**Project title**

NESSI Open Service Framework  
Reference Architecture

**Contact person**

Stefano de Panfilis  
Engineering Ingegneria Informatica S.p.A.  
stefano.depanfilis@eng.it

**Website**

www.nexof-ra.eu

**Total cost / EC contribution**

6,562,157 € / 4,000,000 €

**Start date /end date**

March 2008 / July 2010

### Scope

With the advent of Future Internet and, before that of Internet of Services, a new generation of ICT applications arise. The complexity and heterogeneity of such applications vary so much that the need of new architectural approaches became evident. In the context of Internet of Services, the main goal of NEXOF-RA was to provide a reference architecture for service-based software systems which facilitates the building of this new generation of application through the reuse of well-proven service-oriented concepts.

### Advances and Positioning in global context

NEXOF-RA has delivered a Reference Architecture for the NESSI Open Service Framework (NEXOF) which ranges from the computational and network infrastructure up to the interfaces with services consumers. This Reference Architecture has lev-

eraged on relevant existing achievements either research results and best practices.

The Reference Architecture follows the “NEXOF Independence Principle”: independence from business size, business domain and implementing technology. While the first two elements are normally addressed in RTD projects, the last point is of crucial importance, in particular for SMEs, as it is extremely costly to deal with many different technological approaches. The NEXOF Reference Architecture relies on the adoption of “architectural patterns” as it is more and more evident that a single reference architecture for Internet of Service is unfeasible. On the other hand the choice of a reference architecture “way of thinking” as adopted in the project

- Subsumes recurring and well-proven concepts and patterns facilitating the:
  - identification and clarification of issues and choices
  - concentration on small number of issues at a time
- Allows an overarching architecture strategy able to accommodate:
  - cross domain, business size and technology aspects
  - system level solutions for shared environment: frame choices to ensure integrated, consistent and coherent concepts

In the light of the Internet of Services business ecosystems, the above choices are needed to face challenges such as:

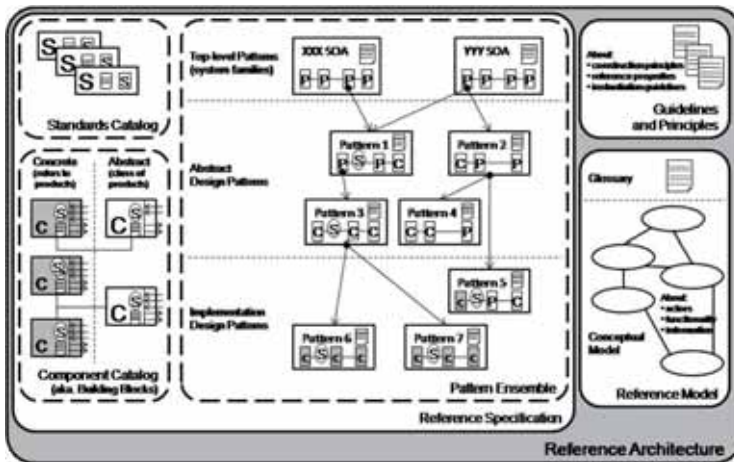
- Incomplete Scenarios and thus requirements to be addressed
- Heterogeneity of current solutions and thus lack of interoperability between different solutions
- Few integrated tools and methodologies aimed at supporting the architect’s problem
- Multiplicity of languages, models, and frameworks with insufficient coherence and consistency

Difficultly to perform an adequate conformance (compliance) test.



## Contribution to standardization and interoperability issues

All the patterns developed within the project were linked to specific standards in the field. This, even if is not a direct contribution to standards, makes the NEXOF-RA approach compliant to best practices and this way provides a relevant evaluation to standard solutions.



## Target users / sectors in business and society

The NEXOF Reference Architecture stakeholders are the service-based systems architects and Internet of Service researchers. In particular:

- **Large Industry** – Main benefits span the need of an overarching architecture strategy to:
  - Accommodate different technological solutions
  - Provide the basis for making incremental specific design decisions and product choices along the system building process
  - Accommodate system level solutions for a shared environment: frame choices to ensure integrated, consistent and coherent concepts
- **SME** – SMEs can benefit from NEXOF-RA as follows:
  - Efficiency through flexibility
  - Interoperability and quality
  - Mastering complex systems
  - Developing novel technologies
  - Fostering citizen related parameters.
- **Public Sector** – This is not homogeneous for what concerns size, domain and responsibilities. Here the approach of NEXOF-RA to build a reference architecture based on the “Independence Principle” could represent a success factor allowing specific instances of ICT service-based systems to be built in agreement with the peculiarities of each administration.

**The Scientific Community** – Due to its integration-nature, the NEXOF Reference Architecture, provides the foundation to align and improve future research activities in the Service Engineering

and SOA fields. It can be used to structure future research activities such as fundamental research and industry-oriented research with practical impact, but also it can be used to integrate existing research results into the NEXOF Reference Architecture. Thereby, NEXOF-RA might also reveal gaps in existing research results including both basic and advanced topics, this fostering an effective alignment and collaboration of further research activities.

## Examples of use cases

**Imagine** a young inventor who just came up with a startling new idea. He is alone, without any means to move from idea to implementation. But the virtual communities that are already emerging today can help: new services will provide the inventor with a personalised access to potential collaborations, across the world – collaborations with pre-defined qualities, partnerships that can take his invention from idea to product. The needs are such as linking different players to deliver an end result; linking complementary stakeholders for a common goal; taking the idea of “community” to a structured industrial approach.

Today, technologies could implement such collaboration – but this is much more – it is the exchange of critical, protected and sensitive data between different actors. It is about sharing information to a level that has yet to be achieved.

## Achievements

The main assets of NEXOF-RA are the concepts and process for specifying, assessing and comparing/selecting architectural patterns. Thus the main results achieved by the project are:

- The Open Architecture Specification processes
- The Architectural Framework, Principles and Model (the Model include an extensive Glossary)
- The Validation process based on fully functional Proof-of-Concepts.
- A methodology to write instantiation guidelines for architectural domains of the NEXOF Reference Architecture
- A sample set of Architectural Patterns proofing the viability of the approach. The patterns (produced with the contribution of NESSI Strategic Projects and the NEXOF Community) belongs to three different domains: Enterprise SOA, Cloud Computing and Internet of Services. The latter provides an architectural design of an internet-based platform to enable anyone to deliver, consume and prosume services at global scale.

All the project results are public domain under the creative common Attribution-NonCommercial-NoDerivs 3.0 Unported licence and are available in the project portal at [www.nexof-ra.eu](http://www.nexof-ra.eu)



OMP



### At a Glance

**Project title**

Open Media Platform

**Contact person**

Stefano Crespi Reghizzi  
Politecnico di Milano  
crespi@elet.polimi.it

**Website**

[www.dei.polimi.it/ricerca/sezioni/progetti/dettaglio.php?&id\\_sottoarea=46&id\\_sezione=3&id\\_progetto=109&action=ricerca&idlang=eng](http://www.dei.polimi.it/ricerca/sezioni/progetti/dettaglio.php?&id_sottoarea=46&id_sezione=3&id_progetto=109&action=ricerca&idlang=eng)

**Total cost / EC contribution**

4,692,936 € / 3,270,738 €

**Start date /end date**

January 2008 / December 2009

### Scope

OpenMediaPlatform introduces true dynamic composition for media services, thereby moving away from traditional static design approaches. OMP combines and extends component based software engineering principles with binary portability (using the Common Language Infrastructure standard), in such a way that efficient runtime engines can optimally exploit parallelism and available resources in the underlying device architecture. OMP adopts and extends open and standardized Application Programmers' Interfaces from the Khronos Group for media components to enable their flexible combination and their dynamic upgradeability. OMP also enables runtime adaptation to different Quality of Service levels, as required by the service, constrained by platform resources or dictated by network characteristics.

### Advances

The OMP project advances the state of the art in the following directions.

It developed and disseminated an innovative and extensible computing tool infrastructure consisting of static and dynamic composition tools and methodologies, providing component-based software engineering tools, a dynamic compiler for parallel architectures, as well as memory profiling tools and proofs of concept in parallel programming models and program control dependency analysis.

It defined and implemented a novel media infrastructure consisting of enhanced standard Application Programming Interface based on OpenMAX to support resource and context awareness, and it prototyped a standards-based run-time environment with associated resource management methodologies.

The research results are being used in current research projects, including the 2PARMA and TERAFLUX projects.

### Positioning in global context

The project has led to top-level publications, as well as providing the base for new projects (see above).

### Contribution to standardization and interoperability issues

A contribution to the OpenMAX standard has been proposed, introducing quality levels and their management.

### Target users / sectors in business and society

The target users of the tools and methodologies provided are system integrators and technology providers, as well as software and application developers.

### Overall Benefits for business and society

The availability of project products as free software provides reference implementations and

tools that are freely available, thus reducing significantly the effort required to enter the market or start up a research project. The cost of producing such software from scratch would be in excess of 4 person/years for the ILDJIT dynamic compiler. Component-based software engineering tools are also able to reduce the development time for parallel implementation of media tools.

### Examples of use cases

The integration of project results into the Bellagio OpenMAX, Scalable Video Coding reference implementations, and dynamic compilation open source tools provide immediate examples of use of project results in the real world.

### Achievements

The following research products are the key achievements of the OMP project.

*ILDJIT* is a dynamic compiler for the bytecode language defined by the ECMA-335 standard. It implements an adaptive compilation strategy to hide compilation latencies.

Current maturity status is *beta* release.

Project home page: <http://sourceforge.net/projects/ildjit>

1. S. Campanoni, G. Agosta, S. Crespi Reghizzi and A. Di Biagio. A highly flexible, parallel virtual machine: design and experience of ILDJIT. In *Software: Practice and Experience*, Volume 40 Issue 2, pages 177-207, January 2010.
2. M. Tartara, S. Campanoni, G. Agosta and S. Crespi Reghizzi. Parallelism and Retargetability in the ILDJIT Dynamic Compiler. In *Proceedings of the 2PARMA Workshop (ARCS2010 Workshop)*, Hannover, February 2010.
3. M. Tartara, S. Campanoni, G. Agosta and S. Crespi Reghizzi. Just-In-Time compilation on ARM processors. In *proceedings of the fourth workshop on the Implementation, Compilation,*

Optimization of Object-Oriented Languages, Programs and Systems (ICOOOLPS 2009), Genova, July 2009.

4. *Octopus* is a media engine for coordinating audio and video streams. It uses Gstreamer or OpenMAX IL components for individual operations.

Current maturity status is *beta* release.

Project home page:

<http://sandbox.movial.com/wiki/index.php/Octopus>

The *Bellagio OpenMAX* open source project has been extended in the project with a set of optional features in relation to the quality levels handling. Current maturity status is *beta* release.

Project home page: <http://omxil.sourceforge.net>

*DKU* is a programming model designed to support the portability of performance in parallel programs across a variety of platforms.

Current maturity status is *beta* release.

Project home page: <http://dku.sourceforge.net>

*Scalable Video Coding* is an encoding technique aiming at video stream quality adaptation to the trasmission conditions that has been included in several video coding standards. H.264/AVC-based SVC is of major practical interest. A component-based implementation has been developed by HHI within the project, and several studies have been conducted on its capabilities.

Current maturity status is *beta* release.

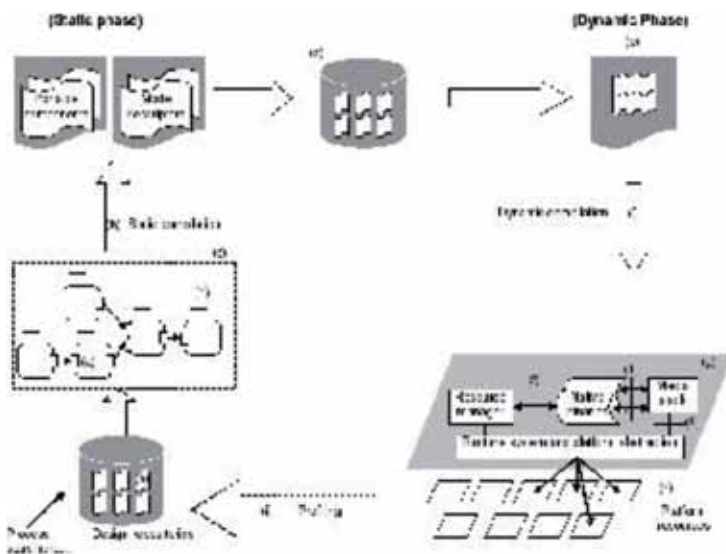
The contacts for this research product are Benno Stabernack and Jens Brandenburg of Fraunhofer HHI ([jens.brandenburg@hhi.fraunhofer.de](mailto:jens.brandenburg@hhi.fraunhofer.de)).

