SEQUOIA PROJECT

"Socio-Economic Impact Assessment for Research Projects"

Contract n° 258346

WP2

SEQUOIA Self-Assessment Methodology Development

Deliverable D2.1

Results from Projects Documentation Inventory

Project funded by the European Commission "Information Society and Media Directorate - General", Support Action

Contract Number: 258346

Project Acronym: SEQUOIA

Deliverable N°: 2.1

Due date: June 2010

Delivery Date (Resubmission): 26 September 2011

Authors: T6, LSE, UKS, ENG

Partners contributed: T6, LSE, EK, ENG

Made available to: Project partners and EC

Versioning		
Version	**Date**	**Name, organisation**
1	15/6/2010	Antonella Passani (T6)
2	21/07/2010	Mauro Navarra (EK), Shenja Van Der Graaf (LSE), Antonella Passani (T6)
3	15/09/2010	Marco Alessi (ENG)
4	20/10/2010	Mauro Navarra (EK), Shenja Van Der Graaf (LSE), Antonella Passani (T6)
5	30/11/2010	Antonella Passani (T6)
6	21/11/2011	Shenja Van Der Graaf (LSE)

Quality check: Louise Newton-Clare - LSE

Internal Reviewers: Francesco Bellini - EK

Table of contents

INTRODUCTION

The aim of this deliverable, as stated in the DoW, is to present the list of all the documents gathered from the projects under analysis (together with other relevant documents) and to provide a concise description of each of them. The documents mapped in this deliverable constitute the basis for the development of the questionnaire and of the focus groups' specific topics.

As we will see this deliverable goes a little further; in fact, the consortium soon understood the need to cluster the projects under analysis in order to develop a deeper understanding from the point of view of their expected impacts. When starting looking at the project information, in fact, it emerged that – besides belonging to the same research field – the projects were very diverse in terms of specific goals, technology developed, and potential beneficiaries. Consequently, in this deliverable, we will list and present the documents gathered and used for developing the two questionnaires developed so far and the clustering exercise undertaken.

Before describing the work done so far, it can be of help to position the present deliverable in the general map of SEQUOIA socio-economic impact assessment development process (Figure 1). As evident from the figure, this deliverable represents the very first step of the methodology definition.

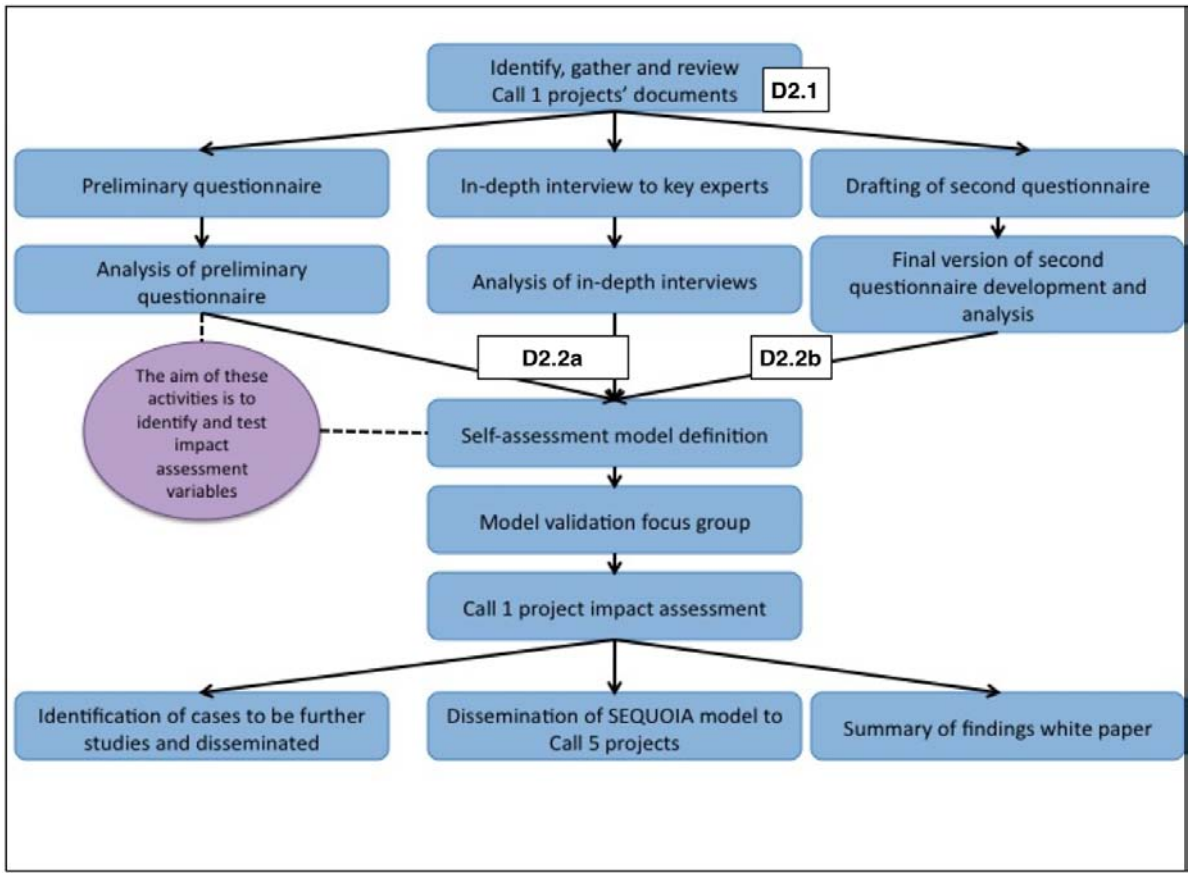

Figure 1 - Process model of SEQUOIA's impact assessment methodology development

The reason for the delay in the deliverable submission is mainly related to the fact that, while searching, reading, and selecting the projects' related documents, the researchers realised that it would be useful to use this deliverable as a living document to be updated, for instance as a result of a better understanding of the projects through the preliminary questionnaire.

The feedback gathered through the questionnaire, in fact, made it possible to add more material and/or better select the material already analysed.

The deliverable is organised in three main chapters and a conclusion. In Chapter 1 an overview of the projects under analysis and their available documents is presented. Chapter 2 presents the clustering exercise performed. Chapter 3 presents a concise description of the documents analysed so far.

CHAPTER 1. CALL 1 PROJECTS: AN OVERVIEW

The SEQUOIA project took in consideration 25 projects (24 IPs/STREPs and one Network of Excellence), which represent the total number of projects funded under Call 1 with the exception of the support actions. Following the organisation of the call in sub-areas we can see the projects grouped in the following 4 areas: service front-end (5 projects), Service architectures (5 projects), Virtualised infrastructures (5 projects), service/software engineering (8 projects), plus a network of excellence (Figure 2).

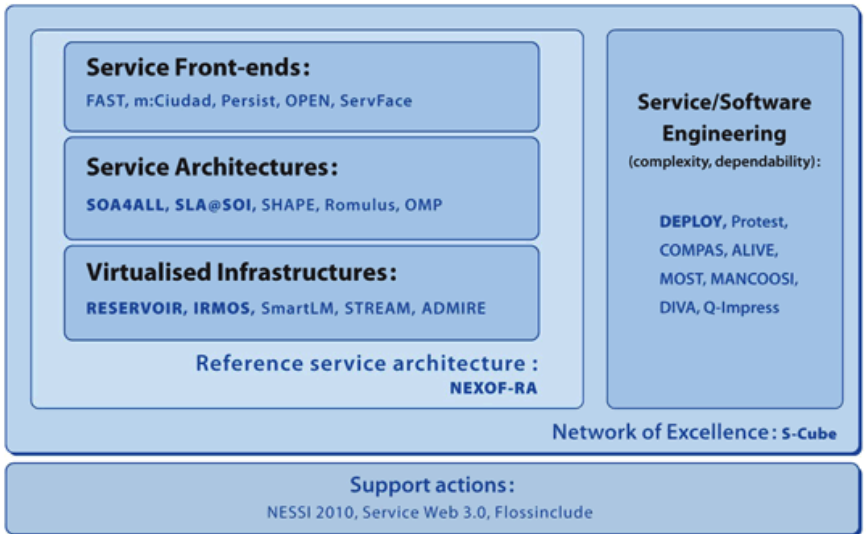

Figure 2 - Call 1 Projects

The list of projects under analysis follows, in alphabetical order:

ADMIRE
ALIVE
Compas
DEPLOY
DIVA
FAST
IRMOS
m:Ciudad
MANCOOSI
MOST
NEXOF-RA
OMP
OPEN
Persists
ProTest
Q-impress
RESERVOIR
Romulus
S-Cube
ServFace
SHAPE
SLA@SOI
SmartLM
SOA4ALL
STREAM

Considering all the projects, there are approximately 450 deliverables to be examined and a consistent amount of papers. The need to use the term “approximately” is due to the difficulty, in some cases, to access the documents that are not always available on projects’ websites and/or are in the same repository as publication and dissemination materials, making the identification more difficult. With reference to publications and papers submitted to conferences, we have around 350 documents. In addition, for each project the website and the factsheet have been taken in consideration.

When looking at the deliverables and other relevant documents we thought it important to consider and try to capture the number of case studies and prototypes developed by the projects. In other words, we look at concrete – even if not completed or fully tested – solutions able to be presented to the users identified. This is because to have concrete realisations and case studies would make the development of variables for the impact assessment easier. From this point of view, we can say that 41 cases where highlighted in project documents and websites. This does not mean that other pilot actions, prototypes and case studies have not been done, but that 41 cases represent a good entry point for learning more about the related projects and the possible ways for mapping Call 1 and Call 5 projects' impacts.

Table 1 includes a short description of each project. We used this preliminary description internally in order to build a common understanding about the projects under analysis. One of the problems the consortium faced, indeed, is represented by the need to quickly develop a common understanding of the projects' subjects especially from a technological point of view. The social science researchers entered the SEQUOIA project with significant knowledge and experience in methodologies of socio-economic impact analysis and in working in ICT-related research projects, but needed help in going deeper in the technological aspects of Call 1 projects. In order to develop a self-assessment methodology, in fact, the variables need to be adjusted to the research field in a very specific way, or at least this should be the expected attitude. The partners of the project with a deeper knowledge of the research area under analysis supported the socio-economic researchers in this sense. During the kick-off meeting, moreover, a brainstorming session was dedicated to the development of a common understanding of the project sub-areas and their main technological aspects.

Project	Brief description
ADMIRE	Data mining tools for extracting information and knowledge from large databases. Example scenarios from meteorological and CRM data. **Output**: Design and development of DMI (*Data Mining and Integration*) platform. The core service infrastructure of the platform is built upon the *Universal Service Management Technologies (USMT).* The USMT, provided by Fujitsu Labs of Europe, will provide a set of common service management capabilities that are required for monitoring and management of services hosted by the DMI platform. Data Mining and Integration processes are implemented as a set of individual processing steps, implemented using GRID services, which are combined to form the end-to-end process. This infrastructure integrates the OGSA (Open Grid Services Architecture)-DAI platform, an enhanced monitoring framework for DMI-oriented processes, and the OGSA-DQP platform, a semantic registry and semantic provisioning services for providing advanced service and resource publishing and discovery mechanisms. The project also provides a set of tools based on the Eclipse platform which advance the power and ease of use of data mining and data integration based on the ADMIRE model and platform.
ALIVE	Formalisation of multiple abstraction layers (service, coordination, organisation) through a model-driven approach for the purpose of coordinating the composition and execution of web services in distributed environments. **Output**: The project: • Develops an advanced framework for application development, deployment and management in service environments,

	• Develops new engineering techniques and tools to help designers to structure service-based systems in such a way that they are easier to understand for non-technical people • Develops a methodology for dynamic, "live" service design and maintenance, and • Layers the framework directly on emerging architectures and toolkits for service-oriented and web services systems.
COMPAS	Model-driven compliance rules for software; policies for all aspects and stages of the software service life-cycle in the SOA context; extension of BPEL to incorporate enforcement of compliance rules in business processes. **Output**: The COMPAS project will design and implement novel models, languages, and an architectural framework including required software components and services to ensure dynamic and on-going compliance of software services to business regulations and design rules. This is achieved using the model-driven software development (MDSD) approach to enable organisations developing custom business compliance solutions more quickly, more economically, and with a lower level of required programming skills. COMPAS will enhance business process languages, such as the Business Process Execution Language (BPEL), with enforceable compliance concepts and policies. Furthermore, COMPAS will develop specification languages and models for expressing typical compliance concerns.
DEPLOY	Formalisation of engineering methods in 5 different sectors (transportation, automotive, space, telecommunications, business information) **Output**: The overall aim of the DEPLOY Project is to make major advances in engineering methods for dependable systems through the deployment of formal engineering methods. Formal engineering methods enable greater mastery of complexity than found in traditional software engineering processes. It is the central role played by mechanically-analysed formal models throughout the system development flow that enables mastery of complexity. As well as leading to big improvements in system dependability, greater mastery of complexity also leads to greater productivity by reducing the expensive test-debug-rework cycle and by facilitating increased reuse of software.
DIVA	This project addresses the problem of reconfiguring the run-time behaviour of a software application in response to changes in the context. It does this through a proposed integration of aspect-oriented programming and model-driven engineering. The former is a systematic way to modularise functionality, whereas the latter provides a consistent formalism for describing systems at multiple levels of abstraction. The integration of these two methodologies should in principle make it easier to substitute different modules dynamically at run-time, based on the evolving needs of the context. **Output**: The goal of DiVA is to provide a new tool-supported methodology with an integrated framework for managing dynamic variability in adaptive systems. It will test its results implementing case studies from two different domains: crisis management and Customer Relationship Management (CRM). The project will provide DiVA studio in order to coordinate and produce an integrated framework for adaptive system engineering. DiVA studio is based on FaMa Tool Suite which is the first tool integrating different solvers for the automated analysis of feature models.
FAST	Visual programming environment to facilitate the development of graphical gadgets and composite GUIs for the visualisation of business processes and the top-down specification of service and application composition. Will support

	mashups of services taken from multiple remote portals and end-points. User-centred orchestration of loosely coupled services will make it easier to modify business workflows bypassing hard-coded syntactical approaches. OSS and community-oriented outputs. **Output**: Create a new visual programming environment that will facilitate the development of complex front-end gadgets, involving execution of relatively complex business processes that rely on back-end Semantic Web Services. This environment is composed by the Gadget Visual Storyboard, a prototype that supports the key element of the FAST Project, the creation of complex gadgets using lower-level building blocks. The prototype will allow users to compose screens in a visual manner in order to create screenflows that can be deployed in current mashup platforms. It is composed by a client and a server side with a semantic catalogue for screen-flow gadgets and back-end services.
IRMOS	Real-time support for service delivery, with implications at different levels of the stack: QoS at the networking level (minimum bandwidth guarantees for streaming), automatic SLA negotiation at the application level. Support for complex inter-organisational transactions (observation: they are probably relying on a centralised transaction coordination server, i.e. WS-Coordination standard). Interestingly, they provide a list of impacts that we can use directly as a starting point in our analysis. Area of relevance is multimedia applications. **Output**: The main objective of IRMOS is to enable real-time interaction and collaboration between people, using distributed multimedia applications running on top of a cloud infrastructure, where processing, storage and networking are delivered with guaranteed levels of service. IRMOS addresses the following cloud service models: • Infrastructure as a Service (IaaS): provided by the IRMOS Intelligent Service Oriented Network Infrastructure (ISONI), including processing, storage and networking resources. • Platform as a Service (PaaS): provided by the IRMOS Service Engineering and Service Management tools. • Software as a Service (SaaS): provided by a set of tools and a specific methodology for the creation and adaptation of applications for virtualized real-time service infrastructures • IRMOS provides a set of tools along with a specific methodology on how to engineer applications for cloud-based platforms. • IRMOS provides a Service Modelling Environment and tools to model service-oriented applications and predict their behaviour when executed on clouds. • IRMOS provides Performance Estimation Services that allow end-users to deal only with high-level application parameters, since these services undertake the task of defining the low-level resource attributes and workload characteristics of the applications. • IRMOS provides Service Management Tools that support resources negotiation, reservation, application execution and monitoring, addressing the real-time related performance requirements in a transparent way to the user while conforming to the Service Level Agreements. • IRMOS incorporates ISONI, an Intelligent Service Oriented Network Infrastructure that provides Virtual Service Networks by virtualising computing, networking and storage resources. • IRMOS enables real-time attributes at all levels of the environment (from

	the application to the platform, down to the infrastructure level).
MANCOOSI	Management of software upgrades in Open Source environments. The focus is the upgrade of individual operating systems. Methodology relies on formal models of the OS and of the upgrade process seen as a set of transactions. Thus, transaction roll-back support is envisaged, in case a step in the upgrade process fails and the previous state of the OS needs to be recovered. **Output**: MANCOOSI will propose virtuous cycles of collaboration among users, distribution editors, and researchers for free/open source software upgrade issues. Mancoosi will develop sophisticated optimisation algorithms to find efficient upgrade paths and high-level request languages which will make software upgrading a simpler process for any user, not only for experienced computer wizards. Mancoosi is also building a transactional layer into end-user package management tools, which will allow to bring your system back to a previous state ("rollback") without further problems, working at the level of individual components, and not on file-system checkpoints.
MCIUDAD	Service development and execution for mobile environments. Support for large numbers of relatively small/simple services is the scenario this project is working towards. Sensors (traffic), blogs, shopping, social networking, games, personal data are examples of the kinds of services addressed. **Output**: A new flexible service description language for mobile user-generated micro services (U+ services). A new creation and editing method to create new U+ services from the mobile device, encouraging user-generated services. A new point of view for knowledge provision and service creation based on the prosumer (producer + consumer) concept. A new embedded execution platform for mobile devices, to execute U+ services. The services are executed (both provided and consumed) in the mobile terminal. An optimised searching environment to locate and access relevant services providing relevant knowledge.
MOST	Ontology-based reasoning combined with model-driven engineering to produce ontology-driven software development. Unification of modelling languages used. Ontologies (semantic networks of terms, or structured and domain-specific vocabularies) and models (formal domain descriptions as the starting point in software development, i.e. class structure etc) are to be integrated. This probably means linking relevant ontologies to specific models, and vice versa. Ontology-based inferences are therefore going to influence the MDA transformations in going from abstract meta-models to specific instances. The implications in terms of functional code generation and data structure derivation are not clear. The project plans to formalise also the process by which the above approach is used; thus this is very much a meta-level perspective. The meta-level is formed by “guidance ontologies”. Sounds very ambitious but it should be interesting to see how far they got. **Output**: One major objective of the MOST approach is to leverage MDSD and other software development processes with ontology technology. This goal will be reached by implementing integration technology for ontologies into MDSD, resulting in ontology-driven software development (ODSD). This integration technology concerns all involved artefacts (ontology and modelling languages, models, tools), as well as the development processes. As a second major objective, the MOST approach creates an integrated ontology-driven, model-aware software process (ODSD). This development process will be guided by process guidance ontologies that formalize the

	rules, conditions and actions a software engineer has to conduct in specific development situations The MOST approach will be applied and evaluated in three case studies in the information systems, ERP, and telecommunications domains.
NEXOF-RA	Comprehensive multi-layer reference architecture for the NESSI Open Service Framework. EzWEB, MASTER, RESERVOIR, SLA@SOI, and SOA4ALL are all supposed to fit within the NEXOF reference architecture. Outputs include also validation implementations and an adoption roadmap. Specifications, templates. The project aims to find a good balance between openness (towards new contributions, standards, components) and coordination/predictability (i.e. how to ensure the reliable integration and execution of the components at any one time). General approach is not to use proprietary components. **Output**: The NEXOF reference architecture defines a catalogue of components and a catalogue of technologies. The specifications cover three specific domains: Internet of Services, Cloud, Enterprise-SOA. For each domain different architectural patters are defined: every pattern is composed by several components coming from the catalogue.
OMP	Dynamic composition of media-rich services for mobile devices to support social networking in mobile environments. OSS approach. QoS and resource management at network level required. Media stack API for app developers. **Output**: Open-source and commercial tools and methodologies that will facilitate component design, integration and deployment. Open-source and commercial tools and methodologies that enable dynamic composition of complex media-rich services on non-PC networked devices. A sample repository of reusable components that are portable across a wide rage of devices and that optimally exploit available resources. A standard schema for information exchange between tools allowing future advances in the dynamic creation of services utilizing innovative tool 'pipelines' A standard schema for run-time information exchange between media-components and service infrastructures that will enable future advances in Resource management and Quality of Service frameworks. Major advances in the Khronos OpenMAX and OpenKODE standards that will allow future networked devices to accommodate the most demanding application use-cases expected. Publicly available open-source reference implementations that will show complex application use cases on a range of platforms. Open source availability will facilitate widespread adoption.
OPEN	Environment supporting migratory services. This actually means that the state of the app or service is persisted so as to be able to be resumed when the user reappears on the network. Thus, rather than 'migratory' this sounds more like 'centralised', although the user perception will be that the service 'moved' or 'migrated' in order to follow him or her. **Output**: The expected output of the project will be an open service platform composed of several software modules able to support migration of services across various types of interaction devices and communication channels in such a way that users are enabled to naturally interact with such services. In particular the functional architecture is composed by the following layers: Presentation layer – Includes support functionality for reconfigurable multi-

	modal interfaces, enabling the users to interact with various types of services (at home and on the move) on different devices in different service environments. Application logic layer – Provides a service platform that enables service components to be dynamically reconfigured as context, device capabilities and migration allow. Middleware layer – OPEN middleware will support the migration process, by several network functionalities, such as trigger management to ensure the migration is triggered at the right time, synchronization for data streams when being migrated, device and service discovery, access to context through a context management system, and many more required to ensure a reliable, secure, and efficient service/application migration.
PERSIST	Smart context-aware environment that moves around with the user. Think of 'wrapping' the user so that user-specific functionality is more easily provided in a continuous manner whilst the 'smart space' takes care of minding the context changes. Sounds a bit like the shock absorber and suspension system of a car. Lots of interesting application scenarios become more easily implementable with this additional layering. **Output:** The vision of PERSIST is of a Personal Smart Space, which is associated with the portable devices carried by the user and which moves around with him/her, providing context-aware pervasiveness to the user at all times and places. The Personal Smart Space will cater for the needs of users, adapting to their preferences and learning new ones as these arise. The objective of PERSIST is to develop Personal Smart Spaces that provide a minimum set of functionalities which can be extended and enhanced as users encounter other smart spaces during their everyday activities. They will be capable of learning and reasoning about users, their intentions, preferences and context. The architecture consists of five different layers, each addressing a well defined part of the PSS (Personal smart space) functionality: Layer 1 – Devices: The PSS definition suggests that a single PSS can span across many different devices. Depending on their processing and networking capabilities, these devices may either implement the PSS stack or part of it, or simply interact with the rest of the PSS framework. Layer 2 - System Run-Time Environment: the System Run-Time layer serves as an abstraction layer between the underlying device operating system and the PSS software in order to achieve as much platform independence as possible. Layer 3 - Overlay Network Management: The Overlay Network Management layer provides the PSS architecture with a Peer-to-Peer (P2P) management and communication layer. The services within this layer will provide functionality for PSS peer group management, peer discovery and message routing between peer networks. Layer 4 - Service Run-Time Environment: The Service Run-Time Environment layer provides a container for the PSS services. It supports service life-cycle management features and provides a service registry, as well as a device registry. Layer 5 - PSS Framework: The PSS Framework layer is the core of the PSS architecture. Its functions include discovering and composing PSS and 3rd-party services, as well as managing context data and user preferences. Moreover, the PSS Framework layer supports automatic learning of preferences and inference of user intentions. This information, together with data from recommender systems, enables the proactive behaviour of the PSS platform. Grouping of context data and preferences, as well as conflict resolution, are also provided by the PSS Framework layer.

PROTEST	Test tools and methods for service-oriented networks. Testing, run-time monitoring, and event logs for off-line analysis. Specification language also needed to test against, where network 'properties' are specified. **Output**: The project will deliver methods and tools to support property-based development of systems. Property-driven development is a powerful new mechanism for gaining assurance of system reliability and functionality. However, in order to deliver its full benefits tools to integrate property-based testing into the development life cycle are needed. Property-based testing will deliver more effective tests, more efficiently: **Property discovery** Since current testing is based on sets of test cases embedded in test suites and current specifications and models are often informal. This project will provide tools to aid the software developers to extract properties from this test data; and develop specialised property languages to facilitate the formalisation of existing specifications. **Test and property evolution** As software systems change and evolve tools to support the evolution of tests and properties will need to be developed in line with the system changes. **Property monitoring** Not all properties can be tested in advance of systems being executed, and not all faults will be found, however thorough the testing. The project will provide tools to support the post hoc examination of trace details for conformance to (or indeed violation of) particular constraints. **Analysing concurrent systems** At the heart of service-oriented systems is concurrency: servers will provide services to multiple clients in a seamlessly concurrent way; services will unite to provide complex functionality through concurrently performing parts of a task. We will provide tools by which concurrent systems can be analysed for fundamental properties. The **Erlang/OTP** Open Telecom Platform has been chosen as our initial implementation vehicle due to its robustness and reliability within the telecoms sector. It is noted for its success in the ATM telecoms switches by Ericsson, one of the project partners.
Q-IMPRESS	QoS in service environments, which means across the whole stack and not only at the network layer. The objective is to enlarge the scope of relevance of service-oriented architecture to environments that have more stringent QoS requirements, such as industrial production, telecommunications, and mission-critical applications. **Output** One of the Q-ImPrESS outcomes is a prototypic integrated development environment, the Eclipse-based Q-ImPrESS IDE, that will assist software engineers during the iterative development and evolution both in existing as well as in newly started software development projects. The central element of a Q-ImPrESS-based development process is the Service Architecture Model (SAM), which is a new abstract design model of a software system, capturing information about a system's static structure, behaviour, deployment and usage scenarios. The Q-ImPrESS IDE offers both textual and graphical model editors for the simple and efficient creation and modification of the Service Architecture Model. The integrated SISSy (Structural Investigation of Software Systems) and SoMoX (SOftware MOdel eXtractor) tools also enable developers to automatically create a basic instance of the SAM based on the source code (Java or C++) of legacy systems. The Q-ImPrESS IDE internally uses meta-models of particular partners for the quality prediction analysis and simulation: the Palladio Component Model

	(PCM), the SOFA component model, ProCom component model and KLAPER. Based on these meta-models, concrete results are determined for the predicted performance, reliability and maintainability of each evolution alternative. Predicted metrics include response time for performance, failure probability for reliability and implementation costs for maintainability. The integrated quality impact and trade-off analysis will provide a complete picture of consequences of design changes at both the service level and the service architecture level. The aim of this approach will be to ensure possibilities for a software engineer to experiment with different service designs and architectures in order to find the one that fulfills given requirements.
RESERVOIR	Service utility delivery framework. Development of technical and legal infrastructure for facilitating cloud computing-based service delivery anywhere. The implication is the need to optimise the resources available through GRID and cloud computing infrastructures and across different administrative domains in order to respond to peak demands. This sounds similar to the load balancing electricity providers need to effect in order to match powerplant availability to user demand. The original term “utility computing” originated form this metaphor. Interesting theoretical problem. Architecture and reference implementation. Technical virtualisation will be aligned with SLAs across different administrative domains. Interface to business users and applications will be provided (Service Management Layer). **Output** In RESERVOIR, we are developing breakthrough system and service technologies that will serve as the infrastructure for Cloud Computing. We aim to achieve this goal by creative coupling of virtualization, grid computing, and business service management techniques Key technologies being developed in RESERVOIR include: Migration enablement of both virtual machines and Virtual Java Service Containers across network and storage boundaries. Distributed management of virtual infrastructures across sites supporting private, public and hybrid cloud architectures. Algorithms for the allocation of resources to conform to SLA (Service Level Agreement) requirements. The creation of a formal Service Definition Language to support service deployment and life cycle management across RESERVOIR sites. Security mechanisms for the safe deployment and relocation of virtual machines across physical machines, and RESERVOIR sites. The development of a business information model as well as business-oriented payment and billing mechanisms to charge for resources used across one or more RESERVOIR sites. Testbed development to benchmark performance of actual industrial use cases in a RESERVOIR environment. Automated Service Lifecycle management for service provisioning and dynamic scalability.
ROMULUS	Improvement of Java web app development through OSS meta-framework integrating different enterprise systems and IDEs to achieve high reliability, productivity, and security. Impact is expected in reduced costs and time for web development. Interesting integration of mashup-oriented development, soft QoS for service environments, and client-server ‘load balancing’. Approach is model-driven and assumes automatic code generation will take care of transformation to technologies of choice.

	Outputs ROMULUS's proposal is based on recognising some of the deficiencies of standard Java Enterprise Edition, and proposing a new paradigm for developing web applications taking advantage of new trends in software engineering, such as domain-driven design combined with agile development methodologies, and some of the principles from Ruby on Rails. In order to have a serious impact, the project does not start from scratch, it is based on two mature open source projects, **Roma** and **LIFERAY**: Roma will make Java application development easy, integrating frameworks and tools using a meta-framework and decreasing the implementation time. With LIFERAY, we are integrating the leading enterprise open source portal framework, with relevant industrial references. Objectives: Increasing productivity of web software development, focused on Java, based on a domain-driven design, MDA techniques and a metaframework. A key objective of the project is the usage of domain-driven design based on a meta-framework that abstracts from the frameworks, together with an MDA approach which generates the application from the meta-framework onto the de-facto standard and popular frameworks. Increasing software productivity by using mashup-oriented development. This objective will research on how web application development can be speeded up thanks to the reuse of existing services and components, as well as the defined methodologies. The project will research several types of mashups: *Mashup web services.* A key objective of this project is to research on how these mashup web services descriptions can be successfully used in the application domain model description, favouring the web service dynamic discovery and composition. *Data-level mashups.* A set of methodologies, best practices and guidelines will be defined in order to help the framework user during complex data mashups processes, while a semantic approach will be used to be able to publish internal data and import external data from the open web using RDF as interchange format, with no programming skills. *Portal mashup services.* A key objective of this project is to research how the definition of portal mashup services can reduce considerably the development time and can increase its reliability. The framework will provide portal services consumed by portal components and scripting languages. *Enterprise mashup services.* A key objective of this project is to research how the definition of an enterprise mashup service can combine information from enterprise search engines, Web services, messaging systems, BI engines and data integration solutions with information from external services. Enhancing the quality of SW development by involving soft goals such as reliability, traceability, security and performance, as well as other quality attributes from the conception of the software. Romulus will research security, traceability, reliability and performance soft goals, and their integration in the development process, based on the identification of soft goals-aspects and the usage of annotations against these aspects in the metaframework. Soft-goals annotations based on aspects, used both in the modelling and the implementation phases contribute to automatic code and tests generation. Balance between client technologies, server technologies and scripting languages. A key research objective of this project is to research the symbiotic combination of client and server technologies and how this can enable more usable application and faster integration or creation of pure client-side technologies. It will also integrate tools for using server-side scripting

	technologies in existing environments, such as portal applications.
S-CUBE	NoE. Integration and alignment of SaaS and IoS research agendas across Europe. The project will address: QoS & SLA, infrastructure, service composition, and business process modelling from two perspectives: engineering design and adaptation & monitoring. **Output** S-Cube will pursue the following objectives which will have a long-lasting impact on European research: Re-aligning, re-shaping and integrating research agendas of key European players from diverse research areas. By synthesising and integrating diversified knowledge, a long-lasting foundation for steering research and for achieving innovation at the highest level will be achieved. Inaugurating a Europe-wide common program of education and training for researchers and industry. This will create a common culture that will have a profound impact on the future of the field. Establishing a pro-active mobility plan to enable cross-fertilisation, which will foster the integration of research communities and the establishment of a common software services research culture. Establishing trust relationships with industry. Via European Technology Platforms (specifically NESSI) a catalytic effect in shaping European research, strengthening industrial competitiveness and addressing main societal challenges will be accomplished. Defining a broader research vision and perspective. This will shape the software service-based Internet of the future and will accelerate economic growth and improve the living conditions of European citizens. From technical point of view work in S-Cube will be guided by the S-Cube research framework, which clearly distinguishes between principles and methods for engineering and adapting service-based systems and the technology which is used to realize those systems while taking into account cross-cutting issues like Quality of Service (QoS) and SLA compliance.
SERVFACE	Interesting idea to extend the web services framework by adding user interfaces to the specification of services. The practical implications in terms of engineering process are considerable: overall methodology, integration of interface definitions in platform-independent models, adaptive user interfaces and UI composition, WS standard extension, IDE. **Output** The ServFace project aims at creating a model-driven service engineering methodology for an integrated development process for service-based applications. ServFace looks at this process from two different perspectives: First, the annotation of services with corresponding user interface annotations that are 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. Together with technical service descriptions like WSDL, 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 is being developed in ServFace will provide tool support for defining service annotations. For the composition of annotated services to complex applications, two alternative modelling approaches are 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 modelled visually by composing the application's UI from parts which are generated using the service annotations. The composition is supported by the ServFace Builder, an online tool that is being developed in the project. The second approach is supported by the MARIAE tool and provides a novel solution able to exploit task models (represented in the ConcurTaskTrees notation) and user interface models (in the MARIA language) for the design and development of interactive applications based on Web services for various types of platforms (desktop, smartphones, vocal, multmodal, …). Four primary objectives have been identified for ServFace: **Service Engineering Methodology.** The objective of ServFace is to design a methodology enabling the development of interactive applications involving the annotation of web services, the presentation-oriented (or, alternatively, task-oriented) composition of annotated services to complex applications, and the generation of executables from these composites. **Model-driven development approach.** Platform-independent models are foreseen to describe service annotations and composition models. Using appropriate inference rules and model-to-code transformations, these models are mapped to various technological platforms like Google Android and Microsoft® Silverlight™. **Service composition.** From the perspective of user interfaces, ServFace directly adopts concepts and technologies developed in the areas of user interfaces design, task-driven software development and end-user development for the composition of annotated services to complex applications. **Integrated development environment**. An integrated tool environment is being created to support all development phases of interactive applications that build upon annotated services. This comprises the annotation tool and authoring tools supporting presentation-oriented and task-based composition approaches for annotated services. Model-to-code transformations map these technology-independent models to executables for concrete run-time environments.
SHAPE	Service modelling based on OMG UPMS standard for semantically-enabled heterogeneous service architecture. A big effort of abstraction to hide the differences in underlying architectures (agent, p2p, Grid) and provide a uniform service development environment that offers an interface to business users. OSS approach. Relevance to oil and steel sectors. Sounds a bit optimistic. **Output** The SHAPE (**S**emantically-enabled **H**eterogeneous Service **A**rchitecture and **P**latforms **E**ngineering) project provides an integrated development environment that brings together Model-driven Engineering (MDE) with the Service-Oriented Architectures (SOA) paradigm. This is complemented with innovative service engineering techniques that appear to be desirable in real-world applications, including support for flexible business modelling, the customization and personalization of services, the integration of agent technology and adaptive systems, and support for employment of semantic technologies. The project defines the necessary metamodels for describing services and related aspects in heterogeneous technology landscapes, develops an integrated tool suite for modelling and deployment, and provides a comprehensive methodology for guiding software engineers and architects through the engineering process of heterogeneous service-based systems.

	This brings together the world of Model-driven Engineering (MDE) with the Service-Oriented Architectures (SOA) paradigms and integrates other technologies like Software Agents, Semantic Web, Grid, and P2P, combining their respective advantages for the effective development and maintenance of high-quality integrated IT systems. The model-driven engineering techniques developed in the project are centred around SoaML, a metamodel for describing service-oriented landscapes that is standardized in OMG, which is extended with metamodels for the other technology platforms and advanced service engineering techniques. The project further develops the necessary infrastructure for employing the modelling techniques in real-world applications. This consists of an integrated tool suite that supports the modelling on the basis of the metamodels and encompasses the necessary model transformations, and a comprehensive methodology framework that provides role-specific guided procedures for creating all model types and supports the creation of customized methodologies for individual engineering projects. The following outlines the relevance of the techniques: **Flexible Business Modelling:** this enables integrated and guided modelling of service-oriented systems at the business level, thereby overcoming the deficiencies of previous not aligned business modelling techniques. A particular benefit of the SHAPE solution is the provision of semi-automated transformations of BPMN business process models to SoaML, which help in bridging the business-IT gap and thereby enable the flexible and rapid adaptation of business solutions to changes in the environment. **Service & SOA Modelling:** this addresses the central demand for enabling model-driven engineering for services and service-oriented systems. In SHAPE, this is facilitated by SoaML, a metamodel for describing service-oriented landscapes that is standardised in OMG. SoaML supports the modelling of services as well as their interaction and usage in service-oriented systems, and its standardization is expected to ensure a wide usage. **Integration of Agent Technology:** this facilitates the integrated modelling and development of systems that employ intelligent software agent technology in order to realize automated planning on top service-oriented landscapes, which is relevant for flexible scheduling or production line optimization in industrial applications. **Integration of Semantic Web Service Technologies:** this enables the integrated usage and modelling of semantic technologies in service-oriented landscapes, which allow the integration of heterogeneous resources that need to interoperate, and support the automation of service consumption tasks in larger applications. **Service Customization:** customization and personalization of services is particularly relevant in business applications where large and complex services need to be handled and consumed. SHAPE provides a novel technique for this that enables the model-driven creation of simplified variants on the basis of explicit variability modelling.
SLA@SOI	Service-Level Agreement Management Framework **Output**: SLA@SOI will provide its results in 3 complementary ways: First, an open source-based SLA management framework will allow realisation of the benefits of predictability, transparency and automation in an arbitrary service-oriented infrastructure. Second, in-depth guidance for industrial stakeholders will be given explaining the best practises on how to transform their service business into an SLA-

	driven one. Finally, SLA@SOI will provide an open reference case which allows for stakeholders to re-run, re-validate and even modify SLA experiments in the context of a concrete application. The technical approach of SLA@SOI is to define a holistic view for the management of service-level agreements (SLAs) and to implement an SLA management framework that can be easily integrated into a service-oriented infrastructure (SOI). The main innovative features of the project are: An automated e-contracting framework, Systematic grounding of SLAs from the business level down to the infrastructure, Exploitation of virtualization technologies at infrastructure level for SLA enforcement, and Advanced engineering methodologies for creation of predictable and manageable services.
SMARTLM	Software licensing for location-independent application execution Traditionally, software licenses have been provided on the basis of a named user, node-locked host, number of concurrent jobs or possibly a floating site license. These models are not sufficiently flexible to support commercial applications that access resources beyond the current administrative domain – possibly as a Utility-like service outside the organisation or Software as a Service (SaaS) model. Software licensing is identified as a particular concern for enterprise IT managers as they start to deploy virtual Grids in any meaningful way. For all the potential benefits of Grids, IT departments cannot afford to buy software licenses for every device in the service-oriented infrastructure that by nature consumes resources dynamically. In addition, the Grid-based Service-Oriented Architecture (SOA) solutions together with other technologies trends such as multi-core and virtualization environment are forcing Independent Software Vendors (ISV) to move away from traditional software licensing models. **Output**: The overall approach consists in treating and implementing software licenses as Grid services thus providing platform independent access just like other virtualized resource. Licenses will become Grid services; a promising approach to overcome the limitations of current monolithic licensing models. Licenses will be managed as agreements, extending the conventional Service Level Agreements (SLAs) which are made today between sellers and buyers in the market. Licenses will be dynamic in order to support agreements that may change over time and where the dynamic negotiation between service provider and consumer is needed.
SOA4ALL	A large IoS project aiming to harmonise the integration of the outputs of the other SOA projects. **Output**: The project will provide a *comprehensive framework and infrastructure* that integrates five complementary and revolutionary technical advances into a coherent and domain independent service delivery platform. The overall architecture of SOA4All can be structured into four main parts: **SOA4All Studio**: A rich Web platform that provides users with a unified view covering the whole lifecycle of services, including design-time, run-time and

	post-mortem analysis. **SOA4All Distributed Service Bus (DSB)**: The infrastructural backbone around which all the SOA4All components communicate and collaborate by combining Semantic Spaces and Enterprise Service Bus. It is is a direct evolution of the open-source PEtALS Enterprise Service Bus (ESB) that is promoted by the OW2 Consortium. To realize this extension, the DSB will benefit from the internal use of the ProActive distributed programming language also promoted by OW2. **SOA4All Platform Services**: The group of services that provide the basic SOA4All functionality and activities, such as Service Ranking and Selection, Service Discovery, Service Adaptation, Service Composition, Service Execution, and the Reasoning Engine. **Business Services** (3rd-party Web services and light-weight processes): The actual services provided by final users. The SOA4All framework will be as technology agnostic and less intrusive as possible.
STREAM	Streaming middleware **Output**: The project aims at providing a highly scalable cloud computing platform to enable a new breed of services. The core is the data streaming platform, StreamCloud, that will be able to parallelize the processing of information flows in large clusters of 100s sites. Current approaches fail to scale for massive information flows. Stream aims at boosting the scalability of current approaches by 1 to 2 orders of magnitude. The Stream platform will provide elastic computing, so the computing resources as used as required by the incoming load. Below the core, there is a high-performance communication layer that enables efficient interaction among sites with access between node memories of tens of microseconds in contrast with tens of milliseconds using current technology. Additionally, this layer will provide parallel IO and low-cost storage for huge amounts of information. Above the core, there is a data-mining layer offering higher-level services to facilitate the development of applications processing the information flows. On top of the data streaming platform there is the application layer in which user applications & services will run.

Table 1 – Brief descriptions of Call 1 Projects

Figure 3) that 12 projects out of 25 will end within 2010. 10 will end within March 2011 and only three projects will continue after March 2011.

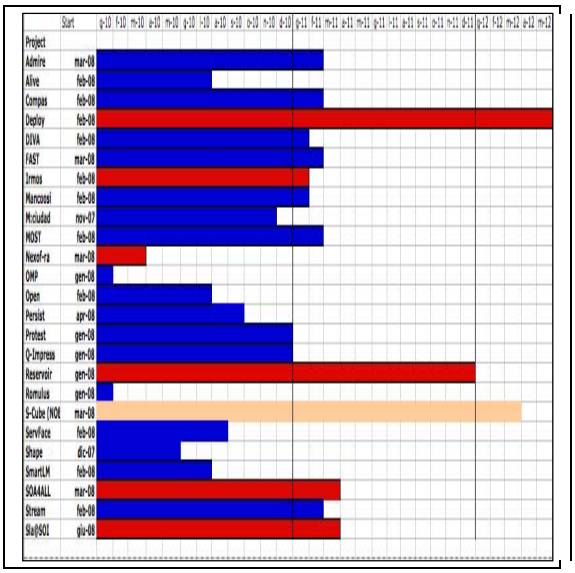

Figure 3 – Project duration

In Annex A the reader will find the first exercise done by SEQUOIA's researchers in terms of projects understanding; the first attempt was to summarise project goals, main technological expected outputs and their beneficiaries and, most importantly, to delineate a complete list of the available documents.

CHAPTER 2. ORGANISING THE PROJECTS UNDER ANALYSIS

In the previous chapter, we briefly described each project under analysis and its main outputs. We took each project separately into consideration, and in this chapter, we will see if it is possible to discover some 'likeness' among the projects in terms of users/beneficiaries and in technological terms. The main objective of this chapter is, therefore, to provide the relational links between the 25 projects under analysis, this for the SEQUOIA Consortium's internal usage. This exercise, in fact, facilitates the understanding of the nature of these projects assisting us to understand the centrality of the various beneficiary categories and the relevance of different technological solutions underpinning Call 1 Projects. It is important to emphasise that this exercise does not impact directly the SEQUOIA methodology developed; rather, it is an instrument useful for the SEQUOIA researchers to have a synthetic view of the projects from the specific prospective of their expected beneficiaries and expected technological outputs.

The resulting project relational exercise builds on the factsheets of the Call 1 Projects. All factsheets were entered in a content analysis package, and based on user or beneficiaries mentioned in the factsheets a thematic analysis for the qualitative data could be conducted. The next figure presents a screenshot of the analytical process for the ADMIRE factsheet including content analysis (yellow) and coding schema.

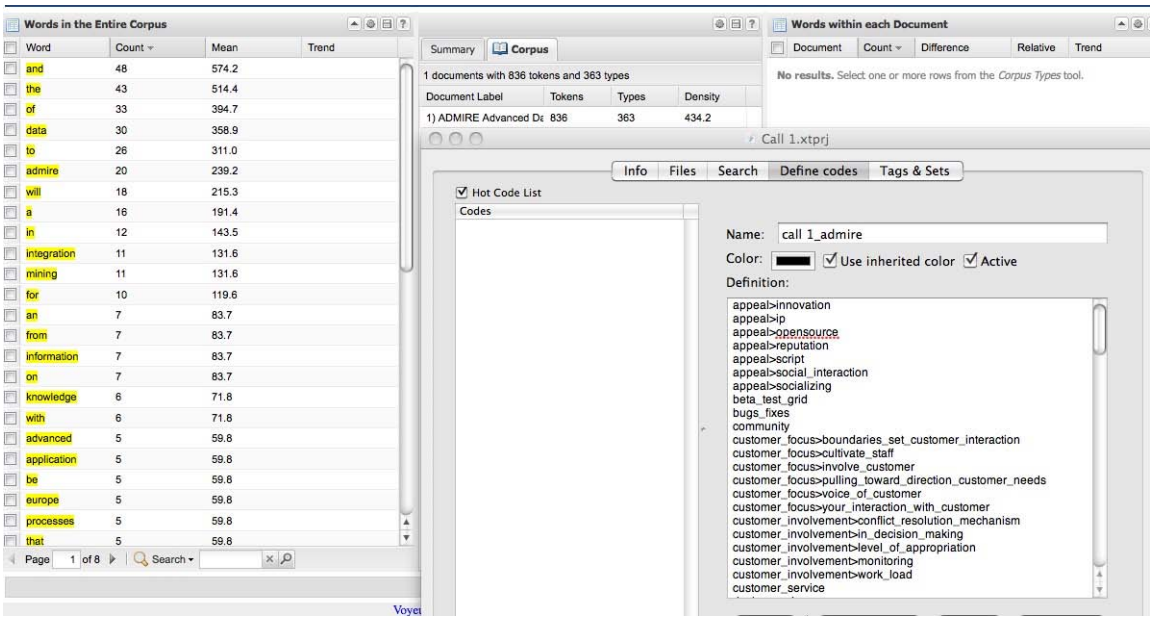

Figure 4 – Screenshot thematic analysis; example ADMIRE project

In other words, this exercise identified common elements (or, words) associated with beneficiaries reported by the participating projects across qualitative data (c.f. Haythornthwaite and Gruzd, 2007)[1]. This resulted in the organisation of the qualitative data into themes or categories associated with users or beneficiaries of the projects (Fereday and Muir-Cochrane, 2006)[2]. Four thematic categories could be developed for describing projects'

[1] Haythornthwaite, C., & Gruzd, A. (2007). A Noun Phrase Analysis Tool for Mining Online Community Conversations. Paper presented at the International Conference on Communications and Technologies, East Lansing, Michigan.

[2] Fereday, J., & Muir-Cochrane, E. (2006). Demonstrating Rigor Using Thematic Analysis: A hybrid approach of inductive and deductive coding and theme development. International Journal of Qualitative Methods, 5(1), 80-92.

main beneficiaries and as a common denominator to illustrate the relationships between and among projects:

1) Academia
2) Developers
3) End Users
4) Industry

The first beneficiary group includes the wider research community, and which can include universities, individual scholars, commercial R&D departments, etc. The second category refers to those involved in the development of software and ICTs in general, such as developers, programmers, software engineers, etc. The third category refers to 'people at home' and can include consumers, citizens, etc. The last group refers to commercial practice or companies, and which can include service providers, infrastructure providers, telcos, SMEs, large corporations, etc.

Table 1 yields insight into the sense-making process of the results of the data analysis involved identifying relationships between the thematic categories, exploring the potential stakeholder targets or beneficiaries for each project. For example, NEXOF-RA targets developers and end-users such as citizens, while developers such as software engineers and researchers are most likely to benefit from MOST.

	Academia	Developers	(End) Users	Industry
ADMIRE		X		X
ALIVE			X	X
COMPAS				X
DEPLOY	X	X		X
DIVA			X	X
FAST			X	
IRMOS			X	
MANCOOSI		X		
MCIUDAD			X	
MOST	X	X		X
NEXOF-RA		X	X	
OMP		X	X	
OPEN			X	
PERSIST			X	
PROTEST		X		X
Q-IMPRESS		X		X
RESERVOIR	X			X
ROMULUS		X		
S-CUBE	X			
SERVFACE	X	X		
SHAPE				X
SLA@SOI				X
SMARTLM			X	X
SOA4ALL			X	X
STREAM		X		

Table 1 - Call 1 projects' users/beneficiary dimension

After this exercise, we thought it would be useful to generate a visualisation for this information in order to facilitate, the understanding about how the projects can be grouped in relation to their users/beneficiaries.

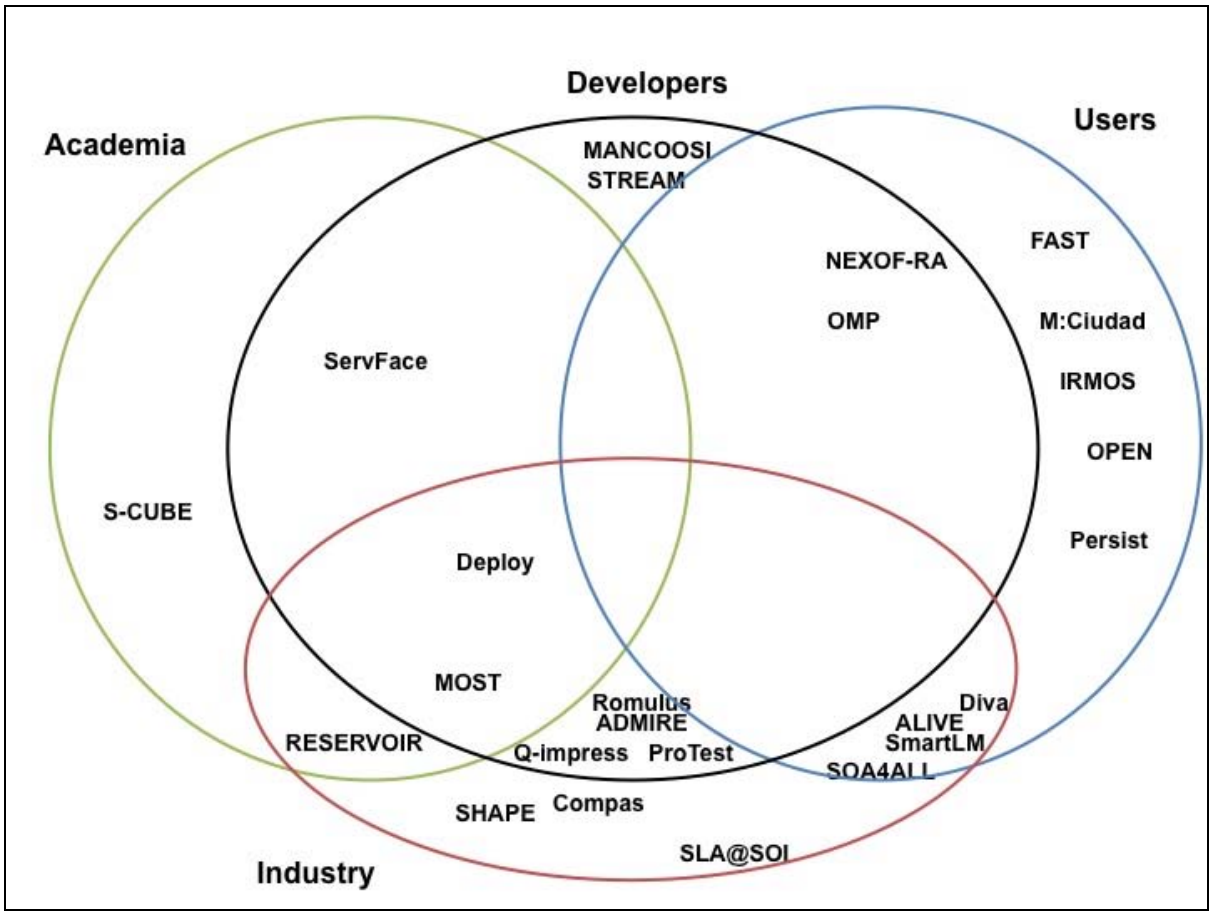

Figure 5 – Venn diagram of Call 1 projects' users/beneficiary dimension

From this figure it becomes clear that the research community is the least targeted beneficiary for Call 1 projects. Five projects specifically target end users, three projects target the industry, and two aim after developers. Most projects have two or three beneficiaries groups, and based on the documents provided by the projects it is not clear-cut whether one is possibly more important than another. For this further research is needed. However, within the scope of the current exercise to make visible the relationships among projects and their beneficiaries, this is of less relevance.

The projects were also considered according to their technological dimension. The following picture shows a tag cloud that resumes different aspects related to the 25 projects, in particular focusing on technologies and architectural paradigms. Describing cross-area aspects, it could be considered complementary to the classification reported in figure 2 presented earlier in this deliverable. It is easy to notice that most of the projects are related to Service Oriented Architecture (SOA), Software as a Service (SAAS) and in particular to web service technologies. On the other hand, Grid and Cloud computing are also two aspects investigated in some projects where the themes of virtualisation and distributed systems are relevant. Moreover, tag cloud shows that information mash-up and mobile technologies are important topics taken into account in some projects.

Figure 6 – Tag Cloud of Call 1 projects' technological dimension

In sum, this chapter has sought to yield insight into the projects under investigation, targeted beneficiaries and technologies involved. This grouping exercise may be particularly relevant for the definition of the variables that are going to be included in the impact assessment method and information gathering instruments at a later stage. We need, in fact, to have a clear map of the similarities among the projects in order to reduce the complexity of the universe under analysis.

CHAPTER 3. PRELIMINARY OVERVIEW OF PROJECTS' AVAILABLE DOCUMENTS

Before introducing the deliverables selected as relevant from the projects under analysis, we will briefly describe a sample of the theoretical documents used in order to orient the methodology development. These documents were of great help in starting to develop the questions for the in-depth interviews with key informants and, generally, to increase our understanding of various implications of the research area represented by the projects under analysis.

3.1 – Theoretical documents

Next-generation connectivity: a review of broadband internet transitions and policy from around the world, 15 February 2010.