1. Heiko Hübert, Benno Stabernack, "Profiling-Based Hardware/Software Co-Exploration for the Design of Video Coding Architectures", *IEEE Transactions on Circuits and Systems for Videotechnology, Special Issue on Algorithm/Architecture Co Exploration of Visual Computing*, Volume 19, Issue 11, Nov. 2009, pages: 1680 – 1691.

2. Benno Stabernack, Jens Brandenburg, Heiko Hübert, Jan Möller, "An Experimental Mobile Terminal for Scalable Video Coding Applications using a H.264/AVC Decoder SOC", *Proc. 13th IEEE International Symposium on Consumer Electronics* May 25-28, 2009, Mielparque-Kyoto, Kyoto, Japan.

In the context of split compilation, two different tools have been produced, demonstrating *split register allocation* and *split vectorization*. The contacts for the split register allocation implementation in Java are Albert Cohen and Boubacar Diouf ([Albert.Cohen@inria.fr](mailto:Albert.Cohen@inria.fr) and [Boubacar.Diouf@inria.fr](mailto:Boubacar.Diouf@inria.fr)).

Current maturity status is *beta* release for both products. The contact for split vectorization in the ECMA-335 bytecode is Erven Rohou ([Erven.Rohou@inria.fr](mailto:Erven.Rohou@inria.fr)).





Open Pervasive Environments for migratory iNteractive services

# OPEN



## At a Glance

### Project title

Open Pervasive Environments for migratory iNteractive services

### Contact person

Fabio Paternò  
ISTI-CNR  
fabio.paterno@isti.cnr.it

### Website

www.ict-open.eu

### Total cost / EC contribution

4,494,329 € / 2,832,080 €

### Start date /end date

February 2008 / August 2010

## Scope

One important aspect of pervasive environments is to provide users with the ability to freely move about and continue the interaction with the applications in use through a variety of interactive devices (e.g. cell phones, PDAs, desktop computers, digital television sets, ...) and communication channels with various characteristics and performance (i.e. WiFi, Bluetooth, sensor networks, UMTS, ...). However, despite such a plethora of types of terminals available in the market, the multi-device experience is missing since people cannot continue to perform their tasks when they move about and change the interaction device. This is due to the lack of platforms supporting the migration of interactive applications in different usage scenarios. The OPEN project provides integrated solutions to such issues able to address three aspects: device change, state persistence and content adaptation. This is obtained through a middleware able to: adapt and preserve the state

of the software application parts dedicated to interacting with end users; support mechanisms for application logic reconfiguration; and define flexible mechanisms from the underlying network layers. The trends analyzed by OPEN indicate a future where people are always online at broadband speeds and with multimedia access, using a range of devices, and engaged in a number of social networks / games. OPEN expands this vision to the simultaneous use of many devices, or the sequential use of one or more, without disruption to the services being used.

## Advances

In order to increase possible services and application scenarios in several contexts (services for citizen, business, games, and so on), OPEN offers a novel infrastructure able to:

- Deliver seamless and transparent support to users in carrying out their tasks when changing available devices, even in multi-user interactive applications, while preserving the interaction state and adapting the user interface to the current computing device considered;
- Provide and coordinate dynamically changing / reconfiguring services in ubiquitous contexts;
- Offer personalized user interaction through an infrastructure providing the context information (e.g. available devices, connectivity, users).

## Positioning in global context

Currently there are some tools that support preserving the application state at server-side (e.g. MS Azure) or navigating on compatible websites offline and synchronize when going back online (e.g. Google Gears) or providing some useful support for distributed users' access across multiple devices (e.g. HTML5, WebSockets). However, none of them is able to provide full support for interactive applications able to dynamically migrate from one device to another in order to follow the users and allow them to continue their tasks across the various devices.



### Contribution to standardization

The results of the project are considered in the W3C Model-Based User Interface Group (<http://www.w3.org/2011/01/mbui-wg-charter>), that develops standards as a basis for creating a market for interoperable authoring tools for context-aware user interfaces for Web-based service front ends. The main contribution of OPEN to this W3C working group has been the application of model-based user interface languages (MARIA), which have been exploited to support adaptation and continuity across devices with varying interaction resources.

### Target users in business and society

Potential target users can be: Software and Application Developers; Systems Integrators, Service Designers, Service and Content Providers, End-users, Policy makers, Entertainment companies, Home automation, or Device manufacturers. OPEN could also introduce new roles in the value chain and new types of business models e.g. considering a 'migration feature' as a new service to be offered in multi-device environments (such as museums, offices, airports).

### Overall Benefits for business and society

The range of opportunities that migratory applications open up can be beneficial in radically different application domains: applications whose tasks require time to carry out (e.g. games, business applications) or applications that have some rigid deadline and thus need to be completed wherever the user is (e.g.: online auctions). Indeed, in current interactive applications there is limited support for state persistence and continuity across various types of devices. With the solutions provided by OPEN, service providers can offer experience continuity to their customers. The OPEN migration platform has been designed to minimize the effort required from developers to turn an application into a migratory application.

### Examples of use cases

There are many applications that can benefit from migratory interactive services. In OPEN we have especially taken into considerations business applications and games.

The Emergency prototype focuses on a scenario where governmental agencies, organizations and companies work together to provide public security in emergency situations (in our case: a flood event). The OPEN solution gathers and integrates all the information on one screen: thus, it enables experts to better analyze and plan response actions in the emergency situation. By migrating the necessary components or even whole applications to one target, experts have all the needed information overlaid on one map at their disposal.

The Social Game prototype offers several functionalities, including chatting, betting, watching IPTV, playing a multiplayer racing game; each of them taking a different area of the screen. It is inspired by a scenario in which the user can compete against real pilots while watching a live grand prix event, can interact with other members of a gaming community through a chat service, and simulate betting on a set of different parameters. The OPEN platform allows users to select some components and migrate them to a mobile device when the user has to move.

### Achievements

- **OPEN Reference Architecture:** derived from top-down analysis, starting from the application scenarios and requirements, and bottom-up analysis starting with the identification of relevant functions based on the experience of the partners.
- **Middleware:** The OPEN Migration Service Platform provides integrated support for migrating interactive applications and application logic reconfiguration based on contextual events.
- **Guidelines (for developers, policy makers, etc.)** supporting effective UI migration in multi-device environments.
- **Publications:** For a list of them please see the OPEN web site (<http://www.ict-open.eu/>).
- **Lessons learned:**
  - Dynamic migration of web applications without posing constraints on the authoring techniques is possible
  - Advantage of separating the reconfiguration of the application logic from the application logic itself
  - The type and behaviour of an application are important aspects that need to be considered when its migration is triggered.



# PERSIST



## At a Glance

### Project title

PERsonal Self-Improving SmarT space

### Contact person

Kevin Doolin  
TSSG  
kdoolin@tssg.org

### Website

www.ict-persist.eu

### Total cost / EC contribution

5,562,523 € / 3,628,006 €

### Start date /end date

April 2008 / September 2010

## Scope

Much research has been devoted to the development of fixed pervasive smart spaces and a major challenge now is how to integrate the ubiquitous mobile user. While the vision of many fixed smart space projects is to provide increasing levels of support for the user who inhabits the fixed space, when he/she steps outside of this island, all of this is lost. The result is that one will end up with islands of pervasiveness separated by voids in which support for pervasiveness is limited or completely absent. In these cases the user experience is all or nothing, with no graceful degradation from the former to the latter.

The challenge for PERSIST was to provide a way of bridging the gap between conventional fixed smart spaces and ubiquitous systems developed for the mobile user so that the user will be constantly covered by their own pervasive Personal Smart Space (PSS).

### Advances

The project developed the Personal Smart Spaces, which provides a minimum set of functionalities that can be extended and enhanced as users

encounter other PSSs during their everyday activities.

In order to address the objective and vision of the project, the consortium identified a number of key features/abilities including the following.

- (1) A PSS is based on a personal area network that can link together a variety of networked components, ranging from mobile or wearable devices to smart dust. Their ability to inter-operate with each other will permit Personal Smart Spaces to automatically adapt environments to satisfy user preferences, to resolve conflicts and to facilitate a migration from smart places to smart regions.
- (2) PSSs can share information and services in an ad-hoc fashion. The possibilities range from simple business card exchanges, to accessing the context sources of others, such as GPS information, to disaster management scenarios.
- (3) A PSS can personalise the configuration of fixed smart spaces. By treating fixed smart spaces as PSSs in exactly the same way as the PSSs for mobile users, the mobile PSSs can interact with the fixed PSSs and personalise them according to individual user preferences whenever a user enters the fixed smart space.
- (4) A PSS can link into pervasive infrastructures. A patient might have a PSS, which incorporates sensors that capture his vital signs and logs them. Upon entering the doctor's surgery this data can be uploaded to the surgery PSS and appended to his medical record, which can in turn be accessed by the doctor's personalised PSS graphical interface.
- (5) A PSS can generate new business models. Consider a not uncommon situation at a bus stop. Andy approaches the bus stop and would like to connect to the Internet to discover when the next bus is due but he has no connectivity. Meanwhile, Betty, another passenger waiting for the bus, is accessing the Internet via her UMTS connection. Betty's PSS automatically offers to act as a broker between Andy and

her Internet service provider. She becomes a micro-operator for Internet service provision.

### Positioning in global context

During the project lifetime no other solutions existed for integrating isolated smart spaces, or “islands of pervasiveness”, with a user’s personal smart space. While significant developments have been made in terms of providing proactive support for users within smart spaces, we believe that we are ahead of the state of the art in terms of providing a complete integrated solution, which brings the areas of pervasive computing and smart space technologies together in this manner. In addition we believe that we have surpassed the state of the art in many distinct areas such as context management, personalization of services and proactivity.

### Contribution to standardization and interoperability issues

The PERSIST project supported the use of open standards and frameworks, and where possible adopted certain standards and frameworks into the integrated prototype development, in particular the OSGi specification.

PERSIST had influence on the two standards; the “Composite Capabilities/Preference Profiles” by the W3C; and the “Human Factors (HF); User Profile Management” by ETSI. Both organisations responded and accepted our comments for review.

### Target users / sectors in business and society

The “users” of the end result of the project include: Software and application developers, Technology providers, Service providers, End-users and consumers.

### Overall Benefits for business and society

The potential social impact of the PERSIST outcome is considerable, should the developed system be deployed in real-world circumstances. It will enable users to interact with each other and with fixed smart spaces in a unique and supportive fashion. Users will be greatly relieved from the burden of manually entering large amounts of information (as the system offers strong learning functionality in support of proactive behaviour), in public or private environments where two or more users are located. PERSIST looks after the configuration of shared resources and services, thereby reducing the amount of user required input. PSSs enable people to form new social interactions and explore new ways of interacting and collaborating with each other. Since PSSs are built around real physical environments, they support the natural desire of human beings to meet and interact in the real world.

Services can be created within the PERSIST environment to address various environments such as home, business, automotive, public spaces (such as hotel, hospital, restaurant, stadium), and others. The overall user experience in these situations is improved thanks to the seamless adaptation of services and user interaction to the current user context. This again relieves the user from the painful tasks of device configuration or complex service searches.

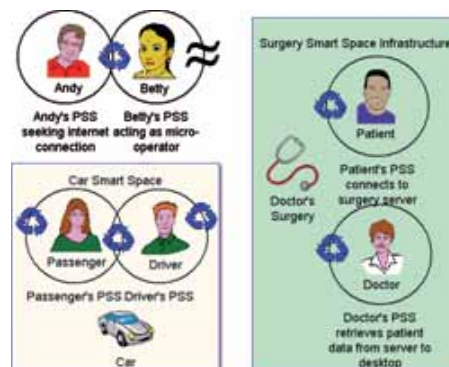
### Examples of use cases

PERSIST provides an open and extensible architecture to enable SMEs to create and use smart spaces. PSSs are inherently open and offer an entry point to the adoption of pervasive systems which is affordable to SMEs. PERSIST explored the potential for new business models based on PSSs whose scale is likely to prove attractive to an SME. PERSIST has provided open source software that will facilitate the implementation of new pervasive services. The project developed a number of demonstrators:

Enhanced navigation – demonstrates the capability of the PERSIST framework to learn about a user’s actions when navigating to / from an airport. Based on learned information from the users GPS device the system recommends locations for the user to visit. Virtual world – demonstrates the capability of the PERSIST framework to recognise a user as they enter a room (simulated by a virtual world). The room automatically reconfigures itself based on the particular users in the room, their context and their personal preferences.