Berkman Center Broadband Final Report

This study analyses the current policies and practices pursued by different countries in the transition to the next generation of connectivity, as well as their past experiences. The review, that concerns different aspects (competition and open access policy, mobile and nomadic access, government investment practices) aims at analysing best practices and identifying the challenges of the next-generation transition, together with the range of possible solutions.

The Socio-economic Impact of Social Computing, Proceedings of a validation and policy options workshop, 2008.

IPTS Exploratory Research on the Socio-economic Impact of Social Computing

This report presents the major outcomes of the workshop discussions held at IPTS in Seville, on February 2008. The objective of the workshop was to validate the results of the research on social computing and to identify key areas for future EU research and innovation policies. In particular the discussion regarded the three key reports presented (one presenting an empirical analysis of the creation, use and adoption of social computing applications, one on social computing and collaborative content and one on social computing and social networking). In addition, brainstorming and open discussion sessions were held on the position of Europe in social computing and on policy options for future EU research and innovation.

Longer term research challenges in Software & Services, 2008.

Future service report

This report outlines open research challenges faced to achieve the transition to a future generation of software services and service-based applications that form the future Internet of Services. In particular, it identifies 6 main software, service and technology-specific issues and research challenges: High-level Service Description Languages, Service Design Theories, Service Domains, Service Platforms, Service Quality, and Software Service Life Cycle.

Guide to measuring the Information Society, 2009.

OECD

The Guide documents the work of the OECD (mainly the Working Party on Indicators for the Information Society - WPIIS) and others in developing statistical standards for measuring key properties of the information society. It is a compilation of concepts, definitions, classifications and methods for information society measurement and analysis. It is articulated around the following chapters: ICT products, ICT infrastructure, ICT supply, ICT demand by businesses, ICT demand by households and individuals, Cross-cutting topics in information society measurement.

Social Computing: Study on the Use and Impact of Online Social Networking, 2008.

IPTS Exploratory Research on the Socio-economic Impact of Social Computing

This report presents the results of a case study on Social Network Sites recently conducted by JRC IPTS, as part of an exploratory research project. The study aims to explore and identify the social and economic implications of SNS in Europe and to suggest policy options and avenues for further research. It provides an analysis of digital applications that facilitate social networking and multimedia interaction amongst individuals, highlighting changes and their implications in how people network, manage and operate their social contacts. Variables of social impact are presented.

Innovation in the Software Sector, 2009.

OECD

This document presents the main aspects that regard innovation in the Software Sector. In particular it deals with the following issues: Measurement challenges, Software innovation processes, the software sector innovation environment and Emerging trends, changing environment and growing challenges.

Consultation on Future Research Priorities on Software & Services.

On-line Consultation process of the European Commission, DG Information Society and Media, Service and Software Architectures and Infrastructures Unit (INFSO D3)

This documents presents the various contributions made during the forum-based consultation process opened by the European Commission in 2009 for the definition of the Work Programme 2011-2013 in the field of software and services.

Trends in connectivity technologies and their socioeconomic impacts – Final report of the study: Policy Options for the Ubiquitous Internet Society, 2009.

Rand Europe

This report suggests a policy framework for addressing the future challenges of the 'Ubiquitous Internet Society'. It has been written for the European Commission's DG Information Society and Media to support the development of its policies for the period 2010-2020. It presents the major Connectivity Challenges and it contains trend analysis, scenario development, modelling of socio-economic impacts and a review of changing business models. In addition, an analysis was made of policies in the US, Japan and South Korea to provide a reference for the EU's own policy in the field of ICTs and especially the future of the Internet (its architecture and socio-economic fall-out).

Study on the Social Impact of ICT (CPP N°55A – SMART N°2007/0068), Final Report D7.1, 30 April 2010.

Universität Siegen, Fachbereich Wirtschaftsinformatik und Neue Medien, Germany

This document presents a summary of the main research findings from the 'Social Impacts of ICT' study that has studied the social impacts of ICT, through the analysis of available empirical evidence about developments in some of the most relevant domains in Europe, and beyond. It presents the key results of the study from the seven domains selected for analysis: participation in policy-making; education and lifelong learning; work; consumption; health; community and family; and creation and distributed innovation. A separate section looks into the main differences between developments in Europe and selected non-European countries, including the US, Japan and Korea. The report concludes with some generalisations and recommendations.

Study on the Social Impact of ICT (CPP N°55A – SMART N°2007/0068), Topic Report 3 (D7.2), Final Version.

Conceptual Framework, Vertical domain reports, Horizontal domain report, Report on

Findings from Flash EB.

Universität Siegen, Fachbereich Wirtschaftsinformatik und Neue Medien, Germany

This document presents the specific research results for each one of the seven domains mentioned before (participation in policy-making; education and lifelong learning; work; consumption; health; community and family; and creation and distributed innovation). For each domain, there is a description and an analysis of what has happened in the past 25 years in order to answer the questions: "which impact ICT has had on these events? What would have happened without the advent of this technology?". The conclusion reached is that the transformative potential of ICT is more of an evolutionary than a revolutionary nature, considering ICT as a trend amplifier.

The Role of e-Infrastructures in the Creation of Global Virtual Research Communities, Final Report, February 2010

eResearch2020

The document presents the results of the research conducted by the eResearch2020 consortium on a diverse sample of e-Infrastructures from around the world, involving both developers and users. The research aimed at better understanding the role of e-Infrastructure in the creation of Global Virtual Research Communities in order to improve policy, enhance technology adoption and facilitate the creation of global virtual research communities. On the basis of the empirical findings (obtained through a survey and case studies), a roadmap has been elaborated to inform research policymakers and e-Infrastructure developers about critical issues in e-Infrastructures for research in the European Research Area and beyond that must be addressed in the coming decade. The roadmap presents 4 possible scenarios for the future. Finally, the Roadmap concludes with a series of recommendations for action.

Economic and Social Impact of Software & Software-Based Services (Smart 2009/0041). D2 – The European Software Industry, 30 july 2009.

Pierre Audoin Consultants SAS (PAC)

The purpose of the study will be to identify the potential economic and social impact of software and services within the future Internet; the elements that are determinant for growth and competitiveness; the actions that need to be taken to implement them; the current barriers and the actions needed to remove them. In particular, this report is about the definition of the Software and Software-Based Services (SSBS), its market analysis, the emerging segment, the economic and social contribution, and provide a policy overview.

Economic and Social Impact of Software & Software-Based Services (Smart 2009/0041). D3 – Baseline Scenario for 2020, 21 january 2010.

Pierre Audoin Consultants SAS (PAC)

This report follows the D2 report as part of the "Economic and Social Impact of Software and Software-Based Services" in Europe. The objective of this document is to project the most likely economic and social impact of the EU SSBS industry on Europe (EU27) up to 2020, with a specific focus on the impact of the development of the Internet of Services on this industry.

3.2 – Call 1 project's selected deliverables and papers

In this session we present a short description of the deliverables and paper that we selected as interesting for understanding Call 1 projects and the related research area. These documents, together with the outputs of preliminary questionnaire-based survey and of interviews to key experts (see D2.2a) guided the definition of SEQUOIA assessment indicators.

Project	Deliverable description
ADMIRE	**"D1.6 Report on progress of model, language and ontology research"** The deliverable reports on the progress of the model, language and ontology research conducted within the project and an implementation done using DISPEL implementing a data-mining pattern. Language and model elements are discussed together with IDE integration of DISPEL. Two main ontologies are created for the reference cases of the project. **D2.5 Research Prototype: Parallel Processing a DMI task** Objective of the document is explore a generic parallel processing model through iterative prototyping, identifying, processing Elements in a DMI task, and describing the DMI task in DISPEL. The deliverable reports the results of measuring the performance and scalability of the system by varying numbers of computing nodes and sizes of datasets. The prototype provided gave useful feedback to various work packages: to deal with an automatic optimisation of a DMI task at both logical and physical levels **Supporting Molecular Modeling Workows within a Grid Services Cloud**. M. Köhler, M. Ruckenbauer, I. Janciak, S. Benkner, H. Lischka, W. Gansterer, International Conference of Computational Science and Its Applications (ICCSA), Fukuoka, Japan. 2010. The paper describes a service oriented approach based on the Vienna Grid Environment (VGE) integrating the Ubuntu Service cloud supporting scheduling and dynamic and partioned workflow, tackling scientific workflow accessing HPC infrastructures **An Ant-Colony-Optimization based approach for Determination of Parameter Significance of Scientific Workflows**. Fakhri Alam Khan, Yuzhang Han, Sabri Pllana, and Peter Brezany. International Conference on Advanced Information Networking and Applications, Perth, Australia, 2010 The paper explores a heuristic-guided approach which enables to find a near-optimal solution using a reasonable amount of resources without the need for the evaluation of all possibilities of values and variable of the scientific experiment, for determination of parameter significance of scientific workflows based on Ant Colony Optimization (ACO **Cloud-Enabled Scalable Decision Tree Construction**, Yuzhang Han, Peter Brezany, and Ivan Janciak. International Conference on Semantics, Knowledge and Grid 2009, 12-14 Oct 2009, Zhuhai. The paper describes the steps to build a service-oriented distributed computational system based on the cloud concepts named distributed computational service cloud, implemented with Web-services-resource-framework-compliant (WSRF) Web services, enabling high-performance and distributed computing. The paper evaluates the cloud using scalable decision tree service, providing computational intensive data mining algorithm.

ALIVE	"**Norms, Organisations and Semantic Web Services: the ALIVE approach"**, *Sergio Alvarez-Napagao, Owen Cliffey, Javier Vazquez-Salceda and Julian Padget*, Novembre 2009 ALIVE is an EU FP7 STREP whose goal is the convergence of organisational and normative modelling with and service-oriented architectures (SOAs) using model-driven software engineering. The project provides a framework for designing and implementing systems, taking into account organisational, coordination and service perspectives. A key project aspect is the integration of normative systems with live SOAs, through the distributed monitoring of normative state. Here we give a brief overview of the project, explore of the domain from a service context, outline the architecture under construction and sketch the use-cases that illustrate and inform the project. "**PaTac: Urban, ubiquitous, personalized services for citizens and tourists**," *Luigi Ceccaroni, Víctor Codina, Manel Palau, Marc Pous*, January 2009 This paper presents the general design of an architecture, based on software agents and oriented to the semantic Web, for the development and deployment of urban, ubiquitous services for citizens and tourist. The goal is to create a platform able to provide personalized services based on recommendation algorithms, and users' location, profile and preferences. **Use Case Models and Prototypes (version 1),** September 2009 This document presents models for each use case, together with a description of software prototypes developed to address the use cases. The document includes information on design decisions, modelling issues and validation. **Service-level Tools and Integration Tools**, September 2009 Report summarising the tools developed as part of the ALIVE project in order to support the development of novel semantic web services and integration of existing services, into service oriented architectures based on organisational models. This document describes all tools developed as part of the project service layer (work package 5).

COMPAS	**D2.1 State-of-the-Art in the Field of Compliance Languages** The deliverable provides an overview of the state-of-the-art in compliance languages, particularly focusing on languages for regulatory and legislative provisions. A wide range of compliance legislations is discussed. Specifically, the legislations defined within Basel II, FINRA (NASD/SEC) regulations, HIPAA, IFRS, MIFID, Sarbanes-Oxley and Tabaksblat are analysed. This set of regulations constitutes a faithful representation of the range of compliance requirements that can be found within compliance legislations. A variety of business process modelling languages is surveyed to examine their support for the expression of basic compliance concerns. The document highlights that current industry standards such as WS-BPEL and BPMN are capable of addressing some, but not all, of the basic compliance concerns (most notably regarding locative and temporal requirements) and as such the may require extension in some form. **D2.6. Implementation of an Integrated Prototype handling Interactive User Specified Compliance Requests in a Compliance Language** This deliverable presents the initial version of the prototype D2.6 (Compliance Request Language Tool), which will be finalized in M35. It describes how the prototype is integrated with the COMPAS Architecture and briefly presents the requirements and the design of the prototype. It also describes the current status of the implementation together with the technology choices utilized **D3.1 Specification of a Behavioural Model for Services** This report presents a formal model for unambiguous specification of business process behaviour suitable for verification of process compliance to formally specified legislative norms and requirements. Target is a model checking of various process properties including structural, temporal, transactional, resource-aware and Quality-of-Service (QoS) characteristics. The approach adheres to the principle of model-driven development and establishes the connection of the proposed formal tools with popular business process modelling notations such as BPMN, UML or WS-BPEL. **D1.2Core Meta-models, Transformation Templates, and Languages** Ensuring business compliance in process-driven SOAs is a tedious and error-prone task because the stakeholders confront two challenges: the increasing complexity of process descriptions and the gap between abstraction levels due to the difference of language syntax and semantics, the difference of granularity, and the lack of supporting links between high-level and low-level process languages. In this deliverable, a MDSD Software Framework is presented comprising a view-based, model-driven approach and various DSLs tooling in order to ease the process development in a flexible, extensible manner. The MDSD Software Framework will contribute a foundation for modelling and development processes and business compliance as well as provide appropriate means for integrating COMPAS partners' efforts at both conceptual and technical levels. **D4.2 BPEL Extensions for Compliant Services** This deliverable defines extensions to the Business Process Execution Language (BPEL) in order to enable the implementation of the compliant service composition specifications. The approach of the research aims at ensuring a faster and more consistent specification and integration of compliance checks into business processes.

DEPLOY	**Abrial, Jean-Raymond and Butler, Michael and Joshi, Rajev and Troubitsyna, Elena and Woodcock, Jim C. P. (2010) 09381 Extended Abstracts Collection — Refinement Based Methods for the Construction of Dependable Systems. Dagstuhl Seminar Proceedings, 09381 (09381). Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany.** With the growing reliance on computers, the total societal costs of their failures are hard to underestimate. Nowadays computers control critical systems from various domains such as aerospace, automotive, railway, business etc. Obviously, such systems must have a high degree of dependability — a degree of trust that can be justifiably placed on them. Although the currently operating systems do have an acceptable level of dependability, we believe that they development process is still rather immature and ad-hoc. The constantly growing system complexity poses an increasing challenge on the system developers and requires significant improvement on the existing developing practice. To address this problem, the book investigates how to establish a set of refinement-based engineering methods that can provide the designers with a systematic methodology for development of complex systems. **Jones, Cliff B and Roscoe, A W and Wood, Ken, eds. (2010) Reflections on the work of C.A.R. Hoare. Springer. ISBN 978-1-84882-911-4** Reflections on the Work of C.A.R. Hoare presents a comprehensive edited survey of all aspects of the subjects, with original contributions by more than 30 international leaders in the field. The book, while honoring Hoare's important contributions, assembles a collection of chapters showing the state of the art in programming languages, sequential programs, concurrency, and unified theories of programming. The book provides a scientific biography of Tony Hoare, demonstrates a principled combination of CSP and functional programming, and a CSP semantics for the π -calculus, reviews methods for proving Hoare formulae, investigates developments in game semantics and semantics based on the state monad, examines the satisfiability problem, void safety, and issues of auxiliary variables, Introduces type families in Haskell and a description of Quicksort in Orc, describes an experiment using the Tokeneer archive, and a correctness proof of cryptographic protocols based on the Shadow Security model, presents a representation of CSP as a bigraphical reactive system, and shows how simple entities can be related to CSP processes, discusses the problem-frames approach, and explores algebraic properties of the new programming combinators. This accessible monograph is an ideal overview of theoretical and empirical evolution in programming logic and semantics of programming languages. It will serve as an invaluable resource for professionals, researchers, libraries, and students who are interested in broadening their knowledge in all of the areas covered **Butler, Michael and Jones, Cliff B and Romanovsky, Alexander and Troubitsyna, Elena (2009) Methods, Models and Tools for Fault Tolerance. LNCS, 5454 . Springer.** This book is an outcome of the workshop on Methods, Models and Tools for Fault Tolerance, MeMoT 2007, held in conjunction with the 6th international conference on Integrated Formal Methods, iFM 2007, in Oxford, UK, in July 2007. The authors of the best workshop papers were asked to enhance and expand their work, and a number of well-established researchers working in the area contributed invited chapters in addition. From the 15 refereed and revised papers presented, 12 are versions reworked from the workshop and 3 papers are invited. The articles are organised in four topical sections on: formal reasoning about fault-tolerant systems and protocols; fault tolerance: modelling in B; fault tolerance in system development process; and fault-tolerant applications. The growing complexity of modern software systems makes it

	increasingly difficult to ensure the overall dependability of software-intensive systems. Mastering system complexity requires design techniques that support clear thinking and rigorous validation and verification. Formal design methods together with fault-tolerant design techniques help to achieve this. Therefore, there is a clear need for methods that enable rigorous modeling and the development of complex fault-tolerant systems. **D4.1 Report on Pilot Deployment in Business Information Sector** The pilot deals with modelling service choreographies. The approach taken is to reuse existing development models for service choreographies, which are written using a domain specific diagrammatic modelling language, and translate these models automatically to Event-B. The proofs, validation and exploitation of the models (e.g. to produce test cases) take place in the background. Topics of relevance include how to guide the developer through patterns and improve the degree of automation, as well as how to give modellers feedback on the diagrammatic level.
DIVA	**D1.2 Framework for Identifying and Modelling of Dynamic Variability in Requirements** In this document focuses on answering the question on what (requirements, their justification and context) to model, as well as the issue of how to restrict the proliferation of variants in requirements models. To address these questions, the DiVA RE approach proposes to construct a feature tree from the input requirements text, thus identifying specific requirements that need to be modelled, and refining this tree with variations and adaptation points; to recover the hard- and soft-goals related to the given requirements to provide a justification for them, and also to use a single goal-model for modelling the large set of potential variations; to analyse the resultant goal/feature tree to extract the relevant context and related variability and constraints which restricts the set of potential configurations. integrate support for composing dynamic variability. The research will follow into the requirements models, as well as for early configuration and trade-off consideration. **D2.2 - Transformation framework** This deliverable describes the first version of the model transformation framework developed by the project. Objectives are to develop a domain-specific language for adaptive system specifications, and a composition and transformation framework. This deliverable consists of three parts. The main part is the tools and metamodels developed so far. The second part is video tutorials showing the usage of these tools, while the third part is this document. This document presents the domain-specific modelling language for adaptive system specification and the tools that support specification, model checking and design time simulation of the adaptation model of adaptive systems. The proposed approach combines aspect-oriented and model-driven principles to cope with the combinatorial explosion and provide model level representations of variants, context variables and adaptation rules. An open source Eclipse-based editor has been developed to support system specification according to the meta-model. Based on the implemented adaptation meta-model, model checking and simulation facilities are provided to prepare for safe adaptation. At runtime the adaptation models are used to drive the adaptation. The approach is validated through case studies **D3.3 Reference architecture - final version** This deliverable presents the final version of the reference architecture to support dynamic variability using model-driven engineering techniques and aspect models. This reference architecture leverages the design-models, as well as the

	reasoning techniques developed in the project, at runtime. The purpose of this document is to provide an overview of the reference architecture, to detail the important parts of this reference architecture and give some implementation details. This document is associated with a software system, demonstrating the reference architecture, which is integrated into DiVA Studio. The reference architecture is based on OSGi, the former one was based on Fractal. In this new version, aspect models are now compiled into Java code, which makes it possible to weave aspect model at runtime, in an efficient way. Details are given for the compilation process of SmartAdapters. The reference architecture now provides support for distribution and describes the architectural extensions to support distribution as well as co-existence and co-dependency. Furthermore it provides a distributed consistency framework to support distributed aspect configuration and reconfiguration. Nelly Bencomo, Jaejoon Lee, and Svein Hallsteinsen (2010) **How dynamic is your Dynamic Software Product Line?** In: 4th International SPLC Workshop on Dynamic Software Product Line 2010, vol. Second Volume (Workshops), ACM. ACM International Conference Proceeding Series. Recently, there have been increasing demands for the postponement of decisions on software adaptations and product variations to provide the flexibility required by dynamic environments and users. The goal is that software adaptations and product variations can be chosen even at runtime. As such, a research theme that addresses development issues for reusable and dynamically reconfigurable core assets has emerged and it is called dynamic software product lines (DSPLs) with its consequential need to manage runtime variability. Research on the use of runtime variability, however, is still heavily based on the specification of decisions during design time. That is, a system simply postpones “when to adapt” to runtime but “how to adapt” is already decided at design time. In this paper a brief assessment of the current research in the area and discuss some research issues related to the feasibility of DSPL oriented approaches to build self-adaptive systems is reported. Brice Morin, Jacques Klein, Jörg Kienzle, and Jean-Marc Jézéquel (2010) **Flexible Model Element Introduction Policies for Aspect-Oriented Modeling** in: MODELS'10: 13th ACM/IEEE Conference on Model-Driven Engineering Languages and Systems, Oslo, Norway. Aspect-Oriented Modeling techniques make it possible to use model transformation to achieve advanced separation of concerns within models. Applying aspects that introduce model elements into a base model in the context of large, potentially composite models is nevertheless tricky: when a pointcut model matches several join points within the base model, it is not clear whether the introduced element should be instantiated once for each match, once within each composite, once for the whole model, or based on a more elaborate criteria. This paper argues that in order to enable a modeler to write semantically correct aspects for large, composite models, an aspect weaver must support a flexible instantiation policy for model element introduction. Example models highlighting the need for such a mechanism are shown, and details of how such policies can be implemented are presented. Brice Morin, Olivier Barais, Grégory Nain, and Jean-Marc Jézéquel (2009) **Taming Dynamically Adaptive Systems Using Models and Aspects** in: ICSE'09: 31st Internation Conference on Software Engineering, Vancouver, Canada. Since software systems need to be continuously available under varying conditions, their ability to evolve at runtime is increasingly seen as one key issue. Modern programming frameworks already provide support for dynamic adaptations. However the high-variability of features in Dynamic Adaptive Systems (DAS) introduces an explosion of possible runtime

	system configurations (often called modes) and mode transitions. Designing these configurations and their transitions is tedious and error-prone, making the system feature evolution difficult. While Aspect-Oriented Modeling (AOM) was introduced to improve the modularity of software, this paper presents how an AOM approach can be used to tame the combinatorial explosion of DAS modes. Using AOM techniques, we derive a wide range of modes by weaving aspects into an explicit model reflecting the runtime system. We use these generated modes to automatically adapt the system. We validate our approach on an adaptive middleware for home-automation currently deployed in Rennes metropolis.
FAST	**D2.2.2 Ontology and Conceptual Model for the Semantic Characterisation of Complex Gadgets** FAST is about the development of complex gadgets for end users, and so the conceptual model and ontology provide the terminology and conceptual framework that is needed for specifically this task. In other words, while this document provides the theoretical underpinning for the implementation-oriented work in the project, the implementation-oriented work on the other hand defines the scope and requirements for the model. This deliverable defines the abstract conceptual model and concrete ontology for the semantic characterisation of complex gadgets in the FAST project. The project applies the *Methontology* ontology development methodology. The design and development process is therefore defined by a number of different living documents, which function both as a development tool, as well as documentation for the ontology itself resulting in the final implementation document which formally defines the ontology in OWL DL semantics. Throughout the document, we provide extensive code examples and visual representations of the different aspects of both model and ontology. **D3.1.2 FAST complex gadget architecture** The documents defines a general architecture for the next generation Service Frontend Resources (SFR, aka. gadgets). A gadget (SFR) is a web application, with high user interaction capabilities, allowing non-technical users to interact with back-end services. A gadget should be seen as the face of service-oriented architecture (SOA). Gadgets must deal with aspects like Service Semantic homogenization, Semantic Piping, User Interface Pipe Handling, Screenflow Composition and finally Application Mashup Integration. All this aspects are treated in the FAST architecture, involving not only external aspects, specifying how gadgets will interact in a Mashup platform, but also the gadget's internal structure, in order to allow the composition of gadgets from their containing components. **D4.3.1 Mechanisms for Gadget-Service Connections and Gadget Functionality** This deliverable exposes several mechanisms which allows the connection and interaction between end-user's interfaces to third-party back-end services. These back-end services cannot be directly used; hence they need to be encapsulated or wrapped into the so-called Resource Adapters within the FAST platform. The application of semantics to the back-end, through the corresponding resource adapters, and front-end building blocks ensures a powerful instrument in the task of building new gadgets, improving the search, and enhancing the connection among the different building blocks which compose a gadget. Therefore, the focus of this deliverable is to define how these wrappers will be constructed to allow the FAST platform to exploit web services such as RESTful web services, SOAP-based web services and semantic web services through WSMO, and define mechanisms to connect them within the gadgets. Hoyer, V., Janner, T., Delchev, I., Fuchsloch, A., López, J., Ortega, S., Fernández, R., Möller, K.H., Rivera, I., Reyes, M.,

	Fradinho, M. (2009). **The FAST Platform: An Open and Semantically-enriched Platform for Designing Multi-Channel and Enterprise-Class Gadgets** in Proceedings of the International Conference on Service-Oriented Computing (ICSOC), Stockholm, Sweden The transfer of the mashup paradigm in corporate environments needs additional capabilities beyond those typically associated with consumer mashups. In this paper, we present the architecture of the FAST platform which allows creating enterprise-class and multi-channel visual building blocks (so called gadgets) in an ad-hoc manner. The design of complex enterprise-class gadgets is supported by an integrated semantic concept which hides the complexity from the actual users. The architectural components of the platform, a technical life cycle model for enterprise mashups, and the FAST gadget ontology are presented. By means of a cross-organisational real-world scenario from the marketing/ promotion event area, we demonstrate the value and potential of the FAST platform. Juan J. Hierro, Till Janner, David Lizcano, Marcos Reyes, Christoph Schroth, Javier Soriano (in alph. Order), **Enhancing User-Service Interaction Through a Global User-Centric Approach to SOA**. The Fourth International Conference on Networking and Services, ICNS'08. Gosier, Guadalupe, 16-21 March 2008. Published in IEEE Computer Society Press.ISBN 0-7695-3094-X Considered as enablers of seamless application-to application integration both within company boundaries and on a global scale, Web Services-based SOAs have traditionally focused on automating service-to-service collaboration. However, they have never featured a "face" to human users. This gap between human users and services still prevents enterprises from realizing how innovations at the SOA front-end help to make people more productive. And, ultimately, it hinders the emergence of a real Web of Services driven by a global, user-centric SOA. In this paper, the author revisit the notion of SOA and analyse its major shortcomings with regard to the emergence of a Web of Services enhancing user service interaction and increased service usability. The authors then elaborate on novel, currently emerging technologies that facilitate the establishment of the global mesh of interoperable user-centric services that makes up that Web of Services and shift the roles and characteristics of resource consumers, providers and intermediaries considerably. A novel platform architecture is presented that builds on all the key technical enablers and has the potential to drive the development of a global, user-centric SOA. David Lizcano, Javier Soriano, Marcos Reyes and Juan J. Hierro, EzWeb/FAST: **Reporting on a Successful Mashup-based Solution for Developing and Deploying Composite Applications in the Upcoming Web of Services**. ACM Proceedings of the 10th International Conference on Information Integration and Web-based Applications & Services, iiWAS 2008, Linz, Austria, November 24-26, 2008. ACM ISBN 978-1-60558-349-5 Service oriented architectures (SOAs) based on Web services have attracted a great interest and IT investments during the last years, principally in the context of business-to-business integration within corporate Intranets. However, they are nowadays evolving to break through enterprise boundaries, in a revolutionary attempt to make the approach pervasive, leading to what we call the ubiquitous SOA, i.e. a SOA conceived as a Web of services made up of compositional resources that empowers end-users to ubiquitously exploit these resources by collaboratively remixing them. In this paper we explore the architectural basis, technologies, frameworks and tools considered necessary to face this novel vision of SOA. David Lizcano, Miguel Jiménez, Javier Soriano, José M. Cantera, Marcos Reyes, Juan J. Hierro, Francisco Garijo and