Dementia assistance – demonstrates how PERSIST can facilitate dementia sufferers by notifying them when they meet someone they may have forgotten that they know, or by providing targeted personalised advertising as the user walks down a street.

Personal smart workspace – demonstrates the application of PERSIST in an enterprise context, showing personalised services for conferences and meetings, for example automatic presentation of a users powerpoint based on their authority and place in a meeting agenda.



### Achievements

The PERSIST consortium specified and implemented a reference architecture for a PSS as an open source project which can be found at <https://sourceforge.net/apps/trac/pSMARTSPACE/wiki>. The project provided an API for developers of third-party services that can exploit the PSS architecture and supporting documentation for this can be found at the sourceforge link above.

The PERSIST project officially participated and was presented in 19 international events and it published 38 papers (in journals, book or conferences). It actively collaborated with several EU research projects investigating similar fields and organised some events for dissemination and collaboration purposes.



## At a Glance

### Project title

Property-based testing

### Contact person

Prof. John Derrick  
University of Sheffield  
J.Derrick@dcs.shef.ac.uk

### Website

[www.protest-project.eu](http://www.protest-project.eu)

### Total cost / EC contribution

3,639,030 € / 2,709,821 €

### Start date /end date

May 2008 / November 2011

## Scope

Modern software systems are increasingly complex and interconnected, and a key challenge for the software industry is to deliver robust and reliable systems to customers. **Testing** is central to getting systems right, and in the **ProTest project** we deliver methods and tools for **property-based testing**, which shifts the emphasis from writing single tests to powerful properties that encapsulate aspects of system behaviour. These properties can then be tested in hundreds of randomly generated scenarios, rather than test-case-by-test-case.

The outcomes of the project allow developers to write more effective tests, more efficiently, and so to deliver higher quality software for a lower price, improving their competitiveness and thus benefitting the European software industry. Using property-based testing we have advanced testing practice in the European automotive industry, and also found numerous bugs in well-tested industrial concurrent code that had defeated other attempts to find them.

Industrial consortium members are system builders Ericsson and LambdaStream, the SME Quviq, which has commercialised the QuickCheck tool, and consultants and trainers Erlang Solutions, and academic partners are Chalmers University, the Polytechnic University of Madrid and the Universities of Sheffield and Kent.

## Advances

Property-driven testing and development is a powerful mechanism for gaining assurance of system reliability and functionality. Property-driven development can be used in a variety of programming languages and systems. The particular platform chosen for initial implementation of the project is Erlang/OTP (Open Telecom Platform), but a crucial aspect of our work was to provide for testing systems written in other languages such as C using the same testing framework.

Since current testing is based on sets of test cases embedded in test suites; we have built tools to aid software developers to extract properties from this data. Also, since many current specifications and models are informal: we have developed specialised property languages to ease the formalisation of specifications. Because all software systems are subject to change and evolution; we have built a refactoring tool to support the evolution of tests and properties in line with the evolution of the system itself.

Because not all properties can be tested in advance of systems being executed, and not all faults are found during testing, we have built a range of tools to support the post hoc examination of trace details for conformance to particular constraints.

Some of the most difficult bugs to track down are related to concurrency. While tracing and property based-testing find some of these, to complement them we have built a model checker that is able to systematically verify properties of concurrent systems.

## Positioning in global context

Property-based testing as developed in ProTest offers the promise of more effective testing, as prop-

erties are substantially more expressive of system requirements than sets of test cases. It promises a paradigm-shift in testing, allowing practitioners to isolate and fix bugs more effectively as well as earlier in the software development process.

ProTest consortium members lead the world in the development of properties for large-scale robust systems written not only in Erlang but also in C, and in sectors including automotive, telecoms, e-business, messaging and data storage.

### Contribution to standardization and interoperability issues

There has been an overall project goal to extend the impact of the results outside the Erlang community, both to other programming languages, and also via standards. As one example, we have developed an extension that allows us to test C code with our testing tool QuickCheck. This has been applied in a case study based on AUTOSAR, an open and standardized automotive software architecture developed in the automotive industry. This resulted in car companies promoting QuickCheck testing as the new standard for testing automotive software and Quviq have partnered with SP Technical Research Institute of Sweden to enable software certification in this area.



### Target users / sectors in business and society

Our results will help create new economic opportunities for software developers in a variety of sectors – including automotive, embedded, telecoms, messaging and e-business – using a variety of programming languages and platforms. Our property-based testing tool Quviq QuickCheck has had high impact within a number of companies including Ericsson, Gemini Mobile, T-Mobile, Motorola and Basho. Allowing reliable software to be delivered more rapidly and at lower cost will be of particular benefit to the European SME software developer community.

### Overall Benefits for business and society

The project provides a development process and tools that ensure dependable quality of service through directly verifying properties of the systems. Our results will allow software developers to bring to market more reliable products on a shorter timescale. Thus, their profitability will be

improved, and with that comes an increased ability to compete effectively in a global industry.

These innovations will help to nurture the European software service-provider sector, and help it to compete effectively on the global stage. Initial developments have been in the Erlang/OTP sector, but property-based testing in C is already supported, giving a wide impact to our results.


### Examples of use cases

- A large company was faced with the challenge of testing a parser for their domain specific language used to manipulate SIP headers. LambdaStream and a Japanese customer of Quviq had similar challenges in testing parsers for protocol messages. Quviq developed a library that can automatically generate tests from grammar specifications; which simplifies testing of parsers up to a large extent.
- A French customer of Quviq, was faced with the challenge of testing Ejabberd, the worlds leading open source chat software. The difficulty was to specify that the order of certain events was not really fixed. Specifying this in Message Sequence Charts (MSCs) results in non-deterministic MSCs. A special specification language based upon MSCs was developed to enable the specification of messages for this particular scenario.

### Achievements

Property-based testing has proved itself in practice: by equipping testers with the tools to track down faults that had proved elusive under other approaches, acknowledged bugs have been fixed. The key results of our project include:

- The QuickCheck tool supports property-based testing of systems written in Erlang and C.
- McErlang supports model checking for Erlang programs, and is integrated with the PULSE scheduler to provide integrated testing of concurrent systems.
- Wrangler provides refactorings that allow the evolution of properties and tests as systems evolve, as well as assisting in property extraction from existing test suites.
- The Trace Tool Builder incorporates high-level offline log analysis, first available in Onviso. TTB is a part of the standard Erlang distribution.
- The ProTest tools are widely available: QuickCheck as a product from Quviq; Wrangler, McErlang and Trace Tool Builder are open source products.
- Through proof-of-concept work, the project has established the principles of property extraction from existing test suites and UML documents. This and other work is documented in the papers and case study reports written during the project.
- Training materials for the tools and techniques, including videos, presentations and courses.
- Links to all tools, papers and other deliverables are at the project website.



# Q-IMPRESS



### At a Glance

**Project title**  
Quality impact prediction for evolving service oriented software

**Contact person**  
Dr Trifu Mircea  
Forschungszentrum Informatik  
An Der Universitaet Karlsruhe  
mtrifu@fzi.de

**Website**  
www.q-impress.eu

**Total cost / EC contribution**  
4,680,826 € / 3,333,446 €

**Start date /end date**  
January 2008 / December 2010

foresee the consequences of design decisions and system resource changes on performance, reliability and maintainability through quality impact analysis and simulation.

### Advances

The methods and tools developed in Q-ImPrESS allow for cost-effective development and evolution of service-oriented software. System engineers can try out different design scenarios and choose the best possible alternative with respect to its impact on the quality of the software, before even starting the actual implementation.

Software design always includes trade-offs between the different considered quality attributes, which refer to both external quality (such as performance and reliability) and internal quality (such as maintainability). Choosing a good design from the beginning helps to avoid later project delays because of quality properties such as performance getting out of control. A systematic consideration of quality aspects during the design can easily save millions of euros in large software projects.

### Scope

Today, complex software systems are an integral part of many everyday products and services from diverse application domains. All these systems need to evolve over long periods of time. Service-orientation provides the flexibility which is needed to adapt them to constantly changing user requirements. Unfortunately, the current standards and methods for service-oriented development do not support the analysis of software systems with respect to their quality of service. As a result, many major innovative application domains are not fully exploiting the advantages of service-orientation. The goal of the Q-ImPrESS project was to bring service-orientation to these critical application domains, where predictable and guaranteed end-to-end quality of service is particularly important. The main challenge here was to create a method for quality-driven software development and evolution. This method shall enable system engineers to

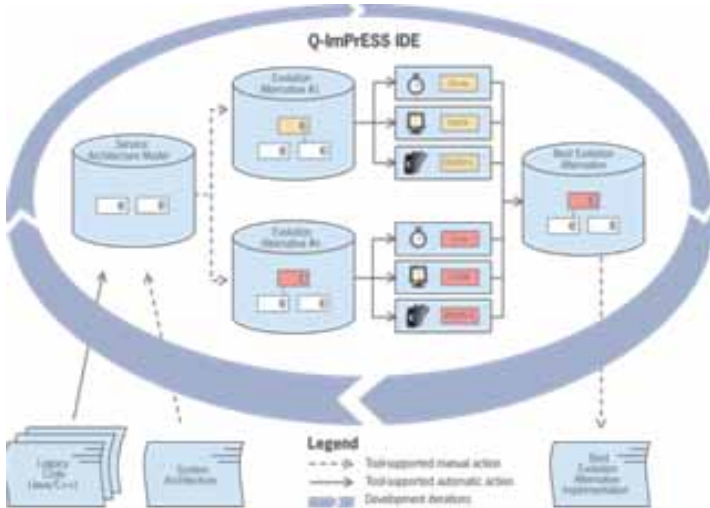
### Positioning in global context

The publicly available Q-ImPrESS Integrated Development Environment (IDE) which implements the Q-ImPrESS method, defines a new abstract design model of a software system augmented with information on service encapsulation and deployment, as shown in the picture. This design model is called the *service architecture model*. It is designed to be fully transformable, supporting the simulation of architectural changes. Such changes may include redistributing functionality between services, adding functionality to a service or service re-composition without changing the existing implementation.

Based on the service architecture model and applying model-driven techniques, individual prediction models for each considered quality attribute can be derived and used to predict the quality attributes for a given service architecture. The inte-



grated quality impact and trade-off analysis provides software engineers with a complete picture of the consequences of their architectural changes. It enables them to experiment with different service designs and architectures and select the best one for their system.



### Contribution to standardization and interoperability issues

The Q-ImPrESS IDE, which implements the Q-ImPrESS method, has been successfully released as OW2 open source project with an implementation based on the Eclipse platform. Furthermore, the project provides the Q-ImPrESS service architecture meta-model (SAMM) which defines a shared standardized, vendor-independent language for modelling service-oriented architectures with predictable quality of service.

### Target users / sectors in business and society

Q-ImPrESS targets critical application domains, where predictable and guaranteed end-to-end quality of service is particularly important. Q-ImPrESS supports several stakeholders including software and application developers, system integrators and architects, decision makers from the software engineering domain, service providers, as well as telecommunication and infrastructure providers, in providing high-grade and cutting edge service-oriented applications.

### Overall Benefits for business and society

The Q-ImPrESS methods help saving costs – it avoids the creation of service-oriented architectures with poor quality (e.g. little reliable applications, slowly responding services) through systematic consideration of quality aspects throughout the design. Trial-and-error cycles during service development are replaced by a well-planned engineering of service-oriented solutions. This increases project timeliness and can dramatically reduce project risks.

### Examples of use cases

The methods and tools of Q-ImPrESS have been successfully applied in the industrial automation and telecommunications domain. The developed methods and tools have become part of the standard tooling of the involved parties used in day-to-day work. Additional showcases for enterprise SOA systems demonstrate the broad range of application domains.

### Achievements

In summary, Q-ImPrESS represents a key element for successful quality-aware software engineering of critical infrastructures in service-oriented software systems.

The project provides the Q-ImPrESS service architecture meta-model (SAMM), a language to model service-oriented architectures with predictable quality of service aspects. Tool support is available to create SAMM instances from existing legacy applications written in C, C++ and Java. Fully automated transformations facilitate performance and reliability predictions for SAMM instances. Maintenance costs are automatically derived for change efforts implied by evolution scenarios. A comprehensive user guide describes the workflow of the Q-ImPrESS methods and tools, and gives corresponding examples. The documentation also includes guidelines on the application in industrial settings, best practices, and possible obstacles. The enterprise SOA showcase provides a publicly available demonstration of a typical SOA application in the supply chain management domain and demonstrates the practical applicability of the Q-ImPrESS method and its tools.

More than 70 peer-reviewed publications (including ten journal papers) illustrate the scientific impact of the project.

# RESERVOIR



## RESERVOIR

### At a Glance

**Project title**

Resources and Services Virtualization without Barriers

**Contact person**

Eliot Salant  
IBM Israel  
salant@il.ibm.com

**Website**

[www.reservoir-fp7.eu](http://www.reservoir-fp7.eu)

**Total cost / EC contribution**

17,169,906 € / 10,528,265 €

**Start date /end date**

February 2008 / March 2011

### Scope

Cloud Computing is rapidly gaining momentum as an alternative to traditional in-house ICT infrastructures – its “pay-as-you-go” model enables businesses to offer their services on the Internet without having to purchase and maintain their own computing infrastructure. However, typically today’s commercial clouds are provided by a relatively small number of large companies and use proprietary technologies and standards. RESERVOIR was a three-year European Union sponsored research project whose primary goal was to develop a reference architecture and technologies for the next of Cloud Computing, and demonstrated pioneering results, showing how virtual machines could be migrated live across physical hosts, cutting across subnetwork boundaries and without requiring shared storage. This technology was taken a step further as RESERVOIR then demonstrated how two or more Clouds potentially run by different organizations could collaborate to

overcome the capacity limitations inherent with a single cloud. Additional innovations added security to the RESERVOIR environment, and made important progress in a wide range of topics, such as statistically based overprovisioning of services while maintaining Service Level Agreements (SLAs), on-demand elastic scaling of virtual resources and monitoring and billing of consumed resources as just a few examples.

### Advances

RESERVOIR has made public its reference architecture – a detailed description of the functional components required to create a RESERVOIR stack, along with the interface specification for each component. In many cases, the developed components have also been released as Open Source, while the other proprietary portions of the stack typically may be licensed for use. Infrastructure providers wishing to build a version of a next generation Compute Cloud use RESERVOIR results to create a commercial Cloud. The RESERVOIR Framework, available from the “Downloads” section of the project web site ([www.reservoir-fp7.eu](http://www.reservoir-fp7.eu)), provides links to both the publically available architecture and design specifications for the project, as well as the code made publically available as Open Source.

### Positioning in global context

While commercial offerings in the world of Cloud Computing are rapidly improving, these solutions are typically proprietary, and lock in the user to a particular Cloud provider. In fact, in many commercial clouds, it is a closely kept secret what technologies are really being used behind the scenes. RESERVOIR believes that no one cloud provider, however large, can meet the spiraling demand for cloud services, and has championed the concept of open standards for federation between disparate clouds as a means for scale-out for clouds– which offers the potential for new European SME cloud ventures.

### Contribution to standardization and interoperability issues

Among RESERVOIR's major contributions to standardization and interoperability issues has been:

- An open standard for a unified VEE Management interface was defined. A new Working Group (OCCI) was created in OGF for its standardisation.
- The project created the first open-source reference implementation of the OCCI interface. This open-source technology supports hybrid cloud computing for federation.
- The RESERVOIR Service Management interfaces, proposed for discussion at the DMTF, under the TCloud label.
- RESERVOIR's results in the management of the virtualisation layer and integration between Service Computing Clouds and Virtual Networks are included in the ITU-T FG FN main recommendations report.

### Target users / sectors in business and society

The main beneficiaries of the RESERVOIR project results are service providers, infrastructure providers, system integrators and the research community. Taken as a whole, RESERVOIR provides a reference implementation for Cloud Computing infrastructure providers and significantly reduce the burden on system integrators in bring up a new Cloud. Service providers will also benefit from the ease of deployment of new services in a RESERVOIR cloud.

### Overall Benefits for business and society

One of RESERVOIR's goals has been to reduce the costs for deploying and hosting ITC services which would result in savings to both businesses and citizens. RESERVOIR's monitoring of actual resource consumption means that service providers will only pay for resources they use, while RESERVOIR's statistical overbooking of resources and elastic scaling means that infrastructure providers can squeeze higher utilization out of their existing equipment, instead of needing to overprovision compute servers.

### Examples of use cases

RESERVOIR picked four actual use cases from its industrial partners in the areas of:

- eBusiness, showing how a legacy ERP system could run over the cloud,
- Utility computing – showing how a typical grid computing application could be run in a cloud computing environment,
- eGovernment – which demonstrated how a typical three-tier application could be run in a Cloud environment, while restricting the migration of data to within geographical boundaries, and meeting peaks in demand on the system,
- Telco use case, demonstrating a Platform as a Service (PaaS) business model where a Telco operator could host owned or third party services on a RESERVOIR Computing Cloud.

### Achievements

RESERVOIR project results span a wide area of achievement including:

- A reference architecture for a next generation Compute Cloud.
- Publication of all component design and interface specifications.
- Over sixty scientific papers on RESERVOIR and Cloud Computing.
- Twelve training sessions across Europe on RESERVOIR technologies and Cloud Computing.
- One hundred and sixteen project presentations, keynote lectures, workshops and demonstrations.
- One hundred and twenty four articles and news clips on RESERVOIR in the press.
- Seven press releases.
- Spinoff technologies currently underdevelopment by a number of the industrial partners.
- Contribution of components to Open Source.
- The proposal of new standards for Cloud Computing.





# ROMULUS

## ROMULUS



### At a Glance

#### Project title

Domain Driven Design and Mashup Oriented Development based on Open Source Java Metaframework for Pragmatic, Reliable and Secure Web Development.

#### Contact person

Mr. Iglesias Carlos  
Informática Gesfor  
cif@germinus.com

#### Website

[www.ict-romulus.eu](http://www.ict-romulus.eu)

#### Total cost / EC contribution

3,111,670 € / 2,214,946 €

#### Start date /end date

January 2008 / December 2009

### Scope

Web Software Development is one of the most active areas in software development in Europe, and Java Enterprise Edition is the preferred option for one million developers. Web development is not yet a mature area, characterised by many technologies and frameworks. Developers are not able to master all these technologies, which affects seriously to their productivity.

Romulus project provides a complete suite based on a new open paradigm for developing web systems with Java technologies. The core of this suite is Roma Framework, focused on improving the Java Web Applications development processes by following a domain driven design approach. The Romulus suite includes utilities interoperating with the framework and covering all the Web development aspects.

### Advances

While many times, developers focused on the details of the technologies they use, Romulus focuses on domain engineering, delegating to a metaframework and making use of code generation techniques for generating most of the web applications from this domain model.

**Roma MetaFramework** – as the core part of the Romulus suite – provides the main functionalities of common Web. It is based on an extensible plug-in architecture that allows developers to include new modules using a metaframework syntax in an aspect oriented way. These modules implement the different aspects of Java Web application like persistence, view, services, users, workflow, reporting, etc. Besides, Romulus speeds up the web application development through the reuse of existing services and components. Applications are exposed as services, allowing developers to extend their applications through mashups with external services, enterprise systems, data sources or portal components such as:

- Data-level mashups. Applications created by Roma can be easily integrated within DERI pipes and Sig.ma semantic tools thanks to the module for RDF generation.
- Portal mashups. Integration with Liferay portal allows to access some portal functionalities from other platforms (Facebook and Netvibes) and provide the ability to build new applications combining specific parts of already existent applications.
- Web services mashups. Applications are integrated in MyCocktail -the Romulus Mashup Builder- through services. The developer can extend his application with mashups in a graphical way.
- Enterprise mashups. The developed components allow users to extend the generated applications within external Enterprise applications thanks to its OpenESB integration.

The Romulus suite is completed with **IDE4Romulus** -plug-in for create Roma applications within integrated development environments-,

**ATP4Romulus** -tool for automatic tests of the applications generated using the Roma framework- and **Wapiti** -black-box testing tool for assessing the security in the generated web applications-.

### Positioning in global context

The impact of Romulus in European software and services development is based on the fact that most software and services are developed in Java Enterprise Edition in Europe.

Romulus improves productivity in Java Web Development. Due to the impact of software in the economy, reducing the cost of software development and improving software quality are important economic objectives in Europe. In addition, Romulus is delivered as an Open Source project, which promotes cost reduction and improvement of innovation. Romulus is carrying out an active dissemination in open source communities (Roma Metaframework, Liferay and OWASP).

### Contribution to standardization and interoperability issues

- W3C Semantic Web Deployment WG: assists users in publishing data and vocabularies describing data.
- W3C RDB2RDF WG: a language for mapping relational data and schemas into RDF and OWL.



### Target users / sectors in business and society

Romulus helps Applications providers to build Java web applications in a more productive, reliable and secure way. The targeted end-users of Romulus are developers. The Romulus is a complete suite of development tools which cover the whole development cycle in order to accomplish these objectives.

### Overall Benefits for business and society

The advantages of Romulus technologies, comparing with conventional Java Web Development tools, are:

- Investment saving: Applications are finally portable across tools and frameworks because business

code and model does not contain any references to them.

- Huge improvement of productivity: Saving at least 50% of effort in comparison to conventional Model-View-Controller approaches and up to 80% for applications doing many database persistence operations.
- Average skills required is lower than for conventional frameworks. Developers have to write less code, making maintenance and testing easier.

### Examples of use cases

To show the applicability of Romulus' technologies in real contexts and demonstrators are provided:

- The European Projects Manager demo is a complete application that is able to manage the costs and the information of European FP7 research projects. Once all information of each project is completed, the project coordinator will be able to make a costs report to send it to the NEF application of the European Commission.
- Cornelius is an application that allows managing projects, company's employees and their hours-worked over projects activities. Its user policy is based on different profiles thanks to the security annotations included in Roma Framework.
- Scrooge is an application for personal financial management. This Rich-Client application has been created by using the Janus wizard. Scrooge allows users to categorize their own expenses and compare them with similar users.

### Achievements

Main software products developed within Romulus:

- Roma Meta Framework, core part of the Romulus suite, provides the main functionalities of common Web applications. [<http://www.romaframework.org/download.htm>]
- IDE4Romulus, NetBeans and Eclipse plugins which provides IDE integration facilities to Roma. [<http://www.ict-romulus.eu/web/ide4romulus>]
- MyCocktail, web application with a graphical user interface for building mashup, it allows to extend Roma applications with mashups. [<http://www.ict-romulus.eu/web/mycocktail/download>]
- ESB4Romulus, Roma Enterprise module allows the exposition of web services defined in Roma over ESB. [<http://www.ict-romulus.eu/web/esb4romulus>]
- Wapiti, an application for auditing the security of web applications. [<http://www.ict-romulus.eu/web/wapiti/download>]
- ATP4Romulus is a tool for the test case generation for the Romulus Framework. [<https://ict-romulus.svn.sourceforge.net/svnroot/ict-romulus/trunk/testing/ATP4Romulus-v03>]
- The complete list can be found at [<http://www.ict-romulus.eu/web/romulus/community/downloads>]

# S-CUBE



### At a Glance

**Project title**

The Software Services and Systems Network

**Contact person**

Dr. Andreas Metzger  
Universitaet Duisburg-Essen  
andreas.metzger@paluno.uni-due.de

**Website**

www.s-cube-network.eu

**Total cost / EC contribution**

11,046,169 € / 8,470,135 €

**Start date /end date**

March 2008 / February 2012

While many organisations across Europe are investigating some of these challenges, so far there has been little or no concerted effort to explore challenges that cut across different research disciplines.

S-Cube, the Network of Excellence on Software Services and Systems, has set out to address that situation by bringing together researchers from various disciplines such as to exploit synergy and learning effects across traditional research boundaries. For example, the software engineering discipline has acquired a huge body of knowledge on how to build and test traditional software systems. However, the distributed ownership of service-oriented systems implies that information required to build and test those systems might not be completely available during design-time, or – if available – can quickly become out-dated. Thus, many development activities need to be extended to run-time. This is where researchers working on self-adaptive computing infrastructures (such as used in grid and cloud computing) bring in their expertise to develop service-oriented systems that are more dynamic and autonomous. Researchers from service-oriented computing provide new languages and protocols to more dynamically and automatically compose and adapt software services. Finally, researchers from the business process management discipline contribute their understanding of how complex networks of organisations can work together to achieve joint goals via shared and distributed ownership of software systems.

### Scope

The rapid evolution of information and communication technology (ICT) means that the opportunities for new ways of computing and interacting are growing. One such opportunity is developing innovative systems through the composition of software services, available over widely distributed infrastructures. Those services have the power to provide utility to users in a much more dynamic and flexible way than is possible with traditional software technology. However, service-oriented systems and their corresponding software services require fundamental changes to the way software is developed, deployed and maintained. Software that constitutes a service-oriented system is no longer owned by on single organization but distributed and shared amongst many organizations. This distributed ownership opens up a whole range of research challenges, including the design, evolution, adaptation and quality assurance of service-oriented systems.