	Nikolaos Tsouroulas, **Leveraging the Upcoming Internet of Services through an Open User-Service Front-end Framework**. In Petri Mähönen, Klaus Pohl, Thierry Priol (editors): Towards a Service-Based Internet, Springer Lecture Notes in Computer Science, nº 5377, ISSN 0302-9743, ISBN-10 3-540-89896-4, ISBN-13 978-3-540-89896-2, ServiceWave 2008 Conference, Madrid, Spain. The Internet of the Future is expected to be composed of a mesh of interoperable Web Services accessed from all over the Web. This approach has not yet caught on since a global user-service interaction is still an open issue. This paper states author position with regard to the next generation front-end technology for the Internet of the Future. This approach will enable the massive deployment of services over the Internet in a user-centric fashion. This paper advocates the full development of front-end technologies to bring services closer to users, empowering them anytime and anywhere. It also outlines all the main gaps and technological challenges that have to be addressed. Finally, a model and architectures are proposed for building these technologies into NESSI's Open Framework Reference Architecture, NEXOF-RA.
IRMOS	**D3.1.2 IRMOS Overall Architecture** The document describes the overall architecture of IRMOS platform as part of IRMOS project. The main objective of IRMOS project is to build a Service Oriented Infrastructure (SOI) for interactive applications with real-time requirements **D5.1.1 Models of Real time Applications on Service Oriented Infrastructures** The report considers the value chain for real time applications hosted by third party service providers. In the context of this value chain, is analysed who might benefit from the use of models, how and why these models might be used, and when during the application lifecycle modelling is most useful. Techniques are presented for building models of real time applications including the use of stochastic process algebras, finite state automata, workflow models (e.g. BPEL, BPMN and YAWL) and specification languages (e.g. UML MARTE). The report discusses how these models need to be supported by estimation of application resource consumption and how mapping techniques allow models to be built for different actors in the value chain. Tool support is also discussed e.g. PRISM for probabilistic model checking and Visual Service Composition Studio for service oriented modelling. A detailed and specific real time application scenario is included and modelled to allow the various techniques presented in this document to be demonstrated and quantitatively evaluated. The scenario also reveals the level of detail needed in order for meaningful modelling to be achieved in practice. The modelling scenario has been carefully engineered to be as representative as possible of the broad range of application characteristics encountered in the three reference IRMOS applications (film postproduction, eLearning, virtual and augmented reality). The modelling techniques discussed and then demonstrated in this report include identifying what resources are necessary to support an application, when those resources will be required during the application workflow, what performance is needed from them (i.e. QoS) and what will happen to the application if the required performance is not delivered. All of these are essential when developing and then agreeing service level agreements between the various entities in a service oriented infrastructure. Tommaso Cucinotta and Dario Faggioli Real-Time Systems Lab, CEIICP, Scuola Superiore Sant'Anna, Pisa (Italy), IEEE Symposium on Industrial Embedded Systems (SIES 2010), **An Exception Based Approach to Timing Constraints Violations in Real-Time and Multimedia Applications.** In this paper, an exception-based programming paradigm is envisioned to deal with timing constraints violations occurring in

	soft real-time and multimedia applications written in the C language. In order to prove viability of the approach, a mechanism allowing to use such paradigm has been designed and implemented as an open-source library of C macros making use of the standard POSIX API (a few Linux-specific optimizations are also briefly discussed). The envisioned approach has been validated by modifying mplayer, one of the most widely used multimedia player for Linux, so as to use the introduced library. Experimental results demonstrate how the exception-based paradigm is effective in improving the audio/video delay exhibited by the player. Tommaso Cucinotta Real-Time Systems Laboratory Scuola Superiore Sant'Anna, Pisa, Italy, Kleopatra Konstanteli and Theodora Varvarigou Advanced Distributed Computing Laboratory National Technical University of Athens, Greece, **Advance Reservations for Distributed Real-Time Workflows with Probabilistic Service Guarantees**, IEEE International Conference on Service-Oriented Computing and Applications (SOCA 2009), December 2009, Taipei, Taiwan, This paper addresses the problem of optimum allocation of distributed real-time workflows with probabilistic service guarantees over a Grid of physical resources made available by a provider. The discussion focuses on how such a problem may be mathematically formalised, both in terms of constraints and objective function to be optimized, which also accounts for possible business rules for regulating the deployment of the workflows. The presented formal problem constitutes a probabilistic admission control test that may be run by a provider in order to decide whether or not it is worth to admit new workflows into the system, and to decide what the optimum allocation of the workflow to the available resources is. Various options are presented which may be plugged into the formal problem description, depending on the specific needs of individual workflows.
m:Ciudad	**D2.1 State of the art in service description** this deliverable studies the state of the art in service description, focusing on microservices (U+) and user generated services. To find a suitable service description language, this deliverable explores existing standards and service-oriented solutions to guarantee a solid background for U+ service description. These existing solutions include different Web 2.0 components (WSMO, WSML, WSDL, SA-WSDL, UDDI, OWL-S, SWSF), the widget model, mashups and choreography mechanisms for U+ services. The survey studies how these existing service descriptions can be expanded or restricted to fit m:Ciudad service requirements, and identifies the main unresolved questions to obtain comprehensive, flexible and efficient service descriptions to serve as a base for automatic service implementation. **D1.4 m:Ciudad Reference Architecture** m:Ciudad is an ecosystem of millions of mobile servers providing user-created services to be used by end users. This unique environment defines the basic pillars that constitute the problem statement for an architecture for m:Ciudad. The m:Ciudad Reference Architecture must support the outsourcing of service provision to terminals owned by end users: specifically mobile phones. This is key because it implies that services are provided from intrinsically limited devices in a mobile context, and therefore the network has to provide a minimum set of supporting functionalities which tie the ecosystem together and cannot be relied on to individual servers. Another important implication is that the m:Ciudad architecture modules in the mobile server side have to take into account the limitations and specificities of mobile terminals. Give users full control over their services. Since services are provided from user-owned equipment, users must be able to control, at all times, which services they provide, how much resources they are willing to devote and which level of privacy or exposure they want to attach to their

	services. In addition, since m:Ciudad allows users to create services, service creators must retain adequate privileges over their created services. Successfully manage the complexity of an ecosystem of millions of potential service providers. Allowing every user to find the service s/he needs in the current context, without compromising the peer-to-peer nature of service provision, is one of the most critical challenges of an architecture for m:Ciudad. Ensuring that communication between consumers and providers is always possible in a mobile environment is also a problem derived from this complexity. All the proposed modules have been implemented and integrated creating a robust and complex ecosystem for the creation, delivery and execution of U+ Services.The technical implications of a p2p U+ services framework and the decisions taken have forced implementation architecture slightly different from the proposed reference. Any module following the specification included in this document could be easily integrated. The dynamic analysis allows to validate initial premises such as that execution processes are between mobiles devices with no intervention of network components, and that creation and execution process are completely independent one another, in favour of platform independence. The architecture is scalable and could be easily upgraded Ville Antila, Jani Mantyjarvi, "**Distributed RESTful Web Services for Mobile Person-to-Person Collaboration**," Next Generation Mobile Applications, Services and Technologies, International Conference on, pp. 119-124, 2009 Third International Conference on Next Generation Mobile Applications, Services and Technologies, 2009 The mobile devices have a key role in aggregating information while people are on the go. Therefore it would be beneficial to use that potential also for sharing information collaboratively. In this paper, one possible solution to this challenge is given. The proposed approach is to use distributed, lightweight Web services for mobile collaboration. The collaboration between the people is direct without the need for centralized information storage. Therefore the approach is described here as mobile person-to-person collaboration. In this paper this concept is described by providing a list of requirements and by presenting an architectural approach. To evaluate the approach, a prototype service was implemented as a case study. To discuss the focus and contribution of this work the architecture is evaluated against the requirements and compared with the relevant state-of-the-art studies in the field. The results indicate that this approach provides a prominent and novel way of leveraging on the existing content for collaboration in a mobile environment as well as refines the requirements for future mobile collaborative applications and services.
MANCOOSI	**D3.1 Survey of the state of the art technologies for** handling versioning, rollback and state snapshot in complex systems Managing the complexity of Open Source software components is the main goal of the Mancoosi project. As in other systems, like databases, transactional upgrades of software components have desirable characteristics. We should be able to revert to a previous working state of the system when a installation / removal has a non expected behaviour and impact. Moreover, this rollback may not only be performed after the problem has occurred, but also any time in the future given that we cannot predict when the error will be discovered. This could be accomplished by storing every change in the system, binaries, documents and configuration files. However, the cost of storing such an amount of information is not feasible in a system in permanent evolution. With the aim of developing and enhancing tools to keep track of system evolution, we proposed to start this workpackage by capturing the state of the art of existing techniques and tools. This deliverable analyses various techniques. The deliverable gives a brief introduction to the work developed and the topics addressed by giving an overview of the structure of the deliverable then describes the concepts of Software component management presenting a chapter with

	main lines of research in academic and scientific fields for component management in the last two decades. The state of the art is presented in three sections: run-time reconfiguration, dynamic adaptation and components selection. Then the concept of Linux packages which are a particular kind of components. The DEB and RPM file formats are also presented in detail. The state of the art of rollback components, and presents an overview of different components critical for setting up a solution of transactional upgrades is given. **D2.1 Metamodel for Describing System Structure and State** The deliverable proposes a model-driven approach to simulate system upgrades in advance and to detect predictable upgrade failures, possibly by notifying the user before the system is affected. The approach relies on an abstract representation of the systems and packages which are given in terms of models that are expressive enough to isolate inconsistent configurations (e.g., situations in which installed components rely on the presence of disappeared sub-components) that are currently not expressible as inter-package relationships. Today's software systems are very complex modular entities, made up of many interacting components that must be deployed and coexist in the same context. Modern operating systems provide the basic infrastructure for deploying and handling all the components that are used as the basic blocks for building more complex systems even though a generic and comprehensive support is far from being provided. In fact, in Free and Open Source Software (FOSS) systems, components evolve independently from each other and because of the huge amount of available components and their different project origins, it is not easy to manage the life cycle of a distribution. Users are in fact allowed to choose and install a wide variety of alternatives whose consistency cannot be checked a priori to their full extent. It is possible to easily make the system unusable by installing or removing some packages that \break" the consistency of what is installed in the system itself. **D5.2 Extension of an existing package manager to produce traces of upgradeability problems in CUDF format** One of the objectives of the Mancoosi project is to resolve some of the problems that users of Free and Open Source Software distributions experience when trying to install, remove, or upgrade packages installed on their machines. The specific goal is to build a database of problem reports generated from such user requests to a meta-installer, which then will be used by the Mancoosi project, and the research community in general, to develop better algorithms to compute upgrade paths. The design of the languages used to produce these error reports have been described in an earlier deliverable. The format of the report produced on a user machine is specific to the software distribution used, but follows a general project-wide scheme called DUDF (Distribution Upgradeability Description Format). Problem reports are uploaded to a server specific to the software distribution and translated by the distribution editor into a common format called CUDF. The Common Upgradeability Description Format (CUDF) is a format for describing upgrade problems independently of a specific FOSS distribution. Reports in CUDF format are then transferred from the distribution editor's server to a central server of the Mancoosi project, where they will be used in the construction of a project-wide problem database. This document summarizes the work done by three different GNU/Linux distributions Debian, Mandriva and Caixa Magica to generate problem reports on user machines in DUDF, transfer them to the distribution editors, and to convert them there into the common format CUDF **DiRec: Diversified Recommendations for Semantic-less Collaborative Filtering**, ICDE 2011, April 11-16, 2011, Hannover, Germany. By Ruby Boim, Tova Milo, and Slava Novgorodov DiRec , a plug-in that allows Collaborative Filtering (CF) Recommender systems to diversify the recommendations that they

	present to users. Di- Rec estimates items diversity by comparing the rankings that different users gave to the items, thereby enabling diversification even in common scenarios where no semantic information on the items is available. Items are clustered based on a novel notion of priority-medoids that provides a natural balance between the need to present highly ranked items vs. highly diverse ones. We demonstrate the operation of DiRec in the context of a movie recommendation system. We show the advantage of recommendation diversification and its feasibility even in the absence of semantic information. **Autocompletion for Mashups**, VLDB 2009, August 24-28, 2009, Lyon, France. By Ohad Greenshpan, Tova Milo and Neoklis Polyzotis. A mashup is a Web application that integrates data, computation and UI elements provided by several components into a single tool. The concept originated from the understanding that there is an increasing number of applications available on the Web and a growing need to combine them in order to meet user requirements. This paper presents MatchUp, a system that supports rapid, on-demand, intuitive development of mashups, based on a novel autocompletion mechanism. The key observation guiding the development of MatchUp is that mashups developed by different users typically share common characteristics; they use similar classes of mashup components and glue them together in a similar manner. MatchUp exploits these similarities to recommend useful completions (missing components and connections between them) for a user's partial mashup specification. The user is presented with a ranking of the recommendations from which she can choose and refine according to her needs. This paper presents the data model and ranking metric underlying a novel autocompletion mecha nism. It introduces an efficient top-k ranking algorithm that is at the core of the MatchUp system and that is formally proved to be optimal in some natural sense. We also experimentally demonstrate the efficiency of our algorithm and the effectiveness of our proposal for rapid mashup construction. **Towards a Model Driven Approach to Upgrade Complex Software Systems**, ENASE 2009, May 9-10 2009, Milan, Italy. By Antonio Cicchetti, Davide Di Ruscio, Patrizio Pelliccione, Alfonso Pierantonio, and Stefano Zacchiroli. Complex software systems are more and more based on the abstraction of package, brought to popularity by Free and Open Source Software (FOSS) distributions. While helpful as an encapsulation layer, packages do not solve all problems of deployment and management of large software collections. In particular upgrades, which often affect several packages at once due to inter-package dependencies, often fail and do not hold good transactional properties. This paper shows how to apply model driven techniques to describe and manage software upgrades of FOSS distributions. It is discussed how to model static and dynamic aspects of package upgrades—the latter being the most challenging aspect to deal with—in order to be able to predict common causes of upgrade failures and undo residual effects of failed or undesired upgrades.
MOST	**D5.1 Definition of the case study requirements** In this document, an analysis of the existing Comarch software product line infrastructure and software development process is provided. Furthermore, 5 concrete scenarios where MOST technology may prove to be useful are provided. Within each scenario, specific use cases for MOST technology are identified and define the requirements for tools that will be created in the scope of the project. **D.2.5.1 Ontology Services for Model-Driven Software Development**

	Today's model-driven software development (MDSD) approaches allow for a more abstract and productive way of developing software systems. However, they still suffer from various challenges. This document reports on the key MDSD challenges addressed in the MOST project. It provides concrete examples of these key challenges in our case studies and gives a service-based description of the capabilities of ontology technology used to address the identified challenges. To document and discuss the differences ontology technology makes for MDSD, presents the vision for Ontology-Driven Software Development (ODSD) that contributes (1) technical solutions to bridge the gap between the modelling and ontology technical spaces, (2) an approach for ontology-aware metamodelling and Language Engineering, (3) an approach for ontology aware modelling and system engineering, and (4) an ontology-aware MDSD process. Finally it presents a generic architecture for ontology-aware MDSD tools that apply ontology services to realise the ODSD vision. **D3.6 Successful Stories and Potential Patterns on Applying Ontologies in Software Engineering** This deliverable aims for an investigation of some successful stories on applying Ontologies and Semantic Web technologies in Software Engineering. The deliverable contains categorization of potential patterns of using ontologies for applications in software technology. **D7.1 Market analysis and preparation of marketing and standardization materials** This report provides and introduction of the main elements captured during the market analysis, in particular a formalization of the value chain, identification of stakeholders, market trends and hype cycle. In the report we present also the proposed business model for MOST using the semiformal representation model as well as a SWOT analysis deriving envisioned strategies for compensating treats and weakness. The second part of the report is dedicated to recommendations for preparation of marketing and standardization material, analysing the mix of variables that directly can influence the success of MOST technology and ensuring that innovations and their way to the market. Gerd Gröoner, Fernando silva Parreiras, and Stefen Staab. **Semantic Recognition of Ontology Refactoring**. In P.F.Patel-Schneider et al., editor, Proc. of the 9th Interational Semantic Web Conference (ISWC), volume 6496 of LNCS, pages 273-288. Springer, 2010. Ontologies are used for sharing information and are often collaboratively developed. They are adapted for different applications and domains resulting in multiple versions of an ontology that are caused by changes and refactorings. Quite often, ontology versions (or parts of them) are syntactical very different but semantically equivalent. While there is existing work on detecting syntactical and structural changes in ontologies, there is still a need in analysing and recognizing ontology changes and refactorings by a semantically comparison of ontology versions. In our approach, we start with a classi_cation of model refactorings found in software engineering for identifying such refactorings in OWL ontologies using DL reasoning to recognise these refactorings. Jens Lemcke, Tirdad Rahmani, and Andreas Friesen. **Semantic business process engineering**. In Uwe Aßmann, Andreas Bartho, and Christian Wende, editors, Reasoning Web, volume 6325 of Lecture Notes in Computer Science, pages 161-181.

	Springer, 2010. In this tutorial, we compare OWL-DL reasoning and Petri net analysis for validating refinement and grounding of business processes. (1) Process refinement: Like in software engineering, the implementation of a business process involves different interacting roles, such as business expert, analyst, process architect, and developer. Each role designs and refines different abstractions of the process until it is sufficiently refined. It is important to verify that the process models of the different abstractions are consistent. (2) Process grounding: A sufficiently refined process has to be mapped on existing IT systems. Ideally, IT systems consist of components with a semantic annotation of their behaviour. The most specific process must respect all IT systems' behaviours. Formally capturing process semantics enables to check automatically for consistent process refinement and grounding. The classic application of semantic techniques in the area of static models is well understood. The analysis of business processes deals with dynamics. Modelling dynamics is a challenge for current approaches of semantic Web services. We compare advantages and shortcomings of Petri net analysis and description logic (DL) reasoning for refinement and grounding validation. Steffen Staab, Tobias Walter, Gerd Grner, and Fernando Silva Parreiras. **Model driven engineering with ontology technologies. In Reasoning Web**. Semantic Technologies for Software Engineering, volume 6325 of LNCS. Springer, 2010 Ontologies constitute formal models of some aspect of the world that may be used for drawing interesting logical conclusions even for large models. Software models capture relevant characteristics of a software artifact to be developed, yet, most often these software models have limited formal semantics, or the underlying (often graphical) software language varies from case to case in a way that makes it hard if not impossible to fix its semantics. In this contribution, we survey the use of ontology technologies for software modeling in order to carry over advantages from ontology technologies to the software-modelling domain. It will turn out that ontology-based metamodels constitute a core means for exploiting expressive ontology reasoning in the software modeling domain while remaining flexible enough to accommodate varying needs of software modellers.
NEXOF - RA	**D7.5 RA Specification 1.0** This document represents the entry point to the set of documents that the deliverable D7.5b is constituted of. The complete D7.5b deliverable is composed of all the following documents: RA Specification – Executive Summary (this document); RA Specification – Pattern Compass; a set of "RA Specification – Pattern" documents. The overall deliverable presents the final version of the NEXOF-RA Specification that is the final version of the related set of architectural patterns specifically built to support the design of NEXOF Compliant (SOA) Infrastructures. In particular, the "RA Specification – Pattern Compass" provides a compass to access all the set of patterns. A part from providing the complete list of patterns that are into the specification, it shows the main relationships among them. This is very helpful to guide a software architect to access, find and usage patterns of the specification in a consistent and effective manner. The compass also provides a short abstract description of each pattern, and some other few information concerning the authors, its category type and level, and a reference to the document where the complete description is presented **D7.2c Definition of an architectural framework & principles** This deliverable is dedicated to lay the principles and the baseline for the creation of *NEXOF-RA Specification*. It mainly fixes rules and restrictions for the formal aspects of the specification, such as the format, the structure and its development

	approach. As far as it concerns architectural solutions, this document does not state any restricted principle, since it is well-known that NEXOF Reference Architecture is domain and technology independent. A part from the restriction on the very general domain of SOA Infrastructures, any concern, problem and solution related to this domain generally is interesting for the project. In the first section this document provides a description of what the NEXOF Reference Architecture is, what it is useful for and what its overall structure is. This section if fundamental to understand all the principles and baseline introduced afterwards. The second section gives a list of principles used for the development of the specification. They have been selected as guidelines for the specification process in order to produce an open and easily evolvable specification. The third section is dedicated to the introduction of the idea of *constructional patterns* as the basic mechanism (baseline) that is used to develop the overall specification. **D2.1b Service Centric System Architecture Contributions to Model and Architecture** The NEXOF-RA work package "Service-Centric Systems Engineering" focuses on the following areas of service based software (SW) systems: specification, discovery, design and composition of services. The goal of the work package "Service-Centric Systems Engineering" is to contribute on these areas to the NEXOF reference architecture specification. The scope of this document is to summarize the activities and the results of the work package "Service-Centric System Engineering". In particular, it describes the contributions to the reference architecture and the overall process that has been followed to contribute to the model and to the specifications. It also describes the most important results that have been achieved. Elaborating the inputs of the external contributors according to the principles, the guidelines and the templates of the project has produced such results. This document contains: a summary of the process that have been followed to obtain the contributions (see section Contribution from the open process); a summary of the achieved results (see section Contribution to the Reference Architecture and Contribution to the PoCs); The aim of NEXOF is to deliver a reference architecture for the NESSI Open Service Framework. In the NEXOF-RA work packages "Reference Architecture: Model" and "Reference Architecture: Specification" nine concerns1 turned out to be a guideline for the whole project. The focus of the work package "Service-Centric System Engineering" has been on the following four concerns: Services, Messaging, Discovery and Composition. Investigations teams (ITs) on the following five different topics (relate to the above 4 concerns) were managed: Service description (see 3.1.2.1), Design time service composition (see 3.1.2.2), Service discovery (see 3.1.2.3), Interoperability of message-based service invocation (3.1.2.4), Service runtime composition (see 3.1.2.5) and contributions from external contributors (including NESSI strategic projects members) were gathered2. The most interesting and mature contributions have been selected (also according to the requirements gathered by the "Requirements and Assessment Criteria" work package), elaborated and made compliant to the NEXOF-RA rules and templates. They have been used to contribute to the NEXOF-RA reference architecture, in particular to the model and to the specifications. The contributions to the Reference Model has been done in terms of: Contribution to the glossary Contribution to the conceptual model by introducing functionality of core service area Service, Message, Discovery, Composition. In particular an important result comes from the Service Description IT. The target of such IT was to provide an answer to the question: "what is a service?". In order to answer this question, a deep analysis about the service

	characterization was needed to avoid ambiguity, vagueness and, more in general, to provide a solution that try to fill the gap bared by existing standards. The result obtained contributes to the NEXOF-RA conceptual model and provides a reference for all the decisions concerning architectural choices of NEXOF Compliant Platforms. The most important contribution to the Reference specification has been done in terms of Architectural patterns: they have been developed by experts and architects of NEXOF-RA project and by external contributors by the open process. The work package “Service-Centric System Engineering” has produced: 5 patterns for Enterprise SOA (ESOA) , 5 patterns for Internet of Service (IoS) 6 patterns for both ESOA and IoS Moreover: 6 additional patterns for ESOA domain has been discussed and initially described in the context of the ITs. Consider that: 17 patterns (of a total of 22) have been produced integrating the results of four ITs (Design time service composition (see 3.1.2.2), Service discovery, Interoperability of message-based service invocation, Service runtime composition The achieved results have been described in details in section Contribution to the Reference Architecture. D6.3 **The NEXOF Reference Model V3.0** A reference architecture subsumes reoccurring and well-proven concepts and patterns of a set of specific software architectures [5]. The main goal of the NEXOF-RA project is to provide a reference architecture for service-based software systems which facilitates the reuse of well-proven service-oriented concepts. The NEXOF Reference Architecture is provided in form of a construction kit that guides the construction of specific service-based infrastructures. The construction kit consists of a set of building blocks implementing architectural patterns which in turn are related to a conceptual model. The NEXOF Reference Model captures the relevant entities and dependencies among them that constitute a service-oriented system on a conceptual level in order to foster the communication about the relevant elements on a higher abstraction level. This document is the final version of the NEXOF Reference Model specification and thus completely replaces the earlier version of the reference model. It is organised in three parts. The first part focuses on defining the syntax and semantics of the model. This includes describing the goals, scope and boundaries of the model and the terminology used within the model as well as across all NEXOF-RA documents. One section is dedicated to the introduction of the structure of the model. It is explained which views, concepts and diagram types are used for the different purposes of the model following the idea of separating structure and behaviour/functionality. Thus, the first part provides a guideline on how to read, use and interpret the model specification by explaining what the syntax and semantics of the different elements constituting the NEXOF Reference Model are. The actual specification of the reference model constitutes the second part of the document. Since the NEXOF Reference Architecture is structured into nine concerns, the specification of the model is organised according to those concerns. It defines the functionalities provided by a NEXOF compliant architecture qualifying the value of these functionalities by specifying the input-output relations along with the actors involved in these functionalities. Each diagram as well as all contained concept constituting the reference model are furthermore explained by textual descriptions and examples. The third part of the document focuses on the application of the NEXOF Reference Model in order to demonstrate how it can be used in practice. For this it is explained how the reference model is embedded into the overall NEXOF Reference Architecture with emphasis on the dependencies with other elements. First of all, the relations with the business- and system

	requirements are described, which form the baseline for the specified functionalities, and then with the patterns, which provide different ways of implementing the described functionalities. Furthermore, it is partially illustrated how the model can help in the derivation of actual NEXOF compliant architecture instances by using an example harbour scenario. The last section of the document draws a conclusion over the achievements of the NEXOF Reference Model during the course of the project. Finally, an evaluation of the reference model and a detailed analysis of potential points for further work in the continuation of building the NESSI Open Service Framework are given. While the three main parts focus on the essential information about the model specification, this document is also accompanied by a set of appendices, which give more technical details of specific topics. This includes, among other things, a list of updated glossary terms in comparison to the last version, the complete set of functionality descriptions enriched with explicit dependencies towards the system requirements and references to the sources of the functionalities as well as the complete definition of the service term as provided by the Service Description Investigation Team.
OMP	**D4.4 OMP Prototype Evaluation report** The report presents a detailed description of the OMP prototypes that have been developed by relevant beneficiaries during the Work Package 4 timeframe. Particular emphasis is put on the integration efforts for the different software tools and technologies, highlighting the challenges that have been encountered and detailing how they have been overcome. Further to this, actual results on OMP performance metrics are given for both OMP tools and OMP software stacks in terms of software productivity figures and runtime overhead assessments with respect to conventional existing tools and software frameworks.
OPEN	**D2.2 Document about Architecture for migratory user interfaces** This deliverable describes the software architecture supporting the migration of the user interface of an interactive application in order to allow users to continue their activities across various devices with different interaction resources. A corresponding prototype is under development. The first version will address migration of Web applications among desktop and mobile systems **D4.1 Solutions for Application Logic Reconfiguration** This document introduces aspects and solutions for the reconfiguration and adaptation of services during their lifecycle. During the migration of one or more services, their look and behavior has to be adapted to the target device. Context information is for example one important aspect, which may influence how to adapt services, like residence of the service user or current battery power and CPU frequency of the devices in use. In this document we will introduce an application scenario which demonstrates various kinds of adaptation and their triggers and according parameters which influence the way services are adapted. Furthermore we will present current solutions for application logic reconfiguration. Finally we will show architectural solutions for the different reconfiguration approaches. **D6.2 Evaluation parameters for enabling the environment programmability** The aim of the document is describing the approach that will be implemented for the programmability evaluation of the OPEN environment. For the scope of this document, the programmability is the capability of defining different rules that describe the migration and adaptation processes depending on the context information. Since the OPEN platform architecture is still an High Level design, the following analysis should be seen as a guidelines collection, to be further evaluated when the