### Advances

To align and bring together the involved disciplines and co-existing, yet disparate, research themes, S-Cube pursues dedicated research and integration activities within the network. Those activities have led to results which are expected to have a sustainable impact on Europe’s research in service-oriented systems.



Figure 1: The S-CUBE Integrated Research Framework:



The most important integration results of S-Cube include:

- **S-Cube Integrated Research Framework (IRF):** The IRF guides S-Cube's research activities and ensures the integration of research results. The framework provides a clear separation of concerns and thus allows handling the complexity involved in aligning and integrating research activities of diverse disciplines. Figure 1 provides a high-level view on the framework.
- **S-Cube Knowledge Model (KM):** The S-Cube KM is a continuously updated on-line encyclopaedia and reference library for the Internet of Services, and is available via the S-Cube web portal at <http://www.s-cube-network.eu/km>. The network realised that the communities involved in S-Cube will not always be able to agree on a common terminology, so the KM is designed to provide interrelated and contextual definitions so researchers can translate between the vocabularies of the various research communities involved in this research.
- **S-Cube Virtual Campus (VC):** The S-Cube VC provides an online space facilitating the sharing of knowledge in the software services and systems community at <http://www.s-cube-network.eu/vc>. The VC provides a collection of learning modules, which provide teaching material based on the research themes of S-Cube. The material is intended to be used during lectures (such as S-Cube's joint master programme) but can also be employed for self-study.

The most important research results of S-Cube include:

- **S-Cube Life-Cycle Model:** The S-Cube life-cycle model defines the relevant activities for adaptation and evolution of service-oriented systems and integrates those into a coherent framework. Figure 2 presents the major activities and phases prescribed by the life-cycle model. In contrast to more traditional life-cycle models, the S-Cube life-cycle model considers the specifics of service-oriented systems, particularly concerning their dynamic adaptation during run-time.
- **S-Cube Quality Reference Model (QRM):** The S-Cube QRM provides a consolidated taxonomy of quality attributes for service-oriented systems, resulting from a thorough analysis of quality models used in software engineering, grid computing, business process management and service-oriented computing. Figure 3 presents the top-level structure of the QRM. The quality reference model can serve as a central access point to understanding the relevant quality attributes of service-oriented systems and how these are related to each other. The S-Cube QRM is accessible through the Knowledge Model via <http://www.s-cube-network.eu/km/qrm>.
- **Quality Prediction and Proactive Adaptation:** Ideally, during run-time the service-oriented system predicts the degradation of quality and thus imminent failures and thus can apply counter-measures to prevent the actual occurrence of failures. Such proactive adaptation addresses key drawbacks of more conventional reactive adaptation approaches. S-Cube's researchers have formed the Quality Prediction Working Group to work on novel approaches for proactive adaptation and to understand how those approaches could mutually benefit from each other. The results are available at <http://www.s-cube-network.eu/qp>.

### Positioning in global context

During the course of the network, S-Cube has placed an emphasis on raising the awareness about its research activities and results in the outside world. This was achieved through intense dissemination activities as reported below.

- A significant number of high-quality chapters, articles and papers were published in books, international journals, conference and workshop proceedings. S-Cube members gave presentations across the world to disseminate their research and that of the network.
- Two books were edited and published by S-Cube members:
  - **Service Research Challenges and Solutions for the Future Internet: Towards Mechanisms and Methods for Engineering, Managing, and Adapting Service-Based**

**Systems.** M. Papazoglou, K. Pohl, M. Parkin, and A. Metzger, (Eds.), Springer, 2010.

- **Service Engineering - European Research Results.** S. Dustdar, F. Li (Eds.), Springer, 2010.
- **International conferences, including Service-Wave, ICSOC and BPSC,** have been co-organized by S-Cube. ServiceWave is the leading EU conference on Future Internet and services systems, while ICSOC is the premier international conference on service-oriented computing.
- S-Cube members have co-organized many workshops, including thematic **workshops** to target industry and specialised communities on Services, Energy and Ecosystems, Engineering Service-Oriented Applications and Business Systems Management and Engineering.
- The organisation of the **annual Service and Software Architectures, Infrastructures and Engineering (SSAIE) Summer School** has brought together students from many countries to receive high-quality training from representatives from academia and industrial partners, such as IBM, Ericsson and BOC.
- S-Cube members lead several of the **SSAIE Collaboration Working Groups (CWGs)** and contributed to the **Future Internet Conference and Assembly (FIA)**.

Further, the S-Cube consortium was extended by taking on board **17 associate members** (including China), leading to fruitful collaborations, research results and creating awareness about S-Cube activities and outcomes.

Finally, S-Cube achieved awareness by positioning itself through the **ETP NESSI** and its **Industry Advisory Board**, which includes representatives from Engineering Ingegneria Informatica (Italy), HP (UK), IBM Research (Israel), SAP (Germany), Siemens (Germany), and Thales (France).

## Target users / sectors in business and society

As a network of excellence, the results of S-Cube are targeted at

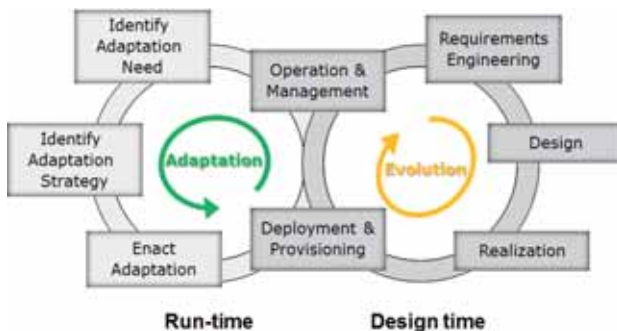
- **researchers and research communities** by providing an understanding of open, multi-disciplinary research to be addressed for the Future Internet and a baseline of novel techniques and methods for service engineering and adaptation (e.g., through the S-Cube integrated research framework);
- **students** benefitting from S-Cube results (e.g., through the joint International Masters in Service Engineering (IMSE), the SSAI summer school on service-oriented computing, the S-Cube Virtual Campus, and the S-Cube Knowledge Model);
- **developers and software engineers** by having access to research results allowing for the more efficient development and evolution of service-oriented systems that need to operate in a highly-dynamic environment;
- **service providers** (in their role of **service integrators**) by understanding how to build composed services that are adaptive and context-aware, leading to improved user satisfaction due to better addressing user needs and dynamically responding to varying user contexts;
- **industry (large enterprises and SMEs)** by providing a catalogue of techniques and methods that can be transferred to industry (such as the quality reference model, knowledge model with key terminology, and learning packages);
- **EU projects** through knowledge exchange, networking, and collaboration (such as the methodology for use case definition defined within S-Cube).

## Overall Benefits for business and society

The Future Internet will emerge through the convergence of software services, 'things' (network enabled devices and sensors), content, and communication networks. The Future Internet is expected to become the key infrastructure, essential to our future society and economy.

In the vision of the Future Internet, software services and service-oriented systems are expected to play a key role as an enabling technology and core building block. These services and systems will provide the correct level of abstraction from hardware and software entities, extending from business functions to data storage, processing and networking, devices and content. Ultimately, current service-oriented systems are expected to evolve into Future Internet Applications, operating on federated, open and trusted platforms, exploit-

Figure 2: The S-Cube lifecycle model:



ing the Internet of Content, Internet of Things and the Networks of the Future.

Future Internet Applications will exhibit an unprecedented level of dynamism, stemming from the need to adapt to broad, changing contexts and requirements. S-Cube has achieved a fundamental understanding of how to engineer adaptive service-oriented systems. S-Cube's research results thus provide an excellent foundation for future research on Future Internet Applications.


### Achievements

Overall, S-Cube has achieved the following important outcomes and results.

#### Joint Research Results

- *Joint Publications*, including over 200 articles and papers and book chapters, as well as 2 books edited by S-Cube members (see <http://www.s-cube-network.eu/results>).
- The *S-Cube Life-cycle Model* (see Figure 2).
- The *S-Cube Quality Reference Model* (see Figure 3).

Figure 3: S-Cube Quality Reference Model



#### Integration Results

- The *S-Cube Integrated Research Framework* (see above and Figure 1).
- The *S-Cube Knowledge Model* (see <http://www.s-cube-network.eu/km>).
- The *S-Cube Virtual Campus* (see <http://www.s-cube-network.eu/vc>).
- *Bonding of research staff and joint supervision*: During the course of the project, S-Cube researchers strengthened their joint research and integration activities by using S-Cube's dedicated mobility programme. In addition, a high number of partners were involved in jointly supervising PhD students leading to a number of jointly supervised PhD theses.
- *Collaborations*: In addition to setting up intensive collaborations with 17 associate members, S-Cube members have actively contributed to shaping ETP NESSI's Strategic Research Agenda, and they have strengthened direct links with industry, e.g., through S-Cube's Industrial Advisory Board.

#### Spreading of Excellence Results

- *Organisation of conferences and workshops*, such as ServiceWave, ICSOC, and BPSC (see <http://www.s-cube-network.eu/events>).
- The organisation of the annual SSAIE *summer school on service-oriented computing* (see <http://www.summersoc.eu/>).
- Setting up the Erasmus Mundus *International Master on Service Engineering* program, IMSE (see <http://www.erasmusmundus-imse.eu/>).



# SERVFACE



### At a Glance

**Project title**  
Service Annotations for User Interfaces

**Contact person**  
Dr. Goebel Steffen  
SAP Research  
steffen.goebel@sap.com

**Website**  
www.servface.eu

**Total cost / EC contribution**  
4,078,172 € / 2,635,098 €

**Start date /end date**  
February 2008 / October 2010

### Scope

Service-oriented Architectures are tailored and used for business applications across several enterprises. Services provide access to business data and allow managing business workflows. They flexibly integrate functionality across platforms and providers by using web service technologies. However, these services are not directly used by end users. User interfaces together with complex control logic must be developed as an additional layer on top of services. This development step is insufficiently supported by methodologies and tools. Instead, user interfaces are designed manually for service interfaces and business processes, which is both time consuming and expensive.

The ServFace projects developed a new methodology together with corresponding tools and models to more easily build user interfaces for service-based applications and this even for non-developers.

### Advances

The ServFace project developed a **model-driven service engineering methodology** for an integrated develop-

ment process for **service-based applications**. ServFace looked at this process from two different perspectives: First, the **annotation of services with corresponding user interface annotations** that were developed in the ServFace project, and, second, the **composition of annotated services** to build interactive service-based applications from annotated services.

The set of Service Annotations identified in the ServFace project are captured in the **ServFace Annotation Model**. It provides the necessary input for an automated user interface inference mechanism that generates high quality user interfaces for the interaction between human users and annotated web services. The **Annotation Tool** that has been developed in ServFace provides tool support for defining service annotations. Service annotations will be published in an annotation repository, from where they are globally available.

For the composition of annotated services to complex applications, two alternative modeling approaches have been investigated in ServFace. The first approach, **presentation-oriented service composition**, uses a novel mashup-oriented integration of annotated services. In presentation-oriented service composition, the application is modeled visually by composing the application's user interface from parts which are generated using the service annotations. The composition is supported by the **ServFace Builder**, an online tool that has been developed in the project. The ServFace Builder integrates an inference engine to generate user interfaces from annotated services. The presentation-oriented composition approach and the ServFace builder will be designed to support end-user development, enabling non-programmers to build service-based applications. The second approach is supported by the **MARIAE tool** and provides a novel solution able to exploit **task models** and **user interface models** for the design and development of interactive applications based on Web services for various types of platforms (desktop, smartphones, vocal, multimodal, ...). In this process the tool is able to automatically import service and annotation descriptions and support interactive association of basic system tasks with Web services operations. Then, a number of semi-automatic transformations are able to exploit the



information in such service and annotation descriptions to derive usable multi-device service front ends.

### Positioning in global context

The ServFace project developed the following three key innovations:

- **User interface annotations for web services** have been developed and supported by corresponding tool support.
- **Development of service-based applications for end users** has been made possible by using user interface annotations and focusing on usability for tool development.
- **Comprehensive evaluation by user studies** ensured that tools and methodologies are usable and accepted by targeted users.

All these innovations have been novel work in the research community, especially in this combination.


### Contribution to standardization and interoperability issues

Task models and abstract user interfaces have been targeted within the W3C standardization body. An incubator group has been active in the area of model-based user interfaces. The final report is available at <http://www.w3.org/2005/Incubator/model-based-ui/XGR-mbui-20100504/> A proposal to W3C for the creation of a group dedicated to produce standards is in preparation, in particular for task models and abstract user interfaces.

### Target users / sectors in business and society

There are three main users of the project results:

- **End-users** are enabled to develop simple service-based applications without involvement of the IT department or external consultants.
- **Service Providers** can simplify the use of their services and thereby increase their market reach by defining service annotations.
- **Software and Application Developers** can more easily and faster develop service-based applications for multiple platforms (e.g., desktop, smartphones, vocal).



### Overall Benefits for business and society

One of the goals of the ServFace project was to provide novel technologies that **improve the competitiveness of enterprises in Europe**. Several important results contribute to that:

**Improved methodology and tools for development of service-based applications:** The ServFace project extend SOA methodology to interactive business-to-business and business-to-consumer scenarios. The tools provided by ServFace will enable engineers to create service-based application faster and with less development efforts.

**User interface annotations for knowledge sharing and rapid UI development:** User interface annotations for web services are a key result of ServFace that capture additional information about web services and thus transfer knowledge from the web service provider to the web service consumer. UI annotations are used to speed up the development of user interfaces and complete applications. UI annotations might create a new ecosystem around web services that lead both to news services based applications and new web services due to the increased demand.

**New business scenarios through service ecosystem:** The ServFace project results will allow for provision of specialized services which can be combined with other services to value added complex services, especially by means of annotations. This enables especially small and specialized service providers (SMEs) with the possibility to provide services in a standardized, quick and easy way directly to customers. It will help to involve SMEs in the process of creating more complex and specialized business processes. Supporting all organizations developing or using software and services, particularly SMEs, to improve their competitiveness and adjust to the emerging global service will be facilitated by ServFace project results.

### Examples of use cases

The three selected scenarios investigated in the project cover different user domains and styles of applications. The “Sales Scenario” defines a use case of a sales order business process for a typical Customer Relationship Management (CRM) application. The “Home Scenario” provides insights into the domain of home appliances and the “Educational Scenario” presents a service-based portal to manage specific university processes. Each scenario description covers the following three aspects: (i) use case description in form of a user story, (ii) concrete and detailed process description of the whole scenario, and (iii) identification of the underlying composed and technical (web-) services.

### Achievements

- ServFace Methodology (final release): Deliverable D1.5
- User interface and service annotation models (final release): Deliverable D2.9
- Annotation Tool (final release): Deliverable D2.12
- ServFace Builder Tool (final release): Deliverable D2.12
- MARIA Tool (final release): Deliverable D2.12
- Publications: all publications are listed on the project website

# SERVICE WEB 3.0

## Service Web 3.0



### At a Glance

#### Project title

Service Web 3.0

#### Contact person

Dr. Elena Simperl

During project execution with Universitaet Innsbruck, now with Karlsruher Institut für Technologie  
elena.simperl@kit.edu

#### Website

[www.serviceweb30.eu](http://www.serviceweb30.eu)

#### Total cost / EC contribution

721,273 € / 482,350 €

#### Start date /end date

January 2008 / December 2009

### Scope

Computing is entering a new era that brings revolutionary changes at many levels, including infrastructure, hardware, middleware, and end-user software. Rapidly advancing ICT trends such as service orientation, smart mobile devices, cloud computing, sensor networks, user-generated content and the Semantic Web transform the Internet into a global platform of knowledge and services. A major challenge in this context is to set up an initiative that coordinates and supports ongoing research and development efforts in the area of Software and Services, specifically those that utilize semantic technologies, towards the realization of this vision of the future Internet of Services.

Service Web 3.0 addressed this challenge through a comprehensive, sustainable program encompass-

ing research coordination and alignment, dissemination and awareness raising, community building and networking activities. In particular, Service Web 3.0 played a pertinent role in the collaborative effort initiated by the European Commission to design and develop the scientific and technological building blocks for the next generation of the Internet, known as the Future Internet. The visionary leadership role of the Service Web 3.0 project formalized through prestigious positions held by senior members of the consortium – such as Future Internet caretaker, conference chair and organizer, standardization technical group and working group chair, as well as steering chair for several other initiatives – which in concert with three roadmapping workshops led to the publication of several roadmaps and collections of roadmaps focusing on future development and adoption of networked services and semantic technologies.

The successful establishment of the Future Internet Symposium and the organization of several academic and industrial conferences (e.g. the Extended Semantic Web Conference and the European Semantic Technology Conference) stand as highlights in reviewing Service Web 3.0's provision of professional forums and channels of dissemination.

Finally, the project's sustainable community networking infrastructure and support efforts were most visible in the established Future Internet Interest Group, the Semantic Technology Forum's Facebook group, the semantic technology YouTube channel broadcasting Service Web 3.0's Future Internet video, and the establishment of the STI Symposium, to be hosted yearly as an international community forum for showcasing achievements and establishing beneficial relationships with fellow researchers and developers in the field of semantic technologies.

## Advances and achievements

The Service Web 3.0 results include:

- (1) Development of a roadmap with key research challenges for Semantics and the Internet of Services;
- (2) Authoring and editing of reference publications, most notably books on the Future Internet, Semantic Web Services and semantic technologies;
- (3) Organization and support for 14 conferences, 14 workshops, 3 tutorials, 5 summer and winter schools, 10 keynote and invited talks, and 3 forums and special sessions on Future Internet related topics;
- (4) Development of promotion and information material comprising 6 scientific publications, and of “The Future Internet Video” (now with 130K views), in order to raise awareness, to encourage technology adoption, and to support community building;
- (5) Driving and maintaining the Future Internet Interest Group.



## Contribution to standardization and interoperability issues

The key standardisation activities led by Service Web 3.0 have been mostly within the field of semantic technologies, particularly semantic web services. These activities can be summarized as follows:

- The project’s identification of and participation in standardisation bodies, including W3C (e.g. a potential contribution towards SA-REST Member Submission and a potential working group to be led by Service Web 3.0 consortium members), OASIS, and the Conceptual Models of Services Working Group (CMS-WG) of STI International;

- The project’s identification of emerging standards and suggestions for how these standards can be improved and exploited, e.g. MicroWSMO, WSMO-Lite, the Semantic Execution Environment, and the Business Process Modelling Ontology;
- The project’s identification of and participation in evaluation programmes, which are important both to research and to the eventual mainstream adoption of semantic technologies, e.g. the Semantic Web Services Challenge and the Semantic Evaluation At Large Scale (SEALS) project.

## Overall Benefits for business and society

Service Web 3.0 targets various audiences, both academia and industry. The community around the Future Internet Assembly, in particular the organizations interested in the Internet of Services, has taken up the results of the project, including the research roadmaps on semantic and services technologies, the promotional Future Internet video, the conferences, workshops and symposia established and organized by Service Web 3.0, and the books authored or edited by project partners. Other communities (e.g., semantic technologies) have leveraged Service Web 3.0 to reach out to potential application areas, such as service-oriented computing.

The success of the Service Web 3.0 video, particularly the generally informative nature of the 6 minute mix of animation and expert interviews, helped broaden the targeted audience of the project’s dissemination activities. The video was shown at several conferences and events throughout 2009 and later. Academic and media organizations alike used the video for various educational and promotional purposes.

To support the diverse community of viewers, measures to ensure beneficial impact included:

- the creation of the Future Internet Interest Group, which keeps its members abreast of upcoming events and activities that are of relevance to the work of the Services Architecture working group of the Future Internet initiative;
- the Service Web 3.0: Services in Industry & Semantics in Services survey;
- the widespread distribution of the Service Web 3.0 public roadmap.



# SHAPE



## At a Glance

### Project title

Semantically-enabled Heterogeneous service Architecture and Platforms Engineering

### Contact person

Dr. Arne-Jørgen Berre  
Stiftelsen Sintef  
arne.j.berre@sintef.no

### Website

www.shape-project.eu

### Total cost / EC contribution

5,712,503 € / 3,907,685 €

### Start date /end date

December 2007 / June 2010

- Model-driven engineering (MDE) approaches. Standardisation of these results as extensions to SoaML.

### Positioning in global context

- SHAPE has contributed to a foundation for service modeling linking the business perspective and the technology perspective to address the challenges of service innovation and the service economy.

### Contribution to standardization and interoperability issues

- The SoaML (Service oriented architecture Modeling Language) OMG specification.

### Target users / sectors in business and society

- Software and Application Developers
- Systems Integrators
- Technology Providers
- Service Providers
- End-users

### Overall Benefits for business and society

The SHAPE project results support a new development paradigm with a higher degree of involvement of joint user and development communities through minimising the gap between business and system modelling, in particular by lifting the system specification models to a higher platform independent level.

## Scope

The objective of the SHAPE project was to support the development and realization of enterprise systems based on a Semantically-enabled Heterogeneous service architecture (SHA). SHA extends Service Oriented Architecture (SOA) with semantics and heterogeneous infrastructures with a unified service oriented approach. SHAPE has contributed to the SoaML (Service oriented architecture Modeling Language) standardization process in OMG and has provide an open source tool-supported methodology implementation based on SoaML.

## Advances

The main advances of SHAPE were to specify, develop and test a tool-supported methodology for designing and implementing:

- Flexible business models and parameterised services on SHA.



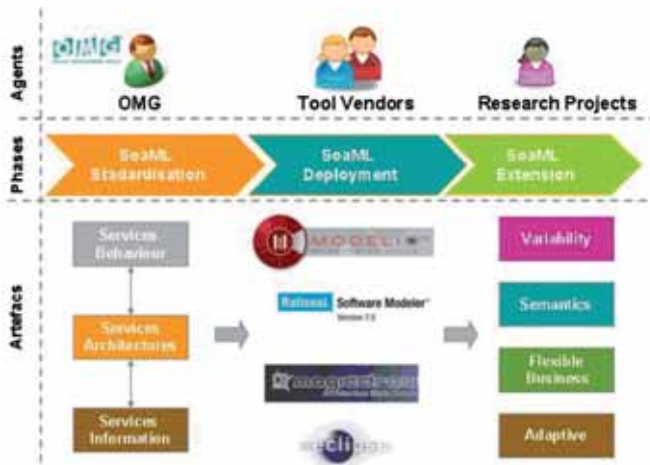
The following strengths and opportunities were identified in the two industrial pilot cases:

- Improvement of complete software development process.
- Improved practices for business and IT modelling.
- Enhanced interoperability of existing IT solutions.
- New system and information integration and interoperability approaches.
- Wrapping of legacy systems behind SOA.
- Improved communication between actors/roles in IT projects.
- Agent technology is a promising approach to scheduling algorithms.

- Increased support for flexible event and action management focused on scheduling algorithms using agent modelling techniques.

### Achievements

- SoaML: Contribution to the OMG specification SoaML (Service oriented architecture Modeling Language). Information about tool support and user experiences can be found at <http://www.soaml.org> and the adopted SoaML specification is available at <http://www.omg.org/spec/SoaML/>
- SHAPE Methodology. The methodology provides support and guidance for end-users in the design and development of service-oriented systems for particular application scenarios using the SHAPE modelling and engineering techniques. Available at: <http://www.shape-project.eu/download/>
- SHAPE Toolsuite: The toolsuite provides the necessary modelling tool and model transformation support to develop Semantically-enabled Heterogeneous service architecture. Available at: <http://www.shape-project.eu/download/>




### Examples of use cases

The SHAPE project derived user requirements and validated results in two industrial pilot cases: Statoil representing the Oil & Gas industry and Saarstahl representing the Steel industry.

The business challenges and related benefits for these two pilots were:

- Improved practices for business and IT modelling to ensure better alignment between business requirements and IT solutions.
- Enable easier system integration for information and services.



# SLA@SOI



### At a Glance

**Project title**  
Empowering the Service Economy with  
SLA-aware Infrastructures

**Contact person**  
Wolfgang Theilmann  
SAP Research  
wolfgang.theilmann@sap.com

**Website**  
www.sla-at-soi.eu

**Total cost / EC contribution**  
15,132,572 € / 9,632,923 €

**Start date /end date**  
June 2008 / July 2011

### Scope

IT-supported service provisioning, including cloud computing, has become of major relevance across practically all industries and domains. However, the service marketplace is still frustrating and cumbersome for both service providers and consumers. Consumers face high risks in service consumption as there are typically no formally agreed SLAs (service level agreements) specifying the quality and delivery conditions for a service. Providers face high costs and risks in the creation, negotiation, provisioning, operation and optimisation of SLA-governed services, and have almost no support in translating SLAs across different stakeholder perspectives and lifecycle phases.