	architecture will be refined. The document is structured as follows: in the second chapter a brief definition of programmability applied to the OPEN migration and adaptation processes is introduced.
Persists	**D2.3 Federation and Deployment Management Report** The process of designing an architecture for a distributed adaptive system requires input from many domains. The domains of federation and deployment are important examples of this. By "federation", we mean the act of joining a group, after the negotiation and the acceptance of some or all of its rules. By "deployment", we mean the act of putting some artefacts in place, defining where and how their exposed features might be accessed and used. A system architecture design for personal self-improving smart spaces highly depends on a comprehensive investigation of all aspects related to federation and deployment, as these two considerations will play an important role in the specification of an adaptive user-centric smart space environment. In order to achieve dynamic federation and seamless deployment, this deliverable will review the current state-of-the-art approaches to federation and deployment, highlighting those that better fit into the heavily distributed and dynamic scenarios of PERSIST. After outlining the available techniques and the best practices, the PERSIST scenarios are analysed, examining the aspects that directly relate to federation and deployment for each scenario. The deliverable concludes with a set of architectural recommendations for federation and deployment. Choosing a "pure" peer-to-peer approach for communication within a federated network is a good architectural choice for PERSIST, as it avoids single point of failure. Federated identity management requires a complex set of technologies and business processes, but the goal behind it is simple: to automatically share identity information across administrative boundaries. Federated identity management technologies improve security by controlling access on an operation-by-operation basis and providing a detailed audit trail. This added security and accountability is especially important for unattended machine-to-machine transactions, where no humans oversee transactions or assume liability, which will be the ultimate outcome of the pro-activity and reasoning components of PERSIST. Deployment of services in a highly dynamic environment as the one depicted by PERSIST requirements is a difficult feature to satisfy. The set of available services is likely to evolve continuously as devices join or leave a Personal Smart Space (PSS). The design of physical deployment architectures involves identifying the device options available and the capability requirements for devices. The analysis of the selected PERSIST scenarios shows that a variety of devices and capabilities are required for different scenarios, therefore the PSS architecture cannot be restricted to a particular device range or supporting technology. The PSS architecture should be device neutral, but it must take into account the existing differences between mobile devices and high capability fixed devices, and attempt to support this. The architecture could also make the assumption that mobile devices of the future will have increased capabilities and then the differences will become negligible. **D2.5 Revised Architecture Design** This deliverable describes the revised and final high-level design of the PSS architecture, which will be used to guide the final PSS prototypes and demonstration applications. Several revisions have been performed on the initial PSS architecture design. This deliverable begins by elaborating on the results of the high-level architecture revision process. The 5- layer PSS architectural approach has not changed, however the need for communication between nonadjacent layers caused a change in the form and accessibility of some layers. Subsequently, the revised functional view of the PSS

	architecture is presented. The following component blocks have been updated: Context Management, PSS Management, Service Management, Proactivity, Decision Support, Security and Privacy Management, Service Run-Time Environment and Overlay Network Management. The revised deployment view of the PSS architecture is provided, where the final mapping of components to the three PSS distributions is described. Then, the document outlines the revised structural view of the PSS architecture, elaborating on the scope and naming of the PSS APIs, on the revised PSS device roles and on how the architecture supports partitioning of a PSS. Finally, the revised PSS data view is presented, with only minor revisions included. Nick Taylor, "**Personal eSpace and Personal Smart Spaces**," Self-Adaptive and Self-Organizing Systems Workshops, IEEE International Conference on, pp. 156-161, 2008 Second IEEE International Conference on Self-Adaptive and Self-Organizing Systems Workshops, 2008 This paper takes the physical notion of Personal Space as the basis for a vision of an analogous logical notion of a Personal eSpace. Convergence between mobile and internet technologies and the future enhancement of both will result in a world which is so complex, so open to abuse and so alien to our common sense reasoning that automated protection mechanisms based on understandable metaphors will be required. The Future Internet and Pervasive Systems are the stage upon which the ideas in this paper are developed. Challenges to personal security in the electronic world, both current and future, are identified. The vision of a Personal Space is presented along with some initial thoughts on how it might be instantiated using intelligent agents and a calculus of trust. The paper concludes with a brief description of the Personal Smart Space paradigm a form of Personal eSpace focusing on smart spaces which is being developed by the EUFP7 project Persist. Elizabeth Papadopoulou, Sarah McBurney, Nick Taylor, and M. Howard Williams. 2008. **Linking Privacy and User Preferences in the Identity Management for a Pervasive System**. In Proceedings of the 2008 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology - Volume 01 (WI-IAT '08) Two important concepts in developing ubiquitous or pervasive computing technologies that are acceptable to the end user are personalisation and privacy. On the one hand it is essential to take account of user needs and preferences to personalise decision making within such a system, on the other hand it is equally important to protect user privacy. One approach to handling user privacy is through the use of virtual identities. This has the advantage that it can also benefit the handling of user preferences. In particular, virtual identities can be used as a substitute for roles. On the other hand user preferences can be used in identity management to assist in selecting a virtual identity to hide the real identity of the user, thereby improving user-friendliness of the system. This paper describes this symbiosis and how it is implemented in the Daidalos pervasive system. Korbinian Frank, Nikos Kalatzis, Ioanna Roussaki, and Nicolas Liampotis. 2009. **Challenges for context management systems imposed by context inference**. In Proceedings of the 6th international workshop on Managing ubiquitous communications and services (MUCS '09) This work gives an overview over the challenges for context management systems in Ubiquitous Computing frameworks or Personal Smart Spaces. Focused on the integration of context inference in today's context management systems (CMSs) we address important design decisions for future frameworks. The inference system we have in mind is probabilistic and relies on the concept of Bayeslets, special inference rules extending Bayesian networks. We show that for inference rule creation,

	storage, inference scheduling and update frequency the best solutions are hybrid, allowing for high flexibility and performance while reducing resource costs. We also see that human expert knowledge cannot be substituted completely in an efficient context-aware system.
ProTest	*No Deliverable available on the website* Huiqing Li School of Computing, University of Kent, UK, Simon Thompson School of Computing, University of Kent, UK, **Testing-framework-aware Refactoring** Testing is the predominant way of establishing evidence that a program meets it requirements. When both test code and the application under test are written in the same programming language, a refactoring tool for this language should be able to refactor both application code and testing code together. However, testing frameworks normally come with particular programming idioms, such as their use of naming conventions, coding patterns, meta-programming techniques and the like. A refactoring tool needs to be aware of those programming idioms in order to refactor test code properly. Meanwhile the particularities of test code also suggest refactorings that are particularly applicable to test code. In this paper we present our experience of extending Wrangler, a refactoring tool for the Erlang programming language, so as to handle the three common testing frameworks for Erlang, as well as discussing the refactoring of test code in its own right. Kirill Bogdanov and Neil Walkinshaw. 2009**. Computing the Structural Difference between State-Based Models**. In Proceedings of the 2009 16th Working Conference on Reverse Engineering (WCRE '09). IEEE Computer Society Software behavior models play an important role in software development. They can be manually generated to specify the intended behavior of a system, or they can be reverse-engineered to capture the actual behavior of the system. Models may differ when they correspond to different versions of the system, or they may contain faults or inaccuracies. In these circumstances, it is important to be able to concisely capture the differences between models, a task that becomes increasingly challenging with complex models. This paper presents the PLTSDiff algorithm that addresses this problem. Given two state machines, the algorithm can identify which states and transitions are different. This can be used to generate a `patch' with differences or to evaluate the extent of the differences between the machines. The paper also shows how the Precision and Recall measure can be adapted to quantify the similarity of two state machines.

Q-ImPrEss	**D2.1: Service Architecture Meta-Model** This document represents deliverable D2.1 of the Q-ImPrESS description of work. It contains the Q-ImPrESS Service Architecture Meta-Model. The Service Architecture Meta-Model specifies how to describe service-oriented architectures in a way that allows latter quality trade-off analyses for evolving software systems. As such, it provides the schema used by the users of the Q-ImPrESS method to describe architectures. For example, the Enterprise SOA showcase of Itemis or the robot control system of ABB should be describable as instance of the Service Architecture Meta-Model. Additionally, the Service Architecture Meta-Model serves as foundation for all academic activities in order to define the Q-ImPrESS method and develop its tools. It defines the storage layout for reverse engineered source code, is the input needed to generate editors and transformations dealing with model instances and finally, allows trade-off analyses to be performed. **D5.2: Experiments with Impact Analysis** This document describes experimental evaluation of the method for checking consistency between service implementation in Java and its behaviour model in TBP. All experiments are performed on the CoCoME application. **D5.3: Trade-off Analysis of QoS Attributes** This document introduces and describes the methods and tools for performing the trade-off analysis of quality attributes in the Q-ImPrESS project.) **D3.3: Resource Usage Modeling** The Q-ImPrESS project deals with modeling of quality attributes in service-oriented architectures, which generally consist of interacting components that share resources. This report analyses the degree to which resource sharing of various omnipresent implicitly shared resources (e.g. memory content caches, memory buses, etc.) affect various quality attributes. The main goal is to identify the resources whose sharing affects the quality attributes significantly, and next propose methods for modeling of these effects **D7.1: Demonstrator Description** This document describes the design and implementation of the industrial demonstrators for the Q-ImPrESS project. For each demonstrator, it includes information about the domain background, an architectural description of the system, and detailed evolution scenarios. The document also sketches the planned application of the Q-ImPrESS method and tools on the demonstrators.
RESERVOIR	**High Level Architectural Specification - Release 2.0** The prime deliverable of the project is the architecture – presented at length in this document – and a reference implementation of a service-oriented infrastructure which, building on open standards and new technologies will provide a scalable, flexible and dependable framework for delivering services as utilities. This document focuses on identifying requirements for a highly dynamic computing cloud and defining a high level architecture for a system that can satisfy those requirements. The goal of this document is to setup the foundation on which the components that make the RESERVOIR compute cloud will be developed. However, this is not intended to be a static document; instead it is a living document that

	will be refined as the project progresses. In this second revision of the document, emphasis is put on a rigorous requirements collection process, and expand on the “vertical functionality” of the system: whereas in the first version of the document a layered architecture with clear separation of abstractions and responsibilities between the different horizontal layers was proposed. In this version, important functions that do not belong solely to any particular layer but need to be dealt at all levels have been identified (hence labelling them “vertical functions”). **D2.2.1 Virtual Execution Environment (VEEH) Design and Open Specification** This document is the design and open specification delivery for RESERVOIR Virtual Execution Environment Infrastructure Enablement sub-activities, Virtual Machine Technology, Relocation Enablement. The document outlines the architecture and high-level design of the VEE Host (VEEH). VEEH is the lowest layer in the RESERVOIR’s layered architecture. VEEH interacts with virtualization technology, e.g. KVM, XEN, etc.. It provides an abstraction upon which the upper layers of the architecture can build. The VEEH Interface (VHI) provides a generic interface to interact with the VEEHs on a RESERVOIR site **D2.3.1 Service Manager Design and Open Specification** The Virtual Execution Host (VEEH) is responsible for control and monitoring of Virtual Execution Environments (VEEs) and their resources. It interacts with the different physical VEE Hosts in the site to do low-level VEE operations required for the service-level management (e.g. create VEE, allocate additional resource to VEE, monitor VEE, migrate VEE, create virtual networks and virtual disks, etc.). Its functionality is exposed to the VEEM through the VEE Host Interface (VHI). This document is the Scientific Report corresponding to the analysis activities performed during the first year of RESERVOIR project within the Service Manager activity of the RESERVOIR project. Celesti, A. Tusa, F. Villari, M. Puliafito, A. Fac. di Ing., Univ. degli Studi di Messina, Messina, Italy “**How to Enhance Cloud Architectures to Enable Cross-Federation**” in: IEEE 3rd International Conference on Cloud Computing (CLOUD), 2010 The near future evolution of the cloud computing can be hypothesized in three subsequent stages: stage 1 "Monolithic" (now), cloud services are based on independent proprietary architectures; stage 2 "Vertical Supply Chain", cloud providers will leverage cloud services from other providers; stage 3 "Horizontal Federation", smaller, medium, and large cloud providers will federate themselves to gain economies of scale and an enlargement of their capabilities. Currently, the major clouds are planning the transition to the stage 2, but how to achieve the stage 3 is unclear because some architectural limitations have to be overcome. In this paper, considering a general cloud architecture, we highlight such limitations and propose some enhancements which add new federation capabilities. In order to address such concerns we propose a solution based on the Cross-Cloud Federation Manager, a new component placeable inside the cloud architectures, allowing a cloud to establish the federation with other clouds according to a three-phase model: discovery, match-making and authentication. Clovis Chapman, Wolfgang Emmerich, Fermn Gal\;n Marquez, Stuart Clayman, and Alex Galis. 2010. **Software architecture definition for on-demand cloud provisioning**. In Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing (HPDC '10). ACM Cloud computing is a promising paradigm for the provisioning of IT services. Cloud computing infrastructures, such as those offered by the RESERVOIR project, aim to facilitate the deployment, management and execution of services across multiple

	physical locations in a seamless manner. In order for service providers to meet their quality of service objectives, it is important to examine how software architectures can be described to take full advantage of the capabilities introduced by such platforms. When dealing with software systems involving numerous loosely coupled components, architectural constraints need to be made explicit to ensure continuous operation when allocating and migrating services from one host in the Cloud to another. In addition, the need for optimising resources and minimising over-provisioning requires service providers to control the dynamic adjustment of capacity throughout the entire service lifecycle. We discuss the implications for software architecture definitions of distributed applications that are to be deployed on Clouds. In particular, we identify novel primitives to support service elasticity, co-location and other requirements, propose language abstractions for these primitives and define their behavioural semantics precisely by establishing constraints on the relationship between architecture definitions and Cloud management infrastructures using a model denotational approach in order to derive appropriate service management cycles. Borja Sotomayor, Ruben S. Montero, Ignacio M. Llorente, Ian Foster, "**Virtual Infrastructure Management in Private and Hybrid Clouds**," IEEE Internet Computing, pp. 14-22, September/October, 2009 One of the many definitions of "cloud" is that of an infrastructure-as-a-service (IaaS) system, in which IT infrastructure is deployed in a provider's data center as virtual machines. With IaaS clouds' growing popularity, tools and technologies are emerging that can transform an organisation's existing infrastructure into a private or hybrid cloud. OpenNebula is an open source, virtual infrastructure manager that deploys virtualized services on both a local pool of resources and external IaaS clouds. Haizea, a resource lease manager, can act as a scheduling back end for OpenNebula, providing features not found in other cloud software or virtualization-based data centre management software.
Romulus	**D2.5 – Open Source Community set up** The purpose of the deliverable D2.5 is to present the Romulus Open source community set up. Romulus' strategy for creating a community has been based on taking advantage of the existing open source projects, contribute to them, and provide collaborative tools for the new projects incubated within project. The new developments of Romulus are delivered as open source projects incubated in Romulus. Thus, Romulus is a community composed of several open source projects focused on Java Web development, whose goal is providing a single access point to Java Web developers, and promote the collaboration between the projects developed within the Romulus project, in order to provide the basis for its collaboration during the project and after its conclusion. **D3.6 – Romulus RAD final** an extension for both Eclipse and NetBeans IDEs has been designed and implemented in order to provide a graphical interface to the Roma Framework [Roma]. Roma is a meta framework that allows rapid development of Web applications. Users that will mainly interact with such a framework could belong to one of the following categories:• Software Developers Software Architects. In order to support their work, many Integrated Development Environments (IDE) are available on the market, both commercial and open source. Since Romulus is completely based on open source software, only open source products have been evaluated for developing Roma graphical support. In this particular field, Eclipse represents a milestone in the market, being the first product that could be considered at enterprise level that was released with an open source

	license. Eclipse Platform has a vast community, a strong development process and is supported by thousands of programmers. NetBeans is the direct competitor of Eclipse and it is open source as well, supported by a big community of programmers. The choice to support both platforms, considering also their capillary distribution, has been taken to cover almost every taste of programmers who want to develop Web application using Roma Framework. **D4.2 - Final report on mashup integration** The approach to Enterprise Integration Mash-Up described in the deliverable "D4.1 Annual Report on mashup integration" and resumed in this section has been fully realized during the project. Using the Romulus components, users are able to create integration services by "mashing-up" existing enterprise services available in the network. The created services are standard, since they are based on BPEL, and technology independent, since they run inside an ESB. The ESB, in fact, plays the role of "technology homogenizer", because through its integration components, it is able to translate most of the protocols currently used at enterprise level into a normalized, technology independent, format. In addition, the ESB plays a very important architectural role in the integration: it represents the "mediator" among the service consumers and the enterprise service providers, allowing the implementation of different integration solutions in a transparent way. The created enterprise integration mash-up services, running inside the ESB, can be invoked by consumers using every supported technology, are described in a standard way by using a WSDL, and are available for other integration mashups once they are available in a public service Registry Repository. The Romulus approach to enterprise integration service mash-up is based on the Roma Framework and on the SOA open source stack provided by Sun that is composed by: OpenESB, GlassFish Application Server and NetBeans IDE. The approach provides all the features and tools necessary to help the developer, involved in complex integration projects, in creating added value starting from the existing information system capabilities. The Roma Framework has been extended by defining the Registry Aspect and the Enterprise Aspect. The first one allows developer to register a web service that is created by using Roma in a remote Service Registry Repository, simply annotating its implementation class. This feature allows the availability and the publication of services in the network, in order to provide their utilization for mash-ups, but also for improving service governance. The second one, instead, provides features for exposing enterprise services inside an ESB, introducing, in a really simple way, this architectural component in it as part of a Roma application. It allows also the creation of BPEL integration processes starting from a web service, simply annotating its implementation class. In this way the developer can concentrate the effort only in the logic, since the framework hides all the technical details. In order to help developers in creating mash-ups, Romulus provides also the WSDL Navigator. It is an extension to the NetBeans IDE that helps the developer in showing the descriptions of every enterprise service present and available in a public registry repository previously configured. The developer, therefore, is able to mash-up services created by using Roma, using the standard SOA stack or any other present in the network previously registered. Finally, a set of enterprise services has been created to provide some utilities for the mash-ups, such as Logging and Accounting. Those services are completely general, thought to be used inside a mash-up project. Every component has been developed, tested and released on the Romulus repository, and it is available for the community. This approach is currently used in several integration projects where Imola Informatica, partner of Romulus and main contributor of the approach, is involved in the Italian and European IT market, by other companies and many single developers that are interested in integration topics. Aftab Iqbal, Oana-Elena Ureche, Michael Hausenblas

	I**ntegrating Linked Data Driven Software Development Interaction into an IDE.** In 5th International Workshop on Semantic Web Enabled Software Engineering Workshop at 8th International Semantic Web Conference (ISWC 2009), With \Linked Data Driven Software Development" (LD2SD) we have introduced a light-weight, linked data-based framework that allows to integrate software artefacts, such as version control systems and issue trackers, as well as discussion forums. The so created interlinked data-space enables uniform query and browsing facilities. In this paper we elaborate on the interaction part of LD2SD, and demonstrate how the LD2SD-interaction can be integrated into an Integrated Development Environment (IDE). We have performed an end-user evaluation and report on our findings and outline future steps
S-CUBE	**CD-JRA-1.1.2 Separate design knowledge models for software engineering and service-based computing** This deliverable presents two distinct bodies of knowledge: the first one is for service oriented computing based on a proposed life cycle that incorporates adaptation-specific phases. Each phase is discussed in depth, and methods, techniques and tools for it are presented. Furthermore, cross-phase aspects are investigated. The other body of knowledge concerns more traditional software engineering and business process methodologies, examined from the perspective of service-based applications. A number of preliminary results on the synergy between the two areas are also presented as a stepping-stone for the following deliverables. **CD-JRA-1.1.4 Coordinated design knowledge models for software engineering and service-based computing** In this deliverable, we discuss the need for the S-Cube life-cycle and the development of enhancements to support its implementation. We introduce the area of Service-Oriented Systems Engineering (SOSE) and discuss on how it is different from Traditional Software Engineering (TSE), while still recognising that both of these disciplines have important inputs to make to the development of Service-Oriented Systems. We progress with further development of the S-Cube life-cycle focusing on require **CD-JRA-2.3.3 Requirements for Service Registries in Dynamic Environments and Evaluation of Existing Service Registries** Web service registries are tools for the implementation of loosely-coupled service-based systems. For instance, business processes query registries in order to find services which implement functionality that is needed in the process, and adaptable service compositions need to be aware of which alternatives are available for each service. Furthermore, there is a clear interrelation between end-to-end quality provisioning and monitoring, and service registries, since SLA monitoring and enforcement is based on the availability of a service repository providing an expressive set of metadata. Even more, with the advent of the Internet of Services, an Internet-scale Web service ecosystem with unique scale and heterogeneity characteristics, a number of new challenges for the next generation of Web service registries will arise. First of all, the sheer size of the ecosystem (in terms of number of clients, providers and services) will cause a need for new scalable service discovery mechanisms built on the notions of the Internet. This includes not only discovery of atomic services, but also of task flows (ad hoc service mashups). Additionally, the distributed and heterogeneous nature of the Internet of Services asks for new data dissemination methods between physically and logically disjoint registry entities, which work in spite of missing, untrusted, inconsistent and wrong data. Further challenging requirements are going to be put forward by mobile, human-provided and ad hoc services, which are common in the Internet of Services. These services are volatile in nature, and need to be actively tracked by the service registries. Finally, another class of challenges is introduced by the human factor in the

	Internet of Services -- since services are often consumed and provided by humans, new means of evaluating service performance based on user-perceived and fuzzy Quality of Experience metrics need to be devised. In this deliverable we describe these requirements for the next generation of service registries for large-scale service environments in detail, and explain why we consider existing registry approaches as not sufficient for these environments. **CD-JRA-2.3.2 Basic Requirements for Self-Healing Services and Decision Support for Local Adaptation** One of the goals of S-Cube is to look for general solutions by integrating research agendas from diverse research areas, such as business processes, service-oriented and grid computing. The world of web services already provides solutions for complex user tasks. The web service model is based on three actors: a service provider, a service requester and a service broker. There are also well established and widely used technologies that enhance the collaboration of these three parties to fulfil service executions required by users. The newly emerging demands of users and researchers call for expanding this service model with business-oriented utilisation (agreement handling), support for human-provided and computation-intensive services. This evolution also affects the service infrastructure; new components appear that need to provide self-* operation. The purpose of this document is to capture the basic requirements for self-healing and decision support in service execution, deployment and runtime management for services including core services such as discovery and registries. Concerning service execution, we describe what kind of functionalities and tools should be provided at the infrastructure level in order to be able to implement a self-healing service. We restrict the scope of this document to the adaption of one service, not of a coordinated set of services. Concerning deployment and run-time management, we envision a conceptual architecture for SLA-based on-demand service provisioning and, based on this framework, three main functionalities are separated: negotiation, brokering and deployment. The document investigates the requirements in details for each of these fields. We investigate context-aware requirements discovery and specification, exploring whether existing models can be applied to improve requirements specification. From the design perspective, we suggest design principles and guidelines that are suitable to enable adaptation. From the adaptation perspective, we investigate SOSE and TSE to present practices for adaptation. Finally, we propose a unified formal model for dealing with the effects of iterative and localized changes between any two interacting service consumers and providers **CD-JRA-2.1.2 Initial Models and Mechanisms for Quantitative Analysis of Correlations Between KPIs, SLAs and Underlying Business Processes** In this deliverable we present initial models and mechanisms for quantitative analysis of correlations between KPIs, SLAs and underlying business processes. We use service network (SN) models for quantitative analysis based on KPIs and SLAs, which enables strategic decisions for participants such as determination of optimal product prices or outsourcing decisions. In order to perform the analysis on the SN abstraction level and implement its results in operational business processes, SNs have to be connected to the BPM stack. We therefore introduce the SN4BPM architecture describing an enhanced BPM layering and lifecycle where SNs constitute a separate layer on top of the established BPM stack. In that context, we describe in particular a model-driven approach to generating abstract business process models from Service Network Models and vice versa. Finally, we deal with monitoring in the cross-organisational setting of service networks. **CD-JRA-2.1.3 Business Transaction Language**