SLA@SOI is delivering and showcasing an open SLA Management Framework that provides holistic support for service level objectives – enabling an open, dynamic, SLA-aware market for

European service providers. It allows SLAs to be managed autonomously throughout the service lifecycle, spanning the entire services stack from the business layer through to infrastructure.

### Advances

Although SLAs have been tackled by many research projects and commercial products, SLAs are typically treated as domain-specific add-ons. However, SLAs can only be consistently managed if all the different stakeholder perspectives and service layers are consistently interlinked and managed. Holistic SLAs are at the core of SLA@SOI. Some of the key advances provided include:

- A machine readable, domain-independent and highly extensible model in which arbitrary SLAs can be comprehensively described.
- A consistent interlinking of SLA models across service hierarchies and with other information models (such as business, software, and infrastructure models).
- A comprehensive open-source SLA management framework automating negotiation, provisioning, monitoring and other operations throughout the service lifecycle.
- A harmonized open-source SLA-aware virtualized infrastructure.

### Positioning in global context

SLA@SOI continues to be unique in its SLA-centric perspective, and is attracting growing interest from international analyst groups, governments and standards organizations. Industry players small and large are adopting results of SLA@SOI to develop new products.

### Contribution to standardization and interoperability issues

SLA@SOI has made significant contributions to standardization bodies.

- The Open Grid Forum's (OGF) **Open Cloud Computing Interface** (OCCI) standard has been



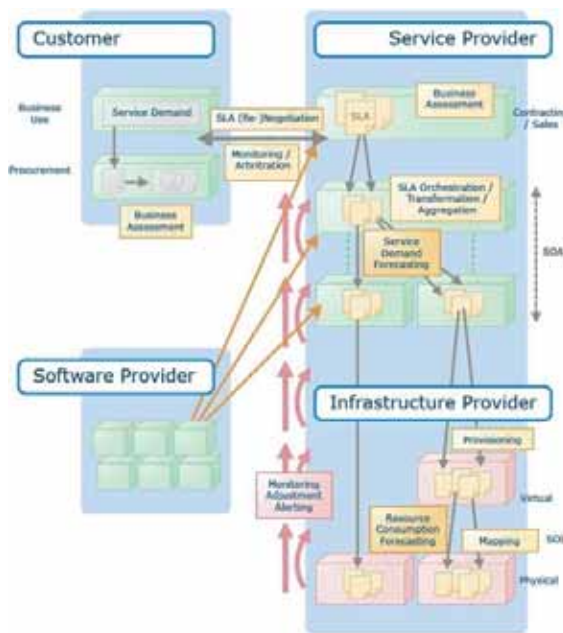
co-chaired and shaped by SLA@SOI to create a generic interface to cloud computing infrastructure through which SLA parameters can be communicated.

- SLA@SOI has also enhanced the **SLA negotiation protocol** as specified by OGF's **WS-Agreement** standard, introducing multi-round negotiation capabilities.
- Researchers from SLA@SOI have also contributed a dedicated **SLA module** to the **Unified Service Description Language (USDL)**, a standard proposed to the W3C.

### Target users / sectors in business and society

The SLA@SOI framework supports

- **Software and Application Developers** in engineering predictable services that can be operated under an SLA.
- **Service Providers and Aggregators** to offer, aggregate and operate services under well-defined, personalised, SLAs.
- **Sales Managers** to understand the conditions, constraints and tradeoffs of possible IT-supported services.
- **Service Consumers** by enabling automatic negotiation and personalised SLAs.



### Overall Benefits for business and society

The usage of the SLA@SOI management framework improves service offerings in the following dimensions:

- Service provisioning becomes more **dynamic** due to reduced preparation / setup times.
- Services become more **dependable**: the SLAs are machine readable and enforceable.
- Service management gets more automated and thus cost **efficient**.

- Service landscapes get more **flexible** due to simplified adjustment or reprovisioning.
- Service provisioning gets more **transparent** due to precise understanding of the tradeoff between cost drivers and service qualities.

### Examples of use cases

The SLA@SOI framework has been successfully adopted and evaluated in four complementary industrial use cases. **ERP Hosting** realizes dynamic provisioning of business application services and demonstrates major business benefits in both dynamic service provisioning and cost reduction. **Enterprise IT** has SLA-enabled internal infrastructure clouds, demonstrating increased agility and dependability. **Service Aggregator** demonstrates SLAs for aggregated telco services, delivering improved customer satisfaction and availability. **eGovernment** applies SLAs to health services, demonstrating increased compliance with governmental regulations. These four use cases prove the relevance of the framework to a broad range of both IT and human-based services.

### Achievements

The SLA@SOI project delivered a comprehensive set of technical results and supporting materials which allow both business and IT-related stakeholders to set up a proper SLA-driven management process for their service offerings. Key technical results include:

- A common machine-readable **SLA model** that can express arbitrary business rules, guarantees, actions and artefacts related to the delivery conditions of a service.
- A **reference architecture** for multi-layer, multi-domain SLA management.
- An **open-source reference implementation** of the SLA@SOI architecture (see <http://sourceforge.net/projects/sla-at-soi/>).

The project has also delivered

- Guidelines, tutorials and lessons learned in order for IT stakeholders to understand how to set up an SLA-driven service business.
- A complete open reference demonstrator that shows the SLA@SOI framework 'at work' and allows the replay, modification and extension of the demonstration.
- Business evaluations, guidelines and lessons learned that explain the business value that can be achieved through proper SLA management in different domains.

All technical and scientific results have been published in high ranked international journals and conference proceedings. A dedicated book summarizing the most relevant results has also been published.



## SMARTLM



### At a Glance

#### Project title

Grid-friendly software licensing for location independent application execution

#### Contact person

Josep Martrat  
Atos Origin  
josep.martrat@atosorigin.com

#### Website

www.smartlm.eu

#### Total cost / EC contribution

4,031,433 € / 2,696,044 €

#### Start date /end date

February 2008 / July 2010

license for the same admin domain of an organization. With Clouds notwithstanding, or when it comes to distributed environments and virtualized infrastructures, we run into trouble. Common software licensing terms are often too restrictive or expensive to run databases and applications on virtualized servers. Grids were an important inflection point in the transformation, but the lack of flexibility for running commercial software licenses in Clouds is still a bottleneck. Although licensing models have evolved with technology innovations, they do not fully satisfy the business issues faced by today's enterprises. The focus of successful licensing and support has to extend beyond cost and technology issues, to achieve software licensing based on business objectives. The goal and challenge is to balance customer needs and vendor business models. The licensing landscape is quite chaotic with many licensing and pricing models around and providers randomly introducing new ones: node locked license, flexible single user license, floating license, score based or token based model, perpetual license, per seat, per CPU, per concurrent user models, pay per use and subscription pricing, hybrid license models, custom contract based models and value oriented pricing.

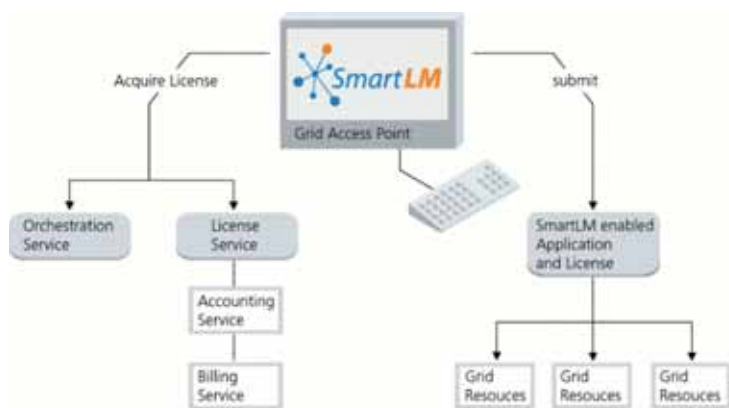
### Scope

Until recent years the use of applications requiring a software license for execution was quite limited in distributed environments but the wide adoption of Grids and Clouds has impacted on the licensing landscape and associated solutions. Due to the mandatory centralised control of license usage at application runtime, e.g. heartbeat control by the license server running at the home site of a user, traditional software licensing practices are not suitable especially when the distributed environment stretches across administrative domains. Therefore we see how the traditional software licensing models are under pressure as they do not satisfy the changing business needs of today's enterprises. Enterprise IT infrastructure is evolving towards hybrid models which harness in house and third party resources. However, licenses are usually bound to hardware rather than fungible resources and are provided on the basis of named users, hostnames, or as a site

### Advances

SmartLM project aims providing new generic licensing virtualization technology based on standards and integrating it in major Grid middleware/cloud solutions. The project has also identified new service-oriented business models for this approach. The main project outcome is the implementation of a novel approach for managing software licenses as web service resources in distributed service oriented environments. The SmartLM architecture is composed by the "license service" which is the core part of the SmartLM software "accounting and billing service" and the "orchestration service". It has been designed as foundation an innovative License Management Service where license resources are treated as web service resources that the end-user can use on-demand, book in advance, aggregate





SmartLM's distributed architecture

with other license resources and more, while having complete control over the costs thanks to the integrated Accounting and Billing Service.

### Positioning in global context

SmartLM is an advanced license enforcement technology with a strong selling proposition for software licensing. SmartLM covers current features of available licensing software and meets the requirements of commercial Grid and Cloud deployments. The most important selling points are:

- Improved security, authentication and authorization
- Embedding in Grid and Cloud environments
- Automated Accounting and Billing

SmartLM overcomes the current licensing mechanisms limitations and allows the development of a token-based mechanism, where the software license is defined and handled as a computing resource.

### Contribution to standardization and interoperability issues

SmartLM has closely cooperated with the “Grid Resource Allocation Agreement Protocol WG” (GRAAP-WG), where the specification of Service Level Agreement Negotiation protocol has been discussed and prepared

### Target users / sectors in business and society

Potential users may be:

- Software and Application Developers
- Systems Integrators
- Application Service Providers

### Overall Benefits for business and society

New licence models that offer access to license protected commercial software for a broader community, e.g. start-ups, entrepreneurs and SMEs who can access the necessary software without huge investments in hard- and soft- ware, will be pos-

sible through the fair and flexible management of software licenses. In addition, SmartLM enables the Software as a Service model, rising new business opportunities to Independent Software Vendors, Application Service Providers and end-user companies. These license models allow users to choose the best suited soft- and hardware for a certain task without a major investment. The flexibility of SmartLM allows the companies to run the software in external infrastructures in an easy way, reducing the need of expensive investments in hardware, adapting the budget of the projects and, as consequence, increasing the competitive position of European companies.

### Examples of use cases

Three license protected commercial applications have been adapted to be executed under control of the new licensing mechanisms and will become part of a show case to convince more code owners to adapt their applications. Concretely the three industrial codes of ISVs that have been adapted for using the SmartLM API are CFX (a computational fluid dynamics simulator) from ANSYS, PERMAS (a general purpose finite element analysis system used in the automotive and aerospace industry) from INTES and OPTIMUS (Process Integration and Design Optimization) from LMS. The current licensing mechanism underlying ANSYS CFX, PERMAS and OPTIMUS are based on licensing mechanisms that assume a local license manager installed in the same local network as the computational resource, where the software runs. These licensing mechanisms restrict the usage of the software product to resources installed in the company network which means less flexibility for customers. The use of the SmartLM with the commercial applications overcomes the current licensing mechanisms limitations and allows the development of a token-based mechanism, where the software license is defined and handled as a computing resource.

### Achievements

- Implemented and integrated a pre-commercial prototype called elasticLM ([www.elasticlm.eu](http://www.elasticlm.eu)).
- Developed new design patterns and technical implementations for a new generic licensing virtualization technology based on standards and integrating it in major Grid middleware solutions.
- Elaborated new business models for commercial applications as pay-per-use model or the specific scenarios for ASP.
- Signed an exploitation agreement that holds the baselines for product exploitation, including IPR management, agreed by all consortium partners.
- Presented project findings to Industrial events to raise awareness of the SmartLM licensing technology.
- Published in various relevant media such as The 451 MIS Euro report, GECON, Euro-Par, eChallenges.



# SOA4All



## At a Glance

### Project title

Service Oriented Architectures for All

### Contact person

Elies Prunés Soler  
Atos Origin R&I  
elies.prunes@atosresearch.eu

### Website

www.soa4all.eu

### Total cost / EC contribution

13,386,634 € / 9,474,415 €

### Start date /end date

March 2008 / April 2011

## Scope

Computer science has entered a new generation. The emerging generation starts by abstracting from software and sees all resources as services in a Service Oriented Architecture (SOA). In a world of services, SOAs have rapidly become the dominant computing paradigm. However, current SOA solutions are still restricted in their application context to being in-house solutions of companies. While service orientation is widely acknowledged for its potential to revolutionize the world of computing by abstracting from the underlying hardware and software layers, its success depends on resolving a number of fundamental challenges that SOA does not address today. SOA4All realizes a world where billions of parties are exposing and consuming services and data via advanced Web technology. To this end, SOA4All aims to bear SOA, Web, Semantic Web, Web of Data and Context adaptation as the core principles able to provide the power, flexibility and simplicity that is necessary for a wider uptake of service-oriented technologies in the Future Internet.