	Application integration remains one of the core drivers of innovation in service engineering. Application integration serves as a means of developing service-enabled applications based on strategic technology capable of creating and successfully executing end-to-end business processes. The trend will be to move from relatively stable, organisation-specific applications to integrated, dynamic, high-value ones where process interactions and trends are examined closely to understand more accurately application needs and dynamics. Such collaborative, complex end-to-end process interactions give rise to the concept of Service Networks (SNs). This deliverable targets the concept of a business transaction and explores how transactional processes and process fragments fit in the context of a running scenario which deals with end-to-end processes in a service network that possess transaction properties. Conventional (ACID) and unconventional (application-based) types of atomicity are introduced, including contract, payment and delivery atomicity, in the frame of a business transaction model. The transaction model provides a comprehensive set of concepts and several standard primitives and conventions that can be utilized to develop complex Service-Based Applications (SBAs) involving transactional process fragments
ServFace	**D2.9 Models for Service Annotations, User Interfaces, and Service-based Interactive Applications** This deliverable introduces a global view on the User Interface (UI) models developed for the ServFace project. The ServFace tool-chain relies on three different steps of development for an interactive application: 1. Service Annotation, 2. Application Composition, 3. Application Runtime Each step of the development produces a model as output, and specifically the corresponding models are the following: 1. Annotation Model, 2. Composite Application Model (CAM) for the Simple Composition Tool, MARIA XML for the Advanced Composition Tool; 3. Prosaic. This Deliverable introduces the Annotation Model, the Composite Application Model and MARIA XML. **D3.7 Discussion of Approaches for platform mappings and runtime support.** This deliverable test the main thesis that the development of Model-to-Model transformations reduces the effort that has otherwise to be spent for the laborious development of Model-to-Code transformations and multiple runtime APIs in situations where multiple design-time Meta-Models exist. Within the ServFace project this means that instead of developing several Model-to-Code transformations for the Composite Application Model in the case of the *ServFace Builder* and for the several Concrete User Interface Models in the case of the Advanced Composition Tool called *Maria,* only Model-to-Model transformations using these Meta-Models as source models have to be provided. This approach has been compared to a direct Model-to-Code transformation realized from the output of one of the design-time tools. **D3.8 Automatic Model-Based Generation of User Interfaces from WSDL Descriptions** This deliverable discusses the problem of automatically deriving a User Interface (UI) starting from the information contained into a WSDL file. The analysis also takes into account the usage of ServFace Annotations in order to enhance the quality of the final result. This document consists of three parts. Section 2 introduces the related work in this area, reporting existing results and solutions for automatic UI generation problem. Section 3 discusses the approach of the prototype included in the Advanced Composition tool. In the first part the overall methodology is introduced. Furthermore, it is explained how WSDL and ServFace Annotations are exploited in the generation process. The second part goes through a case study (the Home Scenario) in order to validate the procedure. In this solution all levels of the CAMELEON Reference Framework are considered using the ConcurTaskTrees notation for the task level and the MARIA language for the abstract and concrete

	descriptions. This allows developers to generate user interfaces adapted to the interaction resources of the target device. Section 4 discusses an approach for generating ad hoc UIs during application runtime by interpreting annotated Web Service description files. The approach has been developed in Android environment and has been validated by the development of a Home application. **D3.9 Specification of the ServFace Platform Mapping** Well-defined platform mappings are crucial to fully exploit the potential of the ServFace approach. Due to the additional information available in the ServFace annotation model and the provisioning of non-generable application logic through operations of annotated services, there is no need to manually write additional code. The process of transforming the model-oriented representation of an interactive service based application into an executable can therefore be completely automated and does not require expert knowledge. The ability to construct executables from metamodel instances is also a good indication for the completeness of the respective metamodel and the maturity of the modelling approach. The deliverable presents the different approaches for platform mappings applied in the ServFace project, namely two model-to-code transformations for generating mobile Google Android applications from CAM instances and multimodal service-based applications from MARIA XML models, and one model interpreter approach for constructing Microsoft Silverlight Rich Internet Applications directly at runtime, thereby avoiding the need for intermediate transformation steps. The target platforms were chosen to represent the three application categories that were identified as the most relevant ones for the project partners. Fabio Paterno', Carmen Santoro, Lucio Davide Spano, ACM Transactions on Computer-Human Interaction (TOCHI) archive Volume 16, Issue 4 (November 2009) table of contents Article No. 19 **MARIA: A Universal, Declarative, Multiple Abstraction Level Language for Service-Oriented Applications in Ubiquitous Environments** One important evolution in software applications is the spread of service-oriented architectures in ubiquitous environments. Such environments are characterized by a wide set of interactive devices, with interactive applications that exploit a number of functionalities developed beforehand and encapsulated in Web services. In this paper, we discuss how a novel model-based UIDL can provide useful support both at design and run time for these types of applications. Web service annotations can also be exploited for providing hints for user interface development at design time. At run-time the language is exploited to support dynamic generation of user interfaces adapted to the different devices at hand during the user interface migration process, which is particularly important in ubiquitous environments. Kritikos K., Paternò F. **Service Discovery Supported by Task Models** - In: EICS'10 - 2nd ACM SIGCHI Symposium on Engineering Interactive Computing Systems (Berlin, Germany, 19-23 June 2010). Proceedings, pp. 261 - 266. ACM, 2010 We propose an approach that takes as input a task model, which includes the user's view of the interactive system, and automatically discovers a set of categorized and ranked service descriptions for each system task of the model. In this way, a set of service operations can be used to implement an application's part or whole functionality so that its development time is significantly reduce
SHAPE	**D2.4 SHAPE Integrated and Tool-supported Methodology**

	The aim of the SHAPE Methodology is to provide 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. These provide a comprehensive collection of tool-supported techniques for the model-driven engineering of service-oriented landscapes with support for various technology platforms. In a concrete application scenario, usually only a subset of the available SHAPE techniques is needed to develop a system landscape that allows meeting the respective business requirements. Thus, the purpose of the SHAPE Methodology is to provide a general purpose framework that supports end-users in the selection and combination of the specific engineering techniques required for the individual application scenario. For this, the SHAPE Methodology provides pre-defined methods and methodology processes along with tool support for creating customized methodologies that define the overall engineering procedure for an individual engineering project. This supports the planning and monitoring of the project execution and provides detailed guided procedures for the distinct engineering tasks that are performed by the actors involved in the project. It provides a comprehensive of the SHAPE Methodology, including a concise overview of the overall framework, detailed specifications of the method content along with the installation- and usage instructions for the methodology tools. The public parts of the SHAPE Methodology are available on the project website (www.shape.eu), and the tools are provided for research and demonstration purposes to the project consortium with respect to the usage restrictions as defined in the consortium agreement. **D3.5 Metamodel and Language Extension for Semantic Web Services, Agents, P2P and Grid** It Describe the UML Profile and Metamodel for Semantically-enabled Heterogeneous serviceoriented (currently renamed to ShaML) which extends the UPMS (currently renamed to SoaML). The description is focused in the support of the metamodel for semantics and adaptive technologies. **D5.4** M**odel transformations from business models to UPMSHA**. This deliverable presents the final status of the model transformations from CIM to PIM of the SHAPE methodology. It builds upon the results of the preceding deliverables which presented (i) the overall model transformation and deployment architecture that is refined in this deliverable to reflect the final status of the SHAPE project and (ii) gives an initial overview on the CIM to PIM transformations of the model transformation architecture. The transformations developed between the CIM and PIM level are the following: CIMFlex to ShaML; BPMN to ShaML; and, CIMFlex to UML. The following main progress and results have been achieved since the interim version of the CIM to PIM model transformations: • Refinement of the model transformation between BPMN and SoaML ; • Final integration into the SHAPE tool suite; **D5.5 Model transformations and deployment – from UPMSHA to WSA, agents, P2P, grid and SWS platforms** This deliverable presents the final status of the model transformations from PIM to PIM/PSM of the SHAPE methodology. It builds upon the results of the preceding Deliverable D5.1 [1], which presented the overall model transformation and deployment architecture that is refined in this deliverable to reflect the current status of the SHAPE project. The model transformations presented in this deliverables are: • PIM to PIM o SoaML Profile to SoaML Metamodel • SoaML Metamodel to SoaML Profile • SoaML Profile to PIM4Agents

	• PIM to PSM o ShaML to WSMO • ShaML to SAWSDL • SoaML to XSD • SoaML to BPEL • SoaML to WSDL • SoaML to J2EE The details of each transformation are further demonstrated using examples partially from the industrial cases of SHAPE. The following main progress and results have been achieved since the interim version of the PIM to PIM and PIM to PSM model transformations as presented in: • Refinement of the existing model transformations • Development of the SoaML to J2EE transformation • Development of the SoaML to SAWSDL transformation • Final integration into the SHAPE tool suite • Publication of the research results in a scientific article Christian Hahn, Stefan Nesbigall, Stefan Warwas, Klaus Fischer, Matthias Klusch, "**Model-driven Approach to the Integration of Multiagent Systems and Semantic Web Services**," Enterprise Distributed Object Computing Conference Workshops, IEEE International, pp. 317-324, 2008 12th Enterprise Distributed Object Computing Conference Workshops, 2008 This paper discusses an innovative mean on model-driven agent-based coordination of Semantic Web services. The general idea is to define a platform independent meta-model for Semantic Web services and integrate it into a platform independent metamodel for multiagent systems. A model-driven Semantic Web services matchmaker agent in combination with model transformations between the platform independent metamodel and existing Semantic Web service formats like OWL-S allow the seamless integration of Semantic Web services into multiagent systems. Federico M. Facca, Srdjan Komazec, Michal Zaremba, "**Towards a Semantic Enabled Middleware for Publish/Subscribe Applications,"** International Conference on Semantic Computing, pp. 498-503, 2008 IEEE International Conference on Semantic Computing, 2008 This paper discusses an innovative mean on model-driven agent-based coordination of Semantic Web services. The general idea is to define a platform independent meta-model for Semantic Web services and integrate it into a platform independent metamodel for multiagent systems. A model-driven Semantic Web services matchmaker agent in combination with model transformations between the platform independent metamodel and existing Semantic Web service formats like OWL-S allow the seamless integration of Semantic Web services into multiagent systems.
SLA@SOI	**D.A3a SLA-aware Service Management**

	SLA@SOI aims at addressing the SLA management problem from a holistic perspective. The objective is to design a multi-layer SLA management framework for SOA based application landscapes. The focus of work package A3 is the software service layer and aims at designing SLA aware service management capabilities for the overall SLA management framework. In this regard, WP A3 has divided the service management into a number of distinct yet complementary and interconnected working tracks. This deliverable document provides description of the activities, progress and achievements during the course of Y2. Building on the work carried out during Y1, work documented in this deliverable is targeted for the various cases. A brief summary of these activities are given in the following paragraph. Firstly, A3 engaged in modelling related activities. The modelling aspects addressed were SOA modelling and software landscape modelling. SOA modelling investigated modelling of service component along with the non-functional properties as well as the service component behaviour which is leveraged by the design time prediction process. Software landscape modelling, on the other hand focused designing a metamodel which can be used to capture information about the service and software related artefacts. Additionally, landscape meta-model incorporated packaging and deployed related aspects to be used by service provisioning process. Secondly, A3 continued work on the dynamic service binding and composition related activities. During Y1, design and architecture of a Dynamic Orchestration Engine was focus of Y1 activities. Building on this an implementation of the DOE was the main focus on Y2 activities. This document presents the DOE implementation specific details. Last but not least, A3 s work focused on runtime monitoring and management of SOA based application landscapes. This document presents details of Monitoring Manager which is responsible for translating SLA guarantee terms into monitoring rules which can be observed during runtime to detect potential SLA violations. Additionally, the document discusses the process monitoring capabilities of the SLA@SOI framework. **D.A6a Predictable / Manageable Service Engineering Methodology and Prediction Services** This work package aims to support the engineering process for predictable software services and components. Predictability is an important feature that helps service and infrastructure providers to make well-informed decisions throughout service design, offering, negotiation, provisioning and run-time. At its core, predictability is the capability to anticipate quality-related properties of services and service hierarchies, before those properties can actually be observed. There are four main contributions to SLA@SOI resulting out of the work within WP A6: • Software Performance and Reliability Prediction: Provides a means for software service providers to evaluate the expected performance and reliability of their services before their run-time. May be used at (i) service offering / SLA template design stage to determine realistic quality parameters to offer, or at (ii) (automated) service / SLA negotiation, to determine realisable quality parameters to agree upon. • Resource Usage Prediction: Provides a means for infrastructure service providers to predict expected infrastructure resource demands at system run-time. May be used at service provisioning stage to determine the best option for deployment of a new VM. • Run-time SLA Violation Prediction: Provides a means for software service providers to anticipate possible near-future software SLA violations at system run-time. May be used at service operations / SLA runtime stage to trigger adjustment

	activities in order to avoid an actual SLA violation. • Manageability Modelling and Design: Provides a design-time methodology and a set of tools for enhancing service components with management capabilities. In particular, service components are augmented with the capabilities needed to (i) extract the runtime data needed to monitor the quality of running services, and to (ii) adjust a service's behaviour and/or its configuration at run time. **D.A2a Business SLA Management** This deliverable is the summary of the work done in second year of the European Project SLA@SOI in the Business SLA Management work package. The objective of this work package is to enrich SLAs, from the business level down to the physical infrastructure, with the business terms and constraints needed in a real eContracting framework. The project envisages developing an automated e-contracting lifecycle that covers the whole interaction between service providers and customers. This spans from the commercial product definition (based on offers and prices) to the negotiation and sales process. The management of the post-sales relationship including SLA monitoring and customer behaviour analysis is also addressed. It is necessary for service providers to define service level agreement templates (SLAT), i.e., template contracts including business terms and constraints. This deliverable presents the architecture used to provide this functionality, and the SLA model customization used to define and describe the Business Terms and Parameters. The Business Terms are based on the work done in the first year, in feedback provided by the use cases of the project and other external sources about contracting. The document also explains product management questions related to business management, and how answers to these questions have been implemented. The business negotiation approach for selling products as well as details on how this can be customized by business people is explained. The business assessment flow used to modify the prices and parameters stored by the framework in the selling process is detailed, as is the translation of violations into business actions like penalties and service terminations. Regarding the customers, this deliverable also describes how their relationships can be managed. This includes the different ways to communicate with customers so that they are able to consume products and services and to receive reports and information about the history of the consumption of their SLAs. It also explains the way they can provide their feedback through the framework, so that the business can study it and take decisions to enhance customer perception and satisfaction. **D.B6a Use Case Specification eGovernment**. This document describes the main requirements and a first architectural hypothesis for the application (eGovernment use case) of the SLA@SOI framework in the context of the Healthcare System of the Italian region of Trentino. The document concerns mainly with the first step of this plan and provide just an high level view of the mobility reservation process. Here SLAs are mainly described at business level as these are the SLAs treated in the current situation. The main contribution of the SLA@SOI framework to this use case are the improvement of the monitoring capability and the optimisation of resource consumption. Marco Comuzzi, Constantinos Kotsokalis, George Spanoudakis, Ramin Yahyapour, "**Establishing and Monitoring SLAs in**

	Complex Service Based Systems," Web Services, IEEE International Conference on, pp. 783-790, 2009 IEEE International Conference on Web Services, 2009 In modern service economies, service provisioning needs to be regulated by complex SLA hierarchies among providers of heterogeneous services, defined at the business, software, and infrastructure layers. Starting from the SLA Management framework defined in the SLA@SOI EU FP7 Integrated Project, we focus on the relationship between establishment and monitoring of such SLAs, showing how the two processes become tightly interleaved in order to provide meaningful mechanisms for SLA management. We first describe the process for SLA establishment adopted within the framework; then,we propose an architecture for monitoring SLAs, which satisfies the two main requirements introduced by SLA establishment: the availability of historical data for evaluating SLA offers and the assessment of the capability to monitor the terms in a SLA offer. Christof Momm, Michael Gebhart, Sebastian Abeck, "**A Model-Driven Approach for Monitoring Business Performance in Web Service Compositions,**" Internet and Web Applications and Services, International Conference on, pp. 343-350, 2009 Fourth International Conference on Internet and Web Applications and Services, 2009 Supporting business services through Web service compositions (WSC) as part of service-oriented architectures (SOA) involves business performance monitoring requirements. Their implementation results in additional development activities. To support these activities, we already contributed a model-driven approach to the development of monitored WSC as part of our preliminary work. In this paper, we present an extension to this approach, which focuses on supporting the specification and transformation of indicators to an executable implementation. To reduce development effort for this particular task, we provide a template-based mechanism for defining performance indicators. In combination with our preliminary work, now fully monitored WSC can be generated automatically from platform-independent design models. We demonstrate the applicability of the overall approach by instantiating an integrated development process for a target platform based on IBM SOA products and showing its application for a sample business process along with monitoring requirements.
SMART LM	**Architectural Overview of the SmartLM Innovative License Management System** Software plays a critical role in business. While conventional, revenue-based business models are still dominating licensing mechanisms, it is evident that the market is getting restless, and the demand for more flexible licensing solutions from customers is growing. Enterprise IT infrastructures are evolving and software licensing needs to evolve with them. Software manufacturers need to change the way licensing works and use flexible and non-hardware based licensing solutions that better fit into distributed and virtual environments. 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. The focus of successful licensing and support has to extend beyond cost and technology issues, the goal is to achieve software licensing based on business objectives that balances customer needs and vendor business models. The SmartLM project has addressed the licensingproblem by working on a framework which delivers impro ved customer choice and that also keeps software vendors happy. SmartLM's offering brings along a model that makes licenses mobile objects. The main approach here is to provide platform-independent access and treat software licenses as services.

	Deliverables are not public information are taken from executive summaries **D4.1 Process Model for Accounting.** This deliverable addresses the process-model for accounting and billing. Two main scenarios are discussed. SmartLM is a token-based system. The accounting process hence will regard the token both as a record and also as a symbolical receipt with an associated price. A portal-based aggregation over the respective license services is employed. **D7.4 Preliminary exploitation plan.** This deliverable provides the preliminary SmartLM exploitation plan. The commercial name is elasticLM which aims at rending mechanisms for managing and using software licenses in a more fair and flexible way. The package consists of three components: license service, accounting and billing service and application interface, all available for download from the website. There are three options to acquire elasticLM. Other services such as implementation support may also be part of the package).
SOA4ALL	**D1.4.1A SOA4All Reference Architecture Specification** The document provides a high level description of the SOA4All platform by showing the relationship between the SOA4All core infrastructure services, the SOA4All Platform Service and the SOA4All Studio, respectively the external third-party business services. Based on a general architecture design methodology, we specify in more details the different components of the SOA4All platform, and their relationships to and dependencies. Another objective of this deliverable is to clearly determine the role of, and the links between different components, and to define that objects that are exchanged and shared; mainly services, ontologies, goals, and processes. These definitions are, from a conceptual point of view, closely aligned with the recent standardization work of the OASIS SEE Technical Committee and their SEE Reference Ontology specification. The last viewpoint that we provide on the architecture is related to functional processes that are realized in SOA4All. This last point of view is partly an outlook to work that will be conducted in the second year of the project, that is to say, the definition and realisation of more complex processes that involve multiple SOA4All Platform Services; for instance, the discovery of service, or the composition of processes to name two. The descriptions of the functional processes are given along the lines of the NEXOF-RA requirements specification for service architectures. This alignment with NEXOF-RA has two advantages: i) the architecture specification in terms of functional processes is well-founded in terms of the definitions and conceptualizations of other NESSI strategic projects, and ii) the SOA4All reference architecture is easily comparable to the NEXOF-RA expectations and requirements, which clearly eases the exploitation of research results in NEXOF-RA. In terms of architectural specifications, there is a short section that outlines the expectations of WP1 in the context of the upcoming implementation task **D2.1.2 Service Modelling Tools Design** The present deliverable complements and is heavily interrelated with deliverable D2.1.1, where the Service Provisioning Platform is described. We focus here on the tools that we will use in order to semantically annotate services, grouped in the so-called “Simple Semantic Web Services Editing Framework”. The main outcomes of this framework are two editors that will allow users to enrich traditional WSDL-based and RESTful services with semantic annotations: The WSMO-Lite Editor and

	the MicroWSMO Editor, respectively. We have identified several characteristics that these tools will need to have in order to achieve the ambitious objective of the project of having an extremely large number of services available for consumption. Concretely, our tools will be lightweight and web-based, useable by both expert and non-expert users, and enabling a community approach towards modelling. We consider these characteristics are key for reaching a scenario of many services being deployed. The detailed design of these tools is then provided, highlighting the functionalities they will cover and the interactions with other architectural components of SOA4All. **D3.2.1 Framework and APIs for integrated reasoning support** This deliverable has two main purposes: First, it is supposed to serve as a requirements document for the integrated reasoner component in SOA4ALL. Secondly, and based on this requirements analysis, it will define a high-level API for the required infrastructure to support reasoning with service annotations by taking the existing WSML2Reasoner tool as a baseline to evolve from. The result of this deliverable is the definition of an high-level API for the reasoning component, which facilitates relevant reasoning tasks in an optimized manner for the underlying formalisms WSML-Core, WSML-Rule and WSML-DL in order to meet the scalability goals of SOA4LL. Deliverable D3.2.1, situated at the base-layer of the SOA4ALL architecture, is of interest to various other components in the Web Enabled Service Platform layer and Service Construction, for which reasoning is basic infrastructure in the process of service discovery and composition. **D6.5.2 Advanced Prototype For Adaptive Service Composition Execution** The Advanced Prototype For Adaptive Service Composition Execution (EE v2) adds further functionalities to the lightweight composition environment presented in the first prototype, the EE v1 presented in [5]. EE v1 is able to adapt service requests to actual service interfaces at execution time, it supports the service selection according to functional, non-functional and contextual information and provide fault handling and dynamic rebinding mechanisms. The new functionality of EE v1 are: RESTful services support: EE v1 was able to work only with Web Services based on the WSDL and SOAP standards; the EE v2 overcomes this limitation allowing a process to use WSDL/SOAP and RESTful Web services at the same time transparently. Automatic deployment of LPML models: when the EE v1 was released, the lightweight process modelling language was still under development, now the EE v2 includes full support for processes described in the LMPL as it allows the deployment of LPML models automatically from the Process Editor1 (described in [24]) by translating them into executable processes/services. Support for Human tasks: EE v2 include the support for processes that contains activities that are executed by humans such as checking requests for completeness and correctness. These activities are called human tasks. As for the integration aspects, the EE v2 is ready to be integrated with some of the other SOA4All components, namely: Process Editor: the EE v2 exposes a Web Service that provides access to the deploy functionality of the EE; this service is invoked by the Process Editor when the user asks for the execution of a newly developed process. Consumption Platform: when a process is deployed, it can be accessible as an HTTP/SOAP Web Service; the SOA4All component in charge of invoking the executable process is the Consumption Platform. Analysis platform: during execution, the EE v2 generates events that trigger messages sent to the Analysis Platform, that is the component in charge of displaying these events to the users. Furthermore, the EE v2 includes some internal improvements, with respect to the previous prototype: 1. Decoupling from BPEL engine implementation: while the EE v1 supported only Active BPEL, now the EE v2 supports any

	BPEL engine implementation2; 2. Update to the latest version of Axis2; 3. General refactoring and re-design for simplifying the architecture of the component. The EE v2 is available with its source code through the SOA4All SVN repository. Finally, a final release of the EE prototype is planned by the end of August 2010 (month 30 of the SOA4All the project). In the final release further improvements will be implemented, if needed, based on the feedbacks coming from the implementation of the SOA4All use cases **D7.3 End User Service Design.** This deliverable serves as a guideline and as a technical specification for implementing the service delivery platform based on the technical components as well as the Web-based front end developed within the SOA4All project. The deliverable also includes the design of a semantic adaptation and integration layer for handling SAP Enterprise Services, a detailed and updated version of the use case story board that demonstrates the features from the user's point of view, and a description of the first evaluation workshop for validating the project's ideas. **D9.3.1 C2C e-Commerce Prototype v1.** The deliverable shows how the SOA4All results can be applied in the e-Commerce domain. The scenario goes clearly beyond a classical use case in a way that it does not only use and apply the results provided by the project. Instead of this, it also adds own ideas and developments to SOA4All, allowing the use case to show the innovation that SOA4All brings to e-Commerce in a future looking and highly flexible web 2.0 environment. The purpose of this scenario is to demonstrate the SOA4All vision by telling a real-world story around the complete set of SOA4All components in a highly practical way so showing the usefulness of the project results.
Stream	Gulisano, V. Jimenez-Peris, R. Patino-Martinez, M. Valduriez, P. Fac. de Inf., Univ. Politec. de Madrid, Madrid, Spain, **StreamCloud: A Large Scale Data Streaming System** in: Distributed Computing Systems (ICDCS), 2010 IEEE 30th International Conference on Data streaming has become an important paradigm for the real-time processing of continuous data flows in domains such as finance, telecommunications, networking, Some applications in these domains require to process massive data flows that current technology is unable to manage, that is, streams that, even for a single query operator, require the capacity of potentially many machines. Research efforts on data streaming have mainly focused on scaling in the number of queries or query operators, but overlooked the scalability issue with respect to the stream volume. In this paper, we present StreamCloud a large-scale data streaming system for processing large data stream volumes. We focus on how to parallelize continuous queries to obtain a highly scalable data-streaming infrastructure. StreamCloud goes beyond the state of the art by using a novel parallelization technique that splits queries into subqueries that are allocated to independent sets of nodes in a way that minimizes the distribution overhead. StreamCloud is implemented as a middleware and is highly independent of the underlying data streaming engine. We explore and evaluate different strategies to parallelise data streaming and tackle with the main bottlenecks and overheads to achieve scalability. The paper presents the system design, implementation and a thorough evaluation of the scalability of the fully implemented system.

	Andrey Brito, Christof Fetzer, Pascal Felber, "**Minimizing Latency in Fault-Tolerant Distributed Stream Processing Systems," Distributed Computing Systems**", International Conference on, pp. 173-182, 2009 29th IEEE International Conference on Distributed Computing Systems, 2009 Event stream processing (ESP) applications target the real-time processing of huge amounts of data. Events traverse a graph of stream processing operators where the information of interest is extracted. As these applications gain popularity, the requirements for scalability, availability, and dependability increase. In terms of dependability and availability, many applications require a precise recovery, i.e., a guarantee that the outputs during and after a recovery would be the same as if the failure that triggered recovery had never occurred. Existing solutions for precise recovery induce prohibitive latency costs, either by requiring continuous checkpoint or logging (in a passive replication approach) or perfect synchronization between replicas executing the same operations (in an active replication approach). We introduce a novel technique to guarantee precise recovery for ESP applications while minimizing the latency costs as compared to traditional approaches. The technique minimizes latencies via speculative execution in a distributed system. In terms of scalability, the key component of our approach is a modified software transactional memory that provides not only the speculation capabilities but also optimistic parallelization for costly operations.

CONCLUSIONS

After the document inventory and the achievement of a better understanding of the projects - both at single level and at comparative level – the consortium developed 2 questionnaires: one for in-depth interviews of key experts and another for the projects. Since this second questionnaire was complete in terms of variables used and very long, the consortium saw a risk in sending this questionnaire to the projects, in view of the lack of responses to an introductory mail sent by the project coordinator. During the July Update project meeting, the consortium agreed on the need to postpone the use of this second questionnaire and to develop a new questionnaire, shorter and more user-friendly. This questionnaire, named "preliminary questionnaire", has been submitted online first; then, having received very few answers, has been sent by email while contacting each project directly also by phone. The need to contact the projects several times and the difficult in opening a direct channel of communication implied a slowing down of the research process. In D2.2a we will describe the outputs of this first survey, together with the result of the in-depth interviews of key experts.

In the Collaboration meeting held in Brussels on 18 and 19 October 2010 the consortium had the opportunity to meet some of the Call 1 projects in person, so that the questionnaire could be distributed and collected shortly thereafter, collecting some more information. On that occasion, we also met some of the Call 5 projects, which showed a lively interest in SEQUOIA's activities. Thus, we were able to start communicating and exchanging information with them. A short collaborative session was organised during the Collaboration meeting and we gathered some interesting information about the participating projects in terms of their users and beneficiaries.

ANNEX A

(Next page)

ADMIRE

Making data-mining easier

ADMIRE (Advanced Data Mining and Integration Research for Europe) aims to deliver a consistent and easy-to-use technology for extracting information and knowledge. The project is motivated by the difficulty of extracting meaningful information by data mining combinations of data from multiple heterogeneous and distributed resources. It will also provide an abstract view of data mining and integration, which will give users and developers the power to cope with complexity and heterogeneity of services, data and processes.

Aims & Objectives

The project will, over the three years, invest two thirds of its effort in developing advances in distributed system technology and one third in developing the precise models that will enable DMI systems to become much easier to use and simultaneously to become scalable and dependable through judiciously engineered flexibility.

The work will be organised as six technical and one management work packages running concurrently for the full three years.