## Advances

The SOA4All project provides a comprehensive **global service delivery platform** that integrates complementary and revolutionary technical advances into a coherent and domain independent service delivery platform, i.e.:

- **SOA** as the dominant paradigm for application development which abstracts from software to the notion of service
- **Web principles** to scale SOA to a world wide Web communications infrastructure
- **Web 2.0** as a means to structure human-machine cooperation in an efficient and cost-effective manner
- **Semantic Web technologies** to enhance service descriptions and automate service discovery, mediation and composition
- **Web of Data** as the ongoing effort of exposing and interlinking on the Web data previously enclosed within silos
- **Context management** to meet local environmental constraints, organizational policies and personal preferences in composing, adapting and delivering services

## Positioning in global context

In a broad context, SOA4All significantly impacts the competitiveness of the European Software and Services industry in the Future Internet. The main success factors for the SOA4All integrated solution are the **open ecosystem** and the **service ecosystem**.

The **open ecosystem** is based on the idea of “a Web of billions of services” in an open environment, where an unlimited number of services can be offered, found, consumed and created, and where the interaction of different services and actors is enabled by SOA4All technology.

SOA4All offers added value in the open ecosystem approach, as it is a flexible and scalable solution for modern architectures that provides functional and nonfunctional properties that no other platform can provide nowadays.

The **service ecosystem** builds on the idea of federated service platforms and service parks, where ‘prosumers’ can find and consume services of different providers through a single platform.

The general value in the service ecosystem is that SOA4All can be used as a federated infrastructure well suited to this kind of service provision. The ‘prosumer’ gets benefit from the community and the ecosystem and can generate its own business value.

### Contribution to standardization and interoperability issues

SOA4All contributions to standardization and interoperability can be summarized as follows:

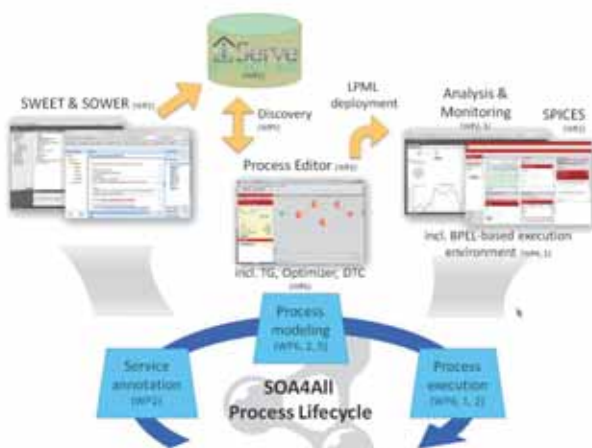
- W3C member submission of WSMO-Lite, a lightweight approach to the semantic annotation of Web service descriptions.
- Integration of lightweight annotation features into latest OASIS Semantic Execution Environments (SEE) Reference Ontology review draft.
- Contribution to Unified Service Description Language (USDL) W3C XG with focus on aligning lightweight service annotation languages and the reuse of existing ontologies with the USDL. The project also acted as consultant to THESEUS/TEXO in regards to USDL standardization, and worked with the SA-REST people to establish a W3C XG/WG.

SOA4All is endorsed by the NESSI Initiative (<http://www.nessi-europe.eu/>) and contributed significantly to the NESSI Open Framework, which is one of the main challenges of the European Platform on Software and Services.

### Target users / sectors in business and society

The **SOA4All Studio** is a fully-fledged Web-based framework that supports end-users throughout the entire life-cycle of services. It can be seen as a key pillar in achieving a Web that allows **everyone** to become a ‘prosumer’, i.e., easily create and access services and to turn the “Web of billions of services” into reality. The SOA4All Studio builds upon advanced Web 2.0 and Semantic Web technologies in order to hide complexity and provide an intuitive graphical interface that allows users to:

- Discover services in a distributed service space
- Annotate services with semantic information
- Compose services in a mashup-like way
- Use (consume) services with a single mouse click
- Monitor and analyze services



### Overall Benefits for business and society

Research and development about SOA related technologies carried on by the SOA4All project can contribute to enable sustainable growth in the field of service creation and delivery on the Web.

As an example, the SOA4All approach to a global service delivery platform can facilitate the creation of service infrastructures and increase the interoperability between large numbers of distributed, heterogeneous and highly dynamic functionalities on the Web.

By exploiting Web 2.0 and semantic technologies in tools offering increased automation and interoperability for service delivery, SOA4All can help new business ideas to be more easily realized and integrated in the upcoming “billions of services” open ecosystem.

### Examples of use cases


The SOA4All project developed three use cases with high market potential to demonstrate SOA4All technologies by means of business scenarios covering different target domains such as eCommerce, telecommunications and Public sector.

For instance, the **C2C Service eCommerce** use case demonstrates how SOA4All solutions can provide an easy way to use existing third-party services, enabling end users to build, publish and execute new eCommerce applications in order to market and sell their own products. E.g., thanks to SOA4All solutions a Webshop provider can combine offers from many different vendors on the fly and integrate them into a new process to be advertised on social platforms.

### Achievements

SOA4All main achievements can be summarized as follows:

- **SOA4All principles and core technologies:** the SOA4All approach to integrating complementary and revolutionary technical advances (the Web, context-aware technologies, Web 2.0 and Semantic Web) into a coherent and domain independent worldwide service delivery platform for the Future Internet of Services.
- **SOA4All reference architecture:** the foundation of the SOA4All framework supporting a world where a massive number of parties expose and consume services by realizing a coherent and domain independent platform.
- **SOA4All open source tools and platform services,** in particular the **SOA4All Studio:** a fully-fledged Web-based distributed system that provides extensive support for completing different tasks along the lifecycle of services.
- **Video tutorials and training material,** illustrating the SOA4All approach and demonstrating SOA4All tools.
- Several **scientific publications,** including journal articles, conference papers, white papers etc.
- **Contribution to standardization bodies,** in particular as for WSMO-Lite and USDL.
- **Collaboration with several research initiatives,** both at the technical level and as joint dissemination activities.



# STREAM

## At a Glance

### Project title

Scalable Autonomic Streaming  
Middleware for Real-Time Processing of  
Massive Data Flows

### Contact person

Ricardo Jimenez-Peris  
Universidad Politécnica de Madrid  
rjimenez@fi.upm.es

### Website

www.streamproject.eu

### Total cost / EC contribution

3,956,379 € / 2,609,597 €

### Start date /end date

February 2008 / April 2011

## Scope

A growing number of applications requires the ability to analyze massive amounts of streaming data in real-time. The state of the art before STREAM was unable to process massive data flows in real-time what has resulted in: 1) processing data offline in batch mode renouncing to the real-time processing requirement; or 2) processing data online but without doing all the required processing (e.g. not correlating data); or 3) discarding the application or not even conceived for considering it unfeasible. STREAM provides a highly scalable and elastic cloud computing platform enabling a whole breed of new services. STREAM enables to process in real-time massive amounts of data. STREAM can transform, correlate, aggregate and mine data at very high rates in the order of millions of messages/events per second.

## Advances

The main Stream contributions are:

- A cloud computing platform, StreamCloud, for providing real-time services over massive data flows, characterized by high expressivity (the same expressiveness as state of the art data streaming engines) and scalability (ability to scale to 100s of nodes).
- A data mining platform, StreamMine, able to perform online and real-time data analysis in combination with the data streaming platform.
- A communication infrastructure with low cost and storage infrastructure characterized by being able to store streaming data at network rates.

## Positioning in global context

Stream processing and complex event processing are emerging computing paradigms, which are gaining more and more attention, both from the academia and from the industry. STREAM outperforms currently available products by a novel parallelization approach. The STREAM platform is able to parallelize the processing of information flows in large clusters of 100s sites. Current approaches fail to scale for massive information flows. STREAM boosts the scalability of current approaches in 2 orders of magnitude. A high performance communication layer enables an efficient interaction among sites with low communication overhead. The storage layer is able to persist data streams at network rates enabling efficient and cost-effective fault-tolerance.

## Target users / sectors in business and society

Potential users of the STREAM platform are software vendors and business sectors who need to process massive amounts of data in real-time fashion, this includes the Internet of Things (e.g. the Smart grid), ICT security, monitoring of large grid and cloud infrastructures, telecommunication



infrastructure, processing of the output of large sensor networks (e.g. instrumented roads and highways), social networks, etc.

### Overall Benefits for business and society

A computing platform like STREAM is the technology enabler for a whole plethora of real-time applications which are currently impossible, such as credit card fraud detection. Credit card fraud is a growing concern of global proportions. Resourceful criminals are finding creative ways to capture private credit card holder account and identification information, and are using this information for fraudulent acquisitions of everything from personal care items to cars to home loans. Because of the universal reach of the Internet, criminals are easily able to perpetrate their crimes from anywhere in the world. The costs of credit card fraud reach nearly €1.8 billion annually in the EU-27. Internet fraud alone accounts for nearly 3% of Internet sales, or 30 times higher than credit card fraud rates in the “physical world.” While consumers are generally held harmless for credit card fraud, the payment industry and merchants absorb the losses from fraudulent purchases, and its participants continually search for ways to detect and prevent them. The STREAM project has demonstrated that real-time fraud detection is feasible using the STREAM platform.

### BUSINESS SCENARIOS

The STREAM platform has been validated in three business scenarios, namely

- Cellular Telephony Anti-Fraud Pilot – According to FCC (Federal Communications Commission), cell phone fraud is defined as the unauthorized use, tampering, or modification of a particular cell phone device or service. Malicious individuals combine complex software and hardware devices to illegally monitor data transmissions from the cell phones of authorized subscribers.

With this stolen data, they are able to clone the compromised phone with one that has been reprogrammed to retransmit the electronic serial number (ESN) and phone number of another authorized cell phone user. As the service provider cannot determine a difference, the fraud goes undetected. The Cellular Telephony Anti-Fraud pilot application detects in real time the cloning fraud and builds and updates customer profiles which comprise customer behavior and which, besides facilitating fraud detection, will give a better understanding of its customers to the Telefónica Company.

- Credit Card Payment Management Pilot – As already mentioned, credit card fraud is a growing concern of global proportions. The STREAM platform has proven to be a powerful tool for credit card detection. The pilot application uses all available data – account, payment and non-monetary – in the fraud decision making process.
- Quality of Service monitoring pilot – Dependable QoS monitoring has become a key requirement in a variety of business applications (virtually all applications which involve multiple parties). The QoS monitoring pilot application detects in real-time violations of the Service Level Agreements which have been agreed between the provider and its customers, and ii) provides unforgeable evidence of such violations.

### Achievements

- StreamCloud: largely scalable parallel-distributed data streaming engine (IPR protected by a filed international patent) featuring elasticity, fault-tolerance, visual development environment, parallel compiler, cloud deployer, real-time graphical performance monitor.
- StreamMine: scalable customizable layer for mining data streams.
- High performance storage tier and low-overhead communication tier.





European Commission

**FP7 Call 1 Achievements in Software & Services –  
Software & Service Architectures, Infrastructures and Engineering**

Luxembourg: Publications Office of the European Union

2010 – 71 pp. – 21 x 29.7 cm

ISBN 978-92-79-20614-6

doi:10.2759/49652

## **HOW TO OBTAIN EU PUBLICATIONS**

### **Free publications:**

- via EU Bookshop (<http://bookshop.europa.eu>);
- at the European Union's representations or delegations. You can obtain their contact details on the Internet (<http://ec.europa.eu>) or by sending a fax to +352 2929-42758.

### **Priced publications:**

- via EU Bookshop (<http://bookshop.europa.eu>).

### **Priced subscriptions (e.g. annual series of the *Official Journal of the European Union* and reports of cases before the Court of Justice of the European Union):**

- via one of the sales agents of the Publications Office of the European Union ([http://publications.europa.eu/others/agents/index\\_en.htm](http://publications.europa.eu/others/agents/index_en.htm)).

# FP7 Call 1 Achievements in Software & Services

**SOFTWARE & SERVICE ARCHITECTURES,  
INFRASTRUCTURES AND ENGINEERING**



Publications Office

ISBN 978-92-79-20614-6



9 789279 206146



European Commission  
Information Society and Media