Deliverables:

- Deliverable report D1.4, On the Systematic Design of DMIL, August 2009
- Deliverable report D1.5, ADMIRE - Report on progress of model, language and ontology research, August 2009
- Deliverable report D2.4, ADMIRE - Internal report on goals, scope and draft definition of AA3, August 2009
- Software platform D3.3, ADMIRE Platform Release 2 - Description, August 2009
- Deliverable report D4.3, ADMIRE - Development Progress Report, August 2009
- Deliverable report D5.3, ADMIRE - Tools Development Progress Report, August 2009
- Deliverable report D6.3, ADMIRE - Report on Progress of ADMIRE Pilot Applications Development and Deployment, August 2009
- Deliverable report D1.2, ADMIRE - DMI Model, Language and Ontology, February 2009
- Research prototype D1.3, ADMIRE - DMI-WF: XML-Schema specification version 1.0, February 2009
- Research prototype D2.2 DMI Using a Pipeline, February 2009
- Software platform D3.2, ADMIRE Platform Release 1 - Description, February 2009
- Deliverable report D4.2, ADMIRE - Development and Deployment Report for USMT V2, February 2009
- Deliverable report D5.2, ADMIRE - Tools Development Report and Requirements Analysis, February 2009

- Deliverable report D6.2, ADMIRE - Report on Validation of Initial ADMIRE Model and Architecture, February 2009
- Deliverable report D1.1, ADMIRE - Towards the High-Level DMI Model, Language and Ontology, August 2008
- Deliverable report D2.1, ADMIRE - Architecture, August 2008
- Deliverable report D3.1, ADMIRE - Platform Delivery and Support Progress Report, August 2008
- Deliverable report D4.1, ADMIRE - Development and Deployment Report for USMT V0 and V1, August 2008
- Deliverable report D5.1, ADMIRE - Tools Development Progress Report, August 2008
- Deliverable report D6.1, ADMIRE - Architecture and Design of the Pilot Applications, August 2008

Publications:

- Supporting Molecular Modeling Workows within a Grid Services Cloud. M. Köhler, M. Ruckenbauer, I. Janciak, S. Benkner, H. Lischka, W. Gansterer, International Conference of Computational Science and Its Applications (ICCSA), Fukuoka, Japan. 2010
- Estimation of Parameters Sensitivity for Scientific Workflows. Fakhri Alam Khan, Yuzhang Han, Sabri Pllana, and Peter Brezany. International Conference on Parallel Processing, Vienna, Austria, 2009.
- An Ant-Colony-Optimization based approach for Determination of Parameter Significance of Scientific Workflows. Fakhri Alam Khan, Yuzhang Han, Sabri Pllana, and Peter Brezany. International Conference on Advanced Information Networking and Applications, Perth, Australia, 2010.
- Workflow Enactment Engine Independent Provenance for e-Science Workflows. Fakhri Alam Khan, Sardar Hussain, Ivan Janciack, and Peter Brezany, International Conference on Research Challenges in Information Science, Nice, France, 2010.
- Towards Large-Scale Scientific Dataspaces for e-Science Applications. Ibrahim Elsayed and Peter Brezany. Data Intensive eScience Workshop, Tsukuba, 2010.
- Towards Optimising Data-intensive Workflows using Parallel Streams. Chee Sun Liew, Liangxiu Han, Jano van Hemert, Malcolm Atkinson. The Third International Workshop on Data Intensive Distributed Computing (DIDC2010), Chicago, Illinois - June 22, 2010.
- Integrating Distributed Data Sources with OGSA-DAI DQP and Views. B. Dobrzelecki, A. Krause, A. Hume, A. Grant, M. Antonioletti, T. Alemu, M. Atkinson, M. Jackson, E. and Theocharopoulos. To appear in Phil. Trans. R. Soc. A.
- Advanced Data Mining and Integration Research for Europe. Malcolm Atkinson, Peter Brezany, Oscar Corcho, Liangxiu Han, Jano van Hemert, Ladislav Hluch_, Ally Hume, Ivan Janciak, Amy Krause, Dave Snelling and Alex Wöhrer. The UK e-Science All Hands Meeting, December 2009, Oxford.
- Cloud-Enabled Scalable Decision Tree Construction, Yuzhang Han, Peter Brezany, and Ivan Janciak. International Conference on Semantics, Knowledge and Grid 2009, 12-14 Oct 2009, Zhuhai.
- Scalable Relational Query Results Handling in Service Oriented Architectures, Alexander Wöhrer, Thomas Lustig and Peter Brezany, International Workshop on High Performance Distributed Data Management (HPDDM'09), 9-12 Nov 2009, London.

- Application of Data Integration and Mining to Environmental Scenarios. Ondrej Habala, Martin _eleng, Viet Tran, Ladislav Hluch_. 5th International Workshop on Grid Computing for Complex Problems, October 2009, Bratislava.
- Tools for Advanced Data Mining. Branislav Simo, Michal Laclavik, Ivan Janciak, Ladislav Hluch_. 5th International Workshop on Grid Computing for Complex Problems, October 2009, Bratislava.
- Unboxing Data Mining via Decomposition in Operators: Towards Macro Optimization and Distribution, Alexander Wöhrer, Yan Zhang, Ehtesam-ul-Haq Dar and Peter Brezany. International Conference on Knowledge Discovery and Information Retrieval, Madeira, 6 - 8 October, 2009, Portugal
- Automating Gene Expression Annotation for Mouse Embryo, Liangxiu Han, Jano van Hemert, Richard Baldock, Malcolm Atkinson. International Conference on Advanced Data Mining and Applications (Springer)
- A Distributed Architecture for Data Mining and Integration, Liangxiu Han, Malcolm Atkinson, Ally Hume, Jano van Hemert, Chee Sun Liew. Second International Workshop on Data-Aware Distributed Computing at the International ACM Symposium on High Performance Distributed Computing (ACM).
- Robust service-based semantic querying to distributed heterogeneous databases, Buil-Aranda C, Corcho O, Krause A. Workshop on Web Semantics (WebS'09) at DEXA 2009, 31 Aug - 4 Sep 2009, Linz.
- Using architectural simulation models to aid the design of data intensive application, Javier Fernández, Liangxiu Han, Alberto Nuñez, Jesus Carretero, Jano van Hemert.

Other Papers

- ADMIRE White Paper (v0.8), Jun 2009 (deprecated).
- ADMIRE – Use Case and Requirements Report. Technical report v1.0, Aug 2008.

ALIVE

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. More specifically the project aims to combine cutting edge Coordination and Organisation mechanisms (providing flexible, high-level means to model the structure of inter-actions between services in the environment) and Model Driven Design (providing for automated transformations from models into multiple target platforms) to create a framework for software and services engineering for “live” open systems of active services.

The paradigm explored and elaborated by ALIVE aims to become a foundational contribution to the future evolution of SOA applications and for networked software applications, improving not only the competiveness in the European software industry but also creating new business opportunities by giving support to a wide variety of distributed business scenarios.

Deliverables :

- Context Tool Prototype (Version 1): OperettA, September 2009
- Theoretical Framework (Version 2), September 2009
- Workflow Tool Prototype (version 1), September 2009
- Service Templates (final), September 2009
- Service Matchmaking Framework (version 1), September 2009
- Service Adaptors, September 2009
- Service-level Tools and Integration Tools, September 2009
- Methodology (version 2), September 2009
- Diagramming Tools (version 1), September 2009
- Usage Guide of Design Tools (version 1), September 2009
- Use Case Models and Prototypes (version 1), September 2009
- Formal Model for Organizational Context, May 2009
- Model for Dynamic Context Update, May 2009
- Operational Model, May 2009
- Computational Mechanisms for Workflow Synthesis and Analysis, May 2009
- Theoretical Framework, January 2009
- Methodology, January 2009
- Service Templates, Novembre 2008
- State of the Art, October 2008
- Technical Requirements, October 2008

Pubblications:

- Possibilistic Stratifies Minimal Model Semantics, Juan Carlos Nieves, Mauricio Osorio, September 2009

- Evaluating Organizational Configurations, Loris Penserini, Frank Dignum, Virginia Dignum, Huib Aldewereld, Davide Grossi, September 2009
- A Preference Meta-Model for Logic Programs with Possibilistic Ordered Disjunction, Roberto Confalonieri, Juan Carlos Nieves, and Javier V´azquez-Salceda, September 2009
- A Query Language for Action Domains Modelled using Answer Set Programming, Luke Hopton, Owen Cliffe, Marina De Vos and Julian Padget, September 2009
- Modelling Normative Frameworks using Answer Set Programming, Owen Cliffe, Marina De Vos and Julian Padget, September 2009
- Norms, Organisations and Semantic Web Services: the ALIVE approach, Sergio Alvarez-Napagao, Owen Cliffey, Javier Vazquez-Salceda and Julian Padget, September 2009
- A Query Language for Virtual Institutions using Answer Set Programming, Luke Hopton, Owen Cliffe, Marina De Vos and Julian Padget, September 2009
- A framework to model norm dynamics in Answer Set Programming, Sofia Panagiotidi, Juan Carlos Nieves and Javier Vázquez-Salceda, September 2009
- Expressing Extension-Based Semantics based on Stratified Minimale Models, Juan Carlos Nieves, Mauricio Osorio and Claudia Zepeda, June 2009
- Combining Organisational and Coordination Theory with Model Driver Approached, Javier V´azquez-Salceda, Luigi Ceccaroni, Frank Dignum,
- Wamberto Vasconcelos, Julian Padget, Siobhan Clarke, Paul Sergeant and Kees Nieuwenhui, June 2009
- Developing Agent-based Organizational Models for Crisis Management, Thomas B. Quillinan, Frances Brazier, Huib Aldewereld, Frank Dignum, Virginia Dignum, Loris Penserini, Niek Wijngaards, May 2009
- A Formal Specification for Organizational Adaptation, Huib Aldewereld, Frank Dignum, Virginia Dignum, and Loris Penserini, May 2009
- Building Multi-Agent Systems for Workflow Enactment and Exception Handling, Joey Lam, Frank Guerin, Wamberto Vasconcelos, and Timothy J. Norman, May 2009
- CROC: a Representational Ontology for Concept, Aris van Dijk, Huib Aldewereld, Virginia Dignum, February 2009
- Enforcing Security in the AgentScape Middleware, Thomas B. Quillinan, Martijn Warnier, Michel Oey, Reinier Timmer and Frances Brazier, February 2009
- Adaptivity within an Organizational Development Framework, Loris Penserini, Huib Aldewereld, Frank Dignum and Virginia Dignum, February 2009
- Towards Goal-Oriented Development of Self-Adaptive Systems, Mirko Morandini, Loris Penserini and Anna Perini, February 2009
- Designing MAS Organisation through an integrated MDA/Ontology Approach? Daniel Okouya,Loris Penserini, Sebastien Saudrais, Athanasios Staikopoulos, Virginia Dignum and Siobh an Clarke, February 2009
- Mutual Dynamic Adaptation of Models and Service Enactment in ALIVE, Athanasios Staikopoulos, Sebastien Saudrais, Siobhn Clarke, Julian Padget, Owen Cliffe and Marina De Vos, January 2009
- Coping with Exceptions in Agent-BasedWorkflow Enactments? Joey Sik-Chun Lam, Frank Guerin, Wamberto Vasconcelos, and Timothy J. Norman, January 2009
- PaTac: Urban, ubiquitous, personalized services for citizens and tourists, Luigi Ceccaroni, Víctor Codina, Manel Palau, Marc Pous, January 2009

- Dynamic orchestration on interactive community approach, Gómez-Sebastiàl, Manel Palau, Salceda and Luigi Ceccaroni, January 2009
- Norm Dynamics in Adaptive Organisations, Huib Aldewereld, Frank Dignum, Loris Penserini, and Virginia Dignum, January 2009
- Regulating Organizations: The ALIVE Approach, Huib Aldewereld, Loris Penserini, Frank Dignum, and Virginia Dignum, January 2009
- Representing and Reasoning about Norm-Governed Organisations with Semantic Web Languages? Joey Sik-Chun Lam, Frank Guerin, Wamberto Vasconcelos, and Timothy J. Norman, January 2009
- ALIVE: A Model Driven approach to Coordination Organisation for Dynamic Services Engineering, Siobhn Clarke, Athanasios Staikopoulos, S_bastien Saudrais, Javier Vazquez-Salceda, Virginia Dignum, Wamberto Vasconscelos, Thomas Quillinan, Luigi Ceccaroni and Chris Reed, January 2009

COMPAS

The COMPAS project will design and implement novel models, languages, and an architectural framework to ensure dynamic and on-going compliance of software services to business regulations and stated user service-requirements. COMPAS will use model-driven techniques, domain-specific languages, and service-oriented infrastructure software to enable organizations developing business compliance solutions easier and faster"

Deliverables:

- D7.1. Report on Public Web-Site
- D8.1. Project Quality Plan
- D2.1. State-of-the-Art in the Field of Compliance Languages
- D5.1. State-of-the-Art in the Field of Adaptive Service Composition Monitoring and Management
- D7.2. Collaboration Activities Plan
- D1.1. Model-driven Integration Architecture for Compliance
- D2.2. Initial Specification of Compliance Language Constructs and Operators
- D3.1. Specification of a Behavioral Model for Services
- D4.1. State-of-the-art report on the existing approaches to improving reusability of processes and service compositions
- D5.2. Initial goal-oriented data model
- D7.4. Report of collaboration activities and updates of the collaboration activities plan
- D2.3. Design of Compliance Language Run-time Environment and Architecture
- D3.2. Visual Environment for Service Description
- D5.3. Final Goal-oriented Data Model
- D1.2. Core Meta-models, Transformation Templates, and Languages
- D1.3. MDSD Software Framework for Business Compliance
- D2.6. Implementation of an Integrated Prototype handling Interactive User Specified Compliance Requests in a Compliance Language
- D3.3. Verification Tools for Service Descriptions
- D4.2. BPEL Extensions for Compliant Services
- D4.4. Supporting Infrastructure – Process Engine, Process Artefact Repository, Process Generation Tool
- D5.4. Reasoning Mechanisms to Support the Identification and the Analysis of Problems Associated with User Requests
- D5.5. Initial Prototype of Compliance Governance Dashboards
- D7.4. Report of collaboration activities and updates of the collaboration activities plan

Publications:

Conference Proceedings

- T. Holmes, U. Zdun, F. Daniel and S. Dustdar: Monitoring and Analyzing Service-based Internet Systems through a Model-Aware Service Environment. Proceedings of the 22nd International Conference on Advanced Information Systems Engineering, Springer, 2010 (to appear).
- T. Holmes, H. Tran, U. Zdun and S. Dustdar: Model-driven and Domain-specific Architectural Knowledge View for Compliance Meta-data in Process-Driven SOAs. Proceedings of the 5th International Workshop on Sharing and Reusing Architectural Knowledge in conjunction with 32nd International Conference on Software Engineering, ACM, 2010 (to appear).
- B. Changizi, N. Kokash and F. Arbab: A Unified Toolset for Business Process Model Formalization. Proceedings of the International Workshop on Formal Engineering approaches to Software Components and Architectures (FESCA), 2010 (to appear).
- N. Kokash, C. Krause and E.P. de Vink: Data-Aware Design and Verification of Service Compositions with Reo and mCRL2. ACM Symposium on Applied Computing 2010, Technical track on Service Oriented Architectures and Programming (ACM SAC), 2010 (to appear).
- D. Schumm, F. Leymann, Z. Ma, T. Scheibler and S. Strauch: Integrating Compliance into Business Processes: Process Fragments as Reusable Compliance Controls. Proceedings of the Multikonferenz Wirtschaftsinformatik (MKWI), 2010.
- E. Mulo, U. Zdun and S. Dustdar: Monitoring Web Service Event Trails for Business Compliance. Proceedings of the IEEE International Conference on Service-Oriented Computing and Applications (SOCA), IEEE Computer Society Press, 2009.
- P. Silveira, C. Rodríguez, F. Casati, F. Daniel, V. D'Andrea, C. Worledge and Z. Taheri: On the Design of Compliance Governance Dashboards for Effective Compliance and Audit Management. Proceedings of the Non Functional Properties and Service Level Agreements Management in Service Oriented Computing Workshop (NFPSLAM-SOC2009), Stockholm, Sweden, 2009.
- T. Holmes, U. Zdun and S. Dustdar: MORSE: A Model-Aware Service Environment. Proceedings of the 4th IEEE Asia-Pacific Services Computing Conference (APSCC), pp. 470-477, IEEE Computer Society Press, 2009.
- T. Anstett, D. Karastoyanova, F. Leymann, R. Mietzner, A. Monakova, D. Schleicher and S. Strauch: MC-Cube: Mastering Customizable Compliance in the Cloud. Proceedings of the 7th International Joint Conference on Service Oriented Computing (ICSOC), Springer, 2009.
- N. Kokash and F. Arbab: Formal Behavioral Modeling and Compliance Analysis for Service-Oriented Systems. Formal Methods for Components and Objects, 7th International Symposium, FMCO 2008, Revised Lectures. Vol. 5751 of LNCS, pp. 21-41, Springer, 2009.
- C. Mayr, U. Zdun, and S. Dustdar: Reusable Architectural Decision Model for Model and Metadata Repositories. In Formal Methods for Components and Objects, 7th International Symposium, FMCO 2008, Revised Lectures. Vol. 5751 of LNCS, Springer, 2009.
- M. Wieland, K. Görlach, D. Schumm and F. Leymann: Towards Reference Passing in Web Service and Workflow-based Applications. Proceedings of the 13th IEEE Enterprise Distributed Object Conference (EDOC 2009), August 2009.

- U. Zdun, M. Strembeck: Reusable Architectural Decisions for DSL Design. Foundational Decisions in DSL Development, In: Proc. of the 14th European Conference on Pattern Languages of Programs (EuroPLoP), Irsee Monastery, Germany, July 2009
- T. Anstett, F. Leymann, R. Mietzner and S. Strauch: Towards BPEL in the Cloud: Exploiting Different Delivery Models for the Execution of Business Processes. International Workshop on Cloud Services (IWCS2009) in conjunction with the 7th IEEE International Conference on Web Services (ICWS 2009), July 2009.
- M. Vasko, E. Oberortner and S. Dustdar: Collaborative Modeling of Web Applications for Various Stakeholders. 5th Model-Driven Web Engineering Workshop in conjunction with ICWE 2009, June 2009.
- F. Daniel, F. Casati, V. D'Andrea, S. Strauch, D. Schumm, F. Leymann, E. Mulo, U. Zdun, S. Dustdar, S. Sebahi, F. de Marchi and M. Hacid: Business Compliance Governance in Service-Oriented Architectures. Proceedings of AINA'09, IEEE Press, May 2009. pdf
- E. Oberortner, U. Zdun, and S. Dustdar: Tailoring a Model-Driven Quality-of-Service DSL for Various Stakeholders. Workshop on Modeling in Software Engineering (MiSE) at ICSE 2009. pdf
- D. Schumm, D. Karastoyanova, F. Leymann and J. Nitzsche: On Visualizing and Modelling BPEL with BPMN. Proceedings of the 4th International Workshop on Workflow Management (ICWM2009), IEEE Computer Society, May 2009.
- N. Kokash and F. Arbab: Applying Reo to Service Coordination in Long-Running Business Transactions. Proceedings of the ACM Symposium on Applied Computing (SAC), 2009, pp. 318-319.
- B. Wetzstein, S. Strauch and F. Leymann: Measuring Performance Metrics of WS-BPEL Service Compositions. In: Proceedings of the Fifth International Conference on Networking and Services (ICNS 2009), Valencia, Spain, April 20-25, 2009.
- G. R. Gangadharan, S. Dustdar, V. D'Andrea et al.: LASS - License Aware Service Selection: Methodology and Framework. ICSOC, 2008.
- C. Mayr, U. Zdun and S. Dustdar: Model-Driven Integration and Management of Data Access Objects in Process-Driven SOAs. ServiceWave, 2008.
- E. Oberortner, U. Zdun and S. Dustdar: Domain-specific Languages for Service-oriented Architectures: An Explorative Study. ServiceWave, 2008.
- F. Arbab, N. Kokash and M. Sun: Towards Using Reo for Compliance-aware Business Process Modelling. In: Proceedings of the International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISOLA'08), vol. 17 of CCIS, Springer, 2008.
- W.-J. v/d Heuvel, M. Papazoglou and B. Orriens: On the risk management and auditing of SOA based business processes. ISoLA, 2008.

Book Chapters

- H. Tran, T. Holmes, U. Zdun and S. Dustdar: Chapter 2: Modeling Process-Driven SOAs – a View-Based Approach. Information Science Reference, Eds: Jorge Cardoso, and Wil van der Aalst, April 2009.

Journal Papers

- H. Tran, U. Zdun and S. Dustdar: VbTrace: Using View-based and Model-driven Development to Support Traceability in Process-driven SOAs. Software & System Modeling, Springer, DOI: 10.1007/s10270-009-0137-0, 2009.

- M. Strembeck and U. Zdun: An Approach for the Systematic Development of Domain-specific Languages. In: Software - Practice and Experience, Wiley InterScience, 2009.
- J. Kramer, Z. Taheri and W. Hutten: Service Oriented Architecture voor interne beheersing. In: PricewaterhouseCoopers Spotlight, volume 3, 2009.

Talk Presentations

- N. Kokash: Compliant Business Process Specifications to Code. iFM Workshop on Formal Methods and SOA (FMSOA), 2009.
- M. P. Papazoglou: Compliance Requirements for Business-process driven SOAs. Keynote address, The IFIP World Computer Congress WCC, 2008.
- F. Arbab and Natallia Kokash: Formal Behavioral Models and Compliance Analysis for Service Oriented Systems. FMCO, 2008.
- S. Dustdar: Towards Compliance-driven Models, Languages, and Architectures for Service-oriented Computing. FMCO, 2008.

Other Publications

- F. Arbab, N. Kokash, Z. Maraikar: Design Time Service Composition with Reo Coordination Tools. NESSI Open Framework - Reference Architecture (NEXOF-RA) - Service Composition, Brussels, Belgium, 2008.
- A. Birukou, V. D'Andrea, N. Kokash: Collaborative Web Service Discovery. NESSI Open Framework - Reference Architecture (NEXOF-RA) - Service Discovery, Brussels, Belgium, 2008.

DEPLOY

The overall aim of the EC Information and Communication Technologies FP7 DEPLOY Project is to make major advances in engineering methods for dependable systems through the deployment of formal engineering methods. Formal engineering methods enable greater mastery of complexity than found in traditional software engineering processes. It is the central role played by mechanically-analysed formal models throughout the system development flow that enables mastery of complexity.

As well as leading to big improvements in system dependability, greater mastery of complexity also leads to greater productivity by reducing the expensive test-debug-rework cycle and by facilitating increased reuse of software.

The work of the project will be driven by the tasks of achieving and evaluating industrial take-up, initially by DEPLOY's industrial partners, of DEPLOY's methods and tools, together with the necessary further research on methods and tools.

Deliverables:

- D12.1 Report on creation of DEPLOY infrastructure
- D14.1 Electronic Newsletter N°1 to N°8
- D14.9 Collaboration Plan
- D15.1 Initial Dissemination / Exploitation Plan
- JD1 Report on Knowledge Transfer
- D9.1 Model construction tools and analysis tools
- D11.1 Measurement Methodology Guide
- D10.1 Teaching material
- D10.2 Design and Implementation of technology transfer plateform
- D11.2 Initial Data collection framework
- D14.10 Collaboration Report
- D15.2 Year 1 Annual Dissemination / Exploitation Report
- D6.1 Advances in Methodological WPs
- D2.1 Pilot Deployment in Transportation
- D10.3 Initial port of technology transfer material
- D1.1 Report on pilot deployment in automotive sector
- D3.1 Report on pilot deployment in space sector
- D4.1 Report on pilot deployment in business information sector
- D5.1 Report on pilot deployment in pervasive telecom
- D9.2 Model construction tools and analysis tools II
- D14.11 Collaboration Report
- D15.3 Year 2 Annual Dissemination / Exploitation Report
- D14.14 Proceedings of the First International Workshop

Publications:

- Olszewska, Marta and Sere, Kaisa (2010) Specification Metrics for Event-B Developments. In: CONQUEST 2010, 20-22 September 2010, Dresden, Germany.
- Sorge, Jennifer and Poppleton, Michael and Butler, Michael (2010) A Basis for Feature-Oriented Modelling in Event-B. Springer, ABZ2010.
- Silva, Renato and Pascal, Carine and Hoang, T. Son and Butler, Michael (2010) Decomposition Tool for Event-B. In: Workshop on Tool Building in Formal Methods - ABZ Conference. (In Press)
- Olszewska, Marta and Sere, Kaisa (2010) Towards Event-B Specification Metrics. In: First Deploy Technical Workshop, 21-23 October 2009, Aix-en-Provence, France.
- Bryans, Jeremy W. and Fitzgerald, John S. and Romanovsky, Alexander and Roth, A. (2010) Patterns for Modelling Time and Consistency in Business Information Systems. In: 15th IEEE International Conference on Engineering of Complex Computer Systems. Oxford, UK. March, 2010. IEEE Computer Society.
- Iliasov, Alexei and Troubitsyna, Elena and Laibinis, Linas and Romanovsky, Alexander and Varpaaniemi, Kimmo and Väisänen, Pauli and Ilic, Dubravka and Latvala, Timo (2010) Verifying Mode Consistency for On-Board Satellite Software. In: The 29th International Conference on Computer Safety, Reliability and Security (Safecomp 2010), 14 - 17 September 2010, Vienna, Austria. Item availablity restricted.
- Lopatkin, Ilya and Iliasov, Alexei and Romanovsky, Alexander (2010) On Fault Tolerance Reuse during Refinement. In: 2nd International Workshop on Software Engineering for Resilient Systems, April 13-16, 2010, Birkbeck College (London, United Kingdom).
- Bicarregui, J. C. and Fitzgerald, J. S. and Larsen, P. G. and Woodcock, J. C. P. (2009) Industrial Practice in Formal Methods: a Review. FM 2009: Formal Methods, LNCS 5 . (In Press)
- Iliasov, Alexei and Dotti, Fernando and Romanovsky, Alexander (2009) Structuring Specifications with Modes. In: Fourth Latin-American Symposium on Dependable Computing (LADC), September 1-4, 2009, Brazil.
- Samia, Mireille and Leuschel, Michael (2009) Pie Tree Visualization. In: Proceedings SEKE 2009. Knowledge Systems Institute Graduate School, pp. 400-405. Item availablity restricted.
- Said, Mar Yah and Butler, Michael and Snook, Colin (2009) Language and Tool Support for Class and State Machine Refinement in UML-B. [Rodin Archive]
- Bryans, Jeremy W. and Fitzgerald, John S. and Romanovsky, Alexander and Roth, Andreas (2009) Formal Modelling and Analysis of Business Information Applications with Fault Tolerant Middleware. Proceedings 14th IEEE International Conference on Engineering of Complex Computer Systems ICECCS 2009. . pp. 68-77.
- Dotti, Fernando and Iliasov, Alexei and Riberiro, Leila and Romanovsky, Alexander (2009) Modal Systems: Specification, Refinement and Realisation. In: International Conference on Formal Engineering Methods - ICFEM 09 , December 9 -12, 2009, Rio de Janeiro, Brazil.
- Iliasov, Alexei and Arief, Budi and Romanovsky, Alexander (2009) Step-wise Development of Resilient Ambient Campus Scenarios. In: Methods, Models and Tools for Fault Tolerance. Springer, pp. 304-330.

- Wieczorek, S. and Kozyura, V. and Roth, A. and Leuschel, Michael and Bendisposto, Jens and Plagge, Daniel and Schieferdecker, I. (2009) Applying Model Checking to Generate Model-based Integration Tests from Choreography Models. 21st IFIP Int. Conference on Testing of Communicating Systems and the 9th Int. Workshop on Formal Approaches to Testing of Software TESTCOM/FATES 2009 . (In Press)
- Wieczorek, S. and Roth, A. and Stefanescu, A. and Kozyura, V. and Charfi, A. and Kraft, F. M. and Schieferdecker, I. (2009) Viewpoints for Modeling Choreographies in Service-Oriented Architectures. Joint Working IEEE/IFIP Conference on Software Architecture 2009 & European Conference on Software Architecture 2009 . (In Press)
- Bryans, Jeremy W. and Fitzgerald, John S. and Alexander, Romanovsky and Andreas, Roth (2008) Formal Modelling and Analysis of Business Information Applications with Fault Tolerant Middleware. Newcastle University, Newcastle University TR, accepted to appear in Proceedings of ICECCS 2009. (In Press)
- Wieczorek, S. and Roth, A. and Stefanescu, A. and Charfi, A. (2008) Precise Steps for Choreography Modeling for SOA Validation and Verification. In: Proceedings of the Fourth IEEE International Symposium on Service-Oriented System Engineering, Dec 18-19, Jhongli, Taiwan.
- Gorbenko, Anatoliy and Kharchenko, Vyacheslav and Romanovsky, Alexander and Mikhaylichenko, Alexey (2008) Experimenting with Exception Propagation Mechanisms in Service-Oriented Architectures. In: 4th International Workshop on Exception Handling (WEH.08), at ACM SIGSOFT 2008/FSE 16, 14 November, 2008, Atlanta, USA. (In Press)
- Butler, Michael (2008) Incremental Design of Distributed Systems with Event-B. In: Marktoberdorf Summer School 2008 Lecture Notes. IoS. (In Press)
- Damchoom, Kriangsak and Butler, Michael and Abrial, Jean-Raymond (2008) Modelling and proof of a Tree-structured File System. In: ICFEM 2008. (In Press)
- Ponsard, Christophe and Dallons, Gautier and Philippe, Massonet (2008) From Rigorous Requirements Engineering to Formal System Design of Safety-Critical Systems. ERCIM News (75). pp. 22-23. ISSN ISSN 0926-4981
- Romanovsky, Alexander (2008) DEPLOY: Industrial Deployment of Advanced System Engineering Methods for High Productivity and Dependability. ERCIM News (74). pp. 54-55.
- Hallerstede, Stefan (2008) On the Purpose of Event-B Proof Obligations. In: ABZ2008, September 2008, London. (In Press)
- Fitzgerald, John S and Bryans, Jeremy W and Greathead, David and Jones, Clifff B and Payne, Richard (2008) Animation-based Validation of a Formal Model of Dynamic Virtual Organisations. Electronic Workshops in Computing,
- Macedo, Hugo Daniel and Larsen, Peter Gorm and Fitzgerald, John (2008) Incremental Development of a Distributed Real-Time Model of a Cardiac Pacing System using VDM. In: Formal Methods 2008, 28-30 May, 2008, Turku, Finland.
- lecomte, thierry and servat, thierry (2007) Formal Methods in Safety-Critical Railway Systems. In: 10th Brasilian Symposium on Formal Methods, 29-31 August 2007, Ouro Preto (Brazil).
- Michael, Jastram and Linas, Laibinis and Felix, Loesch and Manuel, Mazzara, eds. Proceedings of the First Deploy Technical Workshop. Newcastle University.
- Damchoom, Kriangsak and Butler, Michael Applying Event and Machine Decomposition to a Flash-Based Filestore in Event-B. In: SBMF 2009, 19-21 August 2009, Gramado, Brazil

- Mazzara, Manuel Deriving Specifications of Dependable Systems: toward a Method. EWDC 2009 .
- Mazzara, Manuel and Bhattacharyya, Anirban On Modelling and Analysis of Dynamic Reconfiguration of Dependable Real-Time Systems. Technical Report. Newcastle University.
- Mazzara, Manuel and Bhattacharyya, Anirban On Modelling and Analysis of Dynamic Reconfiguration of Dependable Real-Time Systems. Proceedings of DEPEND 2010 . (In Press)
- Mazzara, Manuel and Dragoni, Nicola A Formal Semantics for the WS-BPEL Recovery Framework: The pi-Calculus Way. WSFM 2009 . (In Press)
- Michele, Mazzucco and Manuel , Mazzara and Nicola, Dragoni Design of QoS-aware Provisioning Systems. In: 4th Nordic Workshop on Dependability and Security (NODES 2010), Copenhagen, Denmark. (In Press)

DIVA

DiVA will provide a new tool-supported methodology with an integrated framework for managing dynamic variability in adaptive systems. This goal will be addressed by combining aspect-oriented and model-driven techniques in an innovative way.

Deliverables:

- D1.2. Framework for Identifying and Modelling of Dynamic Variability in Requirements (2010).
- D2.1. Transformation Framework SINTEF, Project deliverable. (2010).
- D3.2. Reference Architecture DiVA Consortium. (2009).
- D6.1 Case Study Specification and Requirements EC. (2009).
- D1.1. Survey and Evaluation Document of the Requirements Engineering for Dynamic Variability Lancaster University, DiVA project, technical report. (2009).
- D4.1. Survey and evaluation of approaches for the adaptation reasoning framework The DiVA project, technical report. (2008).
- D3.1. Survey and evaluation of approaches for runtime variability management INRIA.

Pubblications:

- Morin, B, Fleurey, F, Barais, O, and Jézéquel, J (2010). Aspect-Oriented Modeling to Support Dynamic Adaptation, In: Forum Demo at AOSD'10.
- Acher, M, Collet, P, Fleurey, F, Lahire, P, Moisan, S, and Rigault, J (ed.) (2009). Modeling Context and Dynamic Adaptations with Feature Models, Denver, CO, MODELS 2009. Models@run.time.
- Ali, R, Chitchyan, R, and Giorgini, P (2009). Context for Goal-level Product Line Derivation, In: In Proceedings of 3rd International Workshop on Dynamic Software Product Lines (DSPL09) co-located with the 13th International Software Product Line Conference (SPLC09).
- Ali, R, Yu, Y, Chitchyan, R, Nhlabatsi, A, and Giorgini, P (2009). , In: 3rd International Workshop on Software Product Management (IWSPM09), In conjunction with 17th IEEE International Requirements Engineering Conference (RE09).
- Alves, V, Schneider, D, Becker, M, and Bencomo, N (2009). Comparitive Study of Variability Management in Software Product Lines and Runtime Adaptable Systems, In: Proceedings of the 3rd International workshop on Variability Modelling of Software-intensive Systems, VaMOS 2009.
- Bencomo, N (2009). On the Use of Software Models during Software Execution, In: Proceedings of the Workshop on Modeling in Software Engineering MISE'2009, at ICSE'09, 2009 .
- Bencomo, N and Blair, G (2009). , Springer. LNCS Hot Topics .

- Cetina, C, Haugen, O, Zhang, X, Fleurey, F, and Pelechano, V (2009). Strategies for Variability Transformation at Run-time, In: SPLC 2009 , vol. Proceedings, SPLC.
- Fleurey, F and Solberg, A (2009). A Domain Specific Modeling Language supporting Specification, Simulation and Execution of Dynamic Adaptive Systems. MoDELS/UML, Proceedings.
- Grønmo, R and Møller-Pedersen, B (2009). From Sequence Diagrams to State Machines by Graph Transformation, International Conference on Model Transformation.
- Klein, J, Kienzle, J, Morin, B, and Jézéquel, J (2009). Aspect Model Unweaving, MODELS'09: ACM/IEEE 12th International Conference on Model Driven Engineering Languages and Systems.
- Morin, B, Barais, O, Jezequel, J, Fleurey, F, and Solberg, A (2009), IEEE Computer, pages 46--53.
- Morin, B, Barais, O, Jézéquel, J, Fleurey, F, and Solberg, A (2009). Model@Run.Time to Support Dynamic Adaptation, IEEE Computer, 42(10): 44-51.
- Morin, B, Barais, O, Nain, G, and Jézéquel, J (2009). Taming Dynamically Adaptive Systems Using Models and Aspects , In: ICSE'09: 31st Internation Conference on Software Engineering, Vancouver, Canada.
- Morin, B, Ledoux, T, Ben Hassine, M, Chauvel, F, Barais, O, and Jézéquel, J (2009). Unifying Runtime Adaptation and Design Evolution, In: CIT'09: IEEE 9th International Conference on Computer and Information Technology.
- Morin, B, Nain, G, Barais, O, and Jézéquel, J (2009). Leveraging Models From Design-time to Runtime. A Live Demo, In: Model@Run.Time workshop at MODELS'09.
- Sardinha, A, Chitchyan, R, Weston, N, Greenwood, P, and Rashid, A (2009). EA-Analyzer: Automating Conflict Detection in Aspect-Oriented Requirements, In: Automated Software Engineering (ASD 2009), pp. 530-534.
- Weston, N, Chitchyan, R, and Rashid, A (2009). A framework for constructing semantically composable feature models from natural language requirements, In: 13th International Conference on Software Product Lines (SPLC 2009), pp. 211-220, ACM.
- Weston, N, Chitchyan, R, and Rasihd, A (2009). Formal semantic conflict detection in aspect-oriented requirements, Requirements Engineering Journal, 14(4):247-268.
- Baudry, B and Muñoz, F (2008). Validation challenges in model composition: The case of adaptive systems, In the proceedings of: First International Workshop on Challenges in Model-Driven Software Engineering, 2008 (held in conjunction with the MoDELS'08).
- Bencomo, N (2008). Supporting the Modelling and Generation of Reflective Middleware Families and Applications using Dynamic Variability, PhD Thesis, Computing Department, Lancaster University.
- Bencomo, N, Sawyer, P, Blair, G, and Grace, P (2008). Dynamically Adaptive Systems are Product Lines too: Using Model-Driven Techniques to Capture Dynamic Variability of Adaptive Systems, In: Proceedings of the 2nd International Workshop on Dynamic Software Product Lines (DSPL 2008) .
- Fleurey, F, Dehlen, V, Bencomo, N, Morin, B, and Jézéquel, J (2008). Modeling and Validating Dynamic Adaptation, Models@Runtime Workshop in conjunction with MODELS 2008.
- Grace, P, Blair, G, Flores, C, and Bencomo, N (2008). Engineering Complex Adaptations in Highly Heterogeneous Distributed Systems, In: Proceedings of Autonomics 2008 .

- Morin, B, Barais, O, and Jézéquel, J (2008). K@RT: An Aspect-Oriented and Model-Oriented Framework for Dynamic Software Product Lines, Models@run.time workshop (at MoDELS'08).
- Morin, B, Fleurey, F, Bencomo, N, Jézéquel, J, Solberg, A, Delhen, V, and Blair, G (2008). An Aspect-Oriented and Model-Driven Approach for Managing Dynamic Variability, In: MODELS'08: Proceedings of ACM/IEEE 11th International Conference on Model Driven Engineering Languages and Systems, vol. 5301/2008, pp. 782-796. LNCS.
- Morin, B, Vanwormhoudt, G, Lahire, P, Gaignard, A, Barais, O, and Jézéquel, J (2008). Managing Variability Complexity in Aspect-Oriented Modeling In: MODELS'08: Proceedings of ACM/IEEE 11th International Conference on Model Driven Engineering Languages and Systems, vol. 5301, pp. 797-812, LNCS.
- Munoz, F, Baudry, B, and Barais, O (2008). Improving Maintenance in AOP Through an Interaction Specification Framework, In the proceedings of the 24th International conference on Software Maintentance, ICSM08.
- Perrouin, G, Chauvel, F, DeAntoni, J, and Jézéquel, J (2008). Modeling the Variability Space of Self-Adaptive Applications, In the prodeedings of: 2nd International Workshop on Dynamic Software Product Lines at SPLC'08.
- Weston, N, Chitchyan, R, and Rashid, A (2008). A Formal Approach to Semantic Composition of Aspect-Oriented Requirements, In: In Proceedings of the 2008 16th IEEE international Requirements Engineering Conference, pp. 173-182, IEEE Computer Society.

FAST

FAST aims at providing an innovative visual programming environment that will facilitate the development of next-generation composite user interfaces. It constitutes a novel approach to application composition and business process definition from a top-down user-centric perspective.

Deliverables:

- State of the art in gadgets: semantics, visual design, SWS and catalogs, Feb. 2010
- Ontology and conceptual model for the semantic characterization of complex gadgets, Feb. 2010
- FAST requirements specification, Feb. 2010
- Mediation amongst ontologies: Application to the FAST ontology, Feb. 2010
- FAST Complex Gadget Architecture, Feb. 2010
- The Gadget Development Method, Feb. 2010
- The Gadget Visual Storyboard, Feb. 2010
- The Gadget Visual Storyboard: User Manual, Feb. 2010
- Complex Gadget Studio: Visual tools for the connection of complex gadgets to back-end resources, Feb. 2010
- Complex Gadget Studio: Visual tools for the connection of complex gadgets to back-end resources: User Guide, Feb. 2010
- Mechanisms for gadget-service connections and gadget functionality, Feb. 2010
- User manual of the FAST Catalogue, Feb. 2010
- Prototype semantic catalogue for screen-flow gadgets and back-end services, Feb. 2010
- Scenario Definition, Feb. 2010
- Scenario Design, Feb. 2010
- Implementation of the Scenario, Aug. 2009
- Evaluation of the Implementation, Feb. 2010
- Market and Competition Analysis, Dec. 2008
- Exploitation Strategy, Feb. 2010
- Market Awareness Strategy, Feb. 2010
- Collaboration plan, Aug. 2008

Publications:

- Buechel, B., Janner, T., Schroth, C., Hoyer, V. . Enterprise Mashup vs. Service Composition: What fits to reach the next stage in End-User Development?. Proceedings of the 5th Conference of Professional Knowledge Management, Gesellschaft für Informatik (GI). 2009
- Hoyer, V., Stanoevska-Slabeva . Design Principles of Enterprise Mashups. Proceedings of the 5th Conference of Professional Knowledge Management, Gesellschaft für Informatik (GI). 2009

- Hoyer, V., Janner, T., Schroth, C., Delchev, I., Urmetzer, F. FAST Platform: A Concept for user-centric, enterprise class Mashups. Demo Session, in: Proceedings of the 5th Conference of Professional Knowledge Management, Gesellschaft für Informatik (GI). 2009
- Rafael Fernández, David Lizcano, Javier A. López, Marcos Reyes, and Javier Soriano (in alph. order), Tackling Interoperability in Composite Applications from an Enterprise Mash-up Perspective. Proceedings of the 14th International Conference on Concurrent Enterprising, ICE 2008, Lisboa, Portugal, June 23-25, 2008.
- Juan J. Hierro, Till Janner, David Lizcano, Marcos Reyes, Christoph Schroth, Javier Soriano (in alph. Order), Enhancing User-Service Interaction Through a Global User-Centric Approach to SOA. The Fourth International Conference on Networking and Services, ICNS'08. Gosier, Guadalupe, 16-21 March 2008. Published in IEEE Computer Society Press.ISBN 0-7695-3094-X
- David Lizcano, Javier Soriano, Marcos Reyes and Juan J. Hierro, EzWeb/FAST: Reporting on a Successful Mashup-based Solution for Developing and Deploying Composite Applications in the Upcoming Web of Services. ACM Proceedings of the 10th International Conference on Information Integration and Web-based Applications & Services, iiWAS 2008, Linz, Austria, November 24-26, 2008. ACM ISBN 978-1-60558-349-5
- Hoyer, V., Stanoevska-Slabeva, K., Janner, T., & Schroth, C. (2008). Enterprise Mashups: Design Principles towards the Long Tail of User Needs. In Proceedings of the 2008 IEEE International Conference on Services Computing (SCC 2008) (pp. 2): IEEE Computer Society. Retrieved 2009-03-24, from http://www.alexandria.unisg.ch/publications/44891.
- David Lizcano, Miguel Jiménez, Javier Soriano, José M. Cantera, Marcos Reyes, Juan J. Hierro, Francisco Garijo and Nikolaos Tsouroulas, Leveraging the Upcoming Internet of Services through an Open User-Service Front-end Framework. In Petri Mähönen, Klaus Pohl, Thierry Priol (editors): Towards a Service-Based Internet, Springer Lecture Notes in Computer Science, nº 5377, ISSN 0302-9743, ISBN-10 3-540-89896-4, ISBN-13 978-3-540-89896-2, ServiceWave 2008 Conference, Madrid, Spain.
- David Lizcano, Javier Soriano, Marcos Reyes, and Juan J. Hierro, EzWeb/FAST: Reporting on a Successful Mashup-based Solution for Developing and Deploying Composite Applications in the Upcoming "Ubiquitous SOA", UBICOMM 2008, Valencia, Spain, 2008. ISBN 978-0-7695-3367-4
- Hoyer, V., Stanoevska-Slabeva, K.,. The Changing Role of IT Departments in Enterprise Mashups Environments, Proceedings of the 2nd International Workshop on "Web APIs and Services Mashups", 2008
- Hoyer, V., Fischer, M. *Market Overview of Enterprise Mashup Tools*. In Lecture Notes in Computer Science (LNCS), Vol. 5364, pp. 708-721, 2008

IRMOS

Design, develop and validate 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.

Deliverables:

- D7.3.2 v1.0 "Final Version of Flow Control Architecture"
- D3.1.3 V1.0 "Updated version of IRMOS Overall Architecture"
- D9.5.2 V1.0 " Updated version of Collaboration Plan"
- D6.4.2 V1.0 "Final version of Real-Time Architecture of the Excecution Environment"
- D7.5.1 V1.0"ISONI Proof of Concept with Limited Functionality"
- D7.4.1 V1.0 "Initial Version of Path Supervision Architecture"
- D7.3.1 v1.0 "Inital Version of Flow Control Architecture"
- D4.2.1 V1.0 "Interface definition to the IRMOS SOI"
- D7.2.1 V1.0 "Inital Version of Path Manager Architecture"
- D4.1.1 V1.0 "Definition and implementation of the three scenarios and its real time requirements"
- D5.1.1 V1.0 "Models of Real-time Applications on Service Oriented Infrastructures"
- D3.1.2 V1.0 "IRMOS Overall Architecture"
- D6.4.1 V1.0 "Initial version of Real-Time Architecture of Excecution Environment"
- D7.1.1 V1.0 "ISONI Addressing Schemes"
- D2.4.1 V1.0 “Security Requirements Analysis Report”
- D3.1.1 V1.0 "Prelimenary IRMOS overall archietecture"
- D9.5.1 V1.0 "Initial Version of Collaboration Plan"
- D2.3.1 V1.0 "State of the Art on IRMOS Technologies"
- D2.1.1 V1.0 "Initial version of Requirements Analysis Report"

Publications:

- Exception-Based Management of Timing Constraints Violations for Soft Real-Time Applications 5th International Workshop on Operating Systems Platforms for Embedded Real-Time Applications (OSPERT 2009), Dublin, Ireland, June 2009 16 April 10 Tommaso Cucinotta, Dario Faggioli, Alessandro Evangelista

- QoS Control for Pipelines of Tasks Using Multiple Resources IEEE Transactions on Computers, Vol. 59, no. 3, March 2010, IEEE Computer Society Digital Library 31 March 10 Tommaso Cucinotta, Luigi Palopoli
- Sporadic Server Revisited 25th ACM Symposium On Applied Computing, Sierre, Switzerland, March 2010 22 March 10 Dario Faggioli, Marko Bertogna, Fabio Checconi
- Welcome - The Real-Time QoS Enabled Cloud 2nd EU ICT IRMOS Public Seminar -The Real-Time QoS Enabled Cloud 18 March 10 Karsten Oberle, Alcatel-Lucent, Bell Labs
- Overview - The Real-Time QoS Enabled Cloud 2nd EU ICT IRMOS Public Seminar -The Real-Time QoS Enabled Cloud 18 March 10 Dimosthenis Kyriazis, National Technical University of Athens
- Modelling Interactive Real-time Applications on Service Oriented Infrastructures 2nd EU ICT IRMOS Public Seminar -The Real-Time QoS Enabled Cloud 18 March 10 Luís Costa, SINTEF
- Modelling Interactive Real-time Applications on Service Oriented Infrastructures - Part II 2nd EU ICT IRMOS Public Seminar -The Real-Time QoS Enabled Cloud 18 March 10 Zlatko Zlatev, IT Innovation
- Advancements in Real-Time Virtualized Computing 2nd EU ICT IRMOS Public Seminar -The Real-Time QoS Enabled Cloud 18 March 10 Tommaso Cucinotta, Scuola Superiore Sant'Anna
- How Host Virtualization contributes to Green IT 2nd EU ICT IRMOS Public Seminar -The Real-Time QoS Enabled Cloud 18 March 10 Dominik Lamp, University of Stuttgart
- Approaches to network virtualization and their shortcomings 2nd EU ICT IRMOS Public Seminar -The Real-Time QoS Enabled Cloud 18 March 10 Manuel Stein, Alcatel-Lucent, Bell Labs
- IRMOS Newsletter Issue N°3 15 February 10 IRMOS Consortium
- Advance Reservations for Distributed Real-TimeWorkflows with Probabilistic Service Guarantees IEEE International Conference on Service-Oriented Computing and Applications (SOCA 2009), December 2009, Taipei, Taiwan 14 December 09 Tommaso Cucinotta (SSSA), Kleopatra Konstanteli (NTUA), Theodora Varvarigou (NTUA)
- Real-time QoS Management IRMOS Project Seminar - Real-time QoS Management in Clouds 10 December 09 Mike Boniface (IT Innovation)
- Modelling Interactive Real-time Applications on Service Oriented Infrastructure IRMOS Project Seminar - Real-time QoS Management in Clouds 10 December 09 Zlatko Zlatev (IT Innovation)
- Mapping Rules Creation in SOIs IRMOS Project Seminar - Real-time QoS Management in Clouds 10 December 09 George Kousiouris (NTUA)
- Challenges of Data Storage for Soft Real-time Applications (in SOIs) IRMOS Project Seminar - Real-time QoS Management in Clouds 10 December 09 Darren Golbourn (Xyratex)
- Virtualization of Network and Computational Resources IRMOS Project Seminar - Real-time QoS Management in Clouds 10 December 09 Dominik Lamp (USTUTT)
- Advance Reservations for Distributed Real-time Workflows with Probabilistic Service Guarantees IRMOS Project Seminar - Real-time QoS Management in Clouds 10 December 09 Tommaso Cucinotta (SSSA), Kleopatra Konstanteli, Dora Varvarigou (NTUA)
- Hierarchical Multiprocessor CPU Reservations for the Linux Kernel IRMOS Project Seminar - Real-time QoS Management in Clouds 10 December 09 Fabio Checconi, Tommaso Cucinotta, Dario Faggioli and Giuseppe Lipari (SSSA)

- IRMOS Application Blueprint - A Blueprint for Creating Interactive Applications for Virtualized Real-time Service Infrastructures 02 December 09 Ralf Einhorn, Lars Fürst (DTO), Michael Braitmaier, Dominik Lamp (USTUTT), Kleopatra Konstanteli, George Kousiouris, Andreas Menychtas, Dimosthenis Kyriazis (NTUA), Eduardo Oliveros (TID), Neil Loughran (SINTEF), Bassem I. Nasser (IT Inn)
- A SOAP Performance Comparison of different WSRF The International ACM Conference on Management of Emergent Digital EcoSystems (MEDES), Lyon (France) 27 October 09 Roland Kübert, Hai-Lang Thai, Axel Tenschert, HLRS
- Network Virtualization: The missing piece ICIN 2009 – Beyond the bit pipes, Bordeaux, France 26 October 09 Karsten Oberle, Manuel Stein, Thomas Voith, Marcus Kessler, ALUD; Dominik Lamp, Sören Berger, USTUTT
- The Wizard of OS: a Heartbeat for Legacy Multimedia Applications 7th IEEE Workshop on Embedded Systems for Real-Time Multimedia, Grenoble (ESTImedia 2009), Grenoble, October 2009 15 October 09 Tommaso Cucinotta, Luca Abeni, Luigi Palopoli, Fabio Checconi
- An Intelligent Service Oriented Infrastructure supporting Real-time applications 2nd Japan-EU Symposium on the "Future Internet", Tokyo, Japan 14 October 09 Karsten Oberle (ALUD)
- European Collaborative Research Projects - Xyratex and the IRMOS Project Collaborate to Innovate programme of the Solent Synergy Business Development Group 06 October 09 Stuart Smithson (XY)
- Real-Time Issues in Live Migration of Virtual Machines 4th Workshop on Virtualization in High Performance Computing (VHPC'09), Delft (Netherlands) 25 September 09 Fabio Checconi, Tommaso Cucinotta (SSSA), Manuel Stein (ALUD)
- Multi-level feedback control for Quality of Service Management 14th IEEE International Conference on Emerging Technologies and Factory Automation, Palma de Mallorca (ETFA 2009), Spain, September 2009 22 September 09 Tommaso Cucinotta, Giuseppe Lipari, Luigi Palopoli, Luca Abeni, Rodrigo Santos
- A Real-time Service-Oriented Architecture for Industrial Automation IEEE Transactions on Industrial Informatics, Vol. 5, n. 3, August 2009 01 August 09 Tommaso Cucinotta, Antonio Mancina, Gaetano Anastasi, Giuseppe Lipari, Leonardo Mangeruca, Roberto Checcozzo, Fulvio Rusinà
- Respecting temporal constraints in virtualised services 2nd IEEE International Workshop on Real-Time Service-Oriented Architecture and Applications (RTSOAA 2009), Seattle, Washington, July 2009 20 July 09 Tommaso Cucinotta, Gaetano Anastasi, Luca Abeni
- Real-time guarantees in flexible advance reservations 2nd IEEE International Workshop on Real-Time Service-Oriented Architecture and Applications (RTSOAA 2009), Seattle, Washington, July 2009 20 July 09 Kleopatra Kostanteli (NTUA), Dimosthenis Kyriazis (NTUA), Theodora Varvarigou (NTUA), Tommaso Cucinotta, Gaetano Anastasi
- ISONI Whitepaper V2.0 (V1.0 updated. Improvement of readabiblity compared to V1.0, no major improvement of content) ISONI Whitepaper 01 July 09 ALUD/USTUTT
- Hierarchical Multiprocessor CPU Reservations for the Linux Kernel 5th International Workshop on Operating Systems Platforms for Embedded Real-Time Applications (OSPERT 2009), Dublin, Ireland, June 2009 30 June 09 Fabio Checconi, Tommaso Cucinotta, Dario Faggioli, Giuseppe Lipari

- Network Virtualization: Towards a fully virtualized Service Infrastructure International Supercomputing Conference, ISC'09, Hamburg 23 June 09 M. Kessler, S. Braun, K. Oberle, ALUD; D. Lamp, USTUTT
- IRMOS Project Poster Internet of Services 2009, Brussels 10 June 09 IRMOS Consortium IRMOS Flyer V2 02 June 09 IRMOS Consortium
- Virtual Cities of Art and Immersive Experience in Geo-referenced Communities TiceMED - UBIQUITOUS LEARNING, Milano-Bicocca 28 May 09 Alessandro Mazzetti, Giunti Labs S.r.l.
- Towards Personal Ambient Learning. Blending Mobile, Location & Virtual Worlds based learning for a better Education Immersive Education Summit - London School of Economics in London 23 April 09 Fabrizio Cardinali
- Guaranteering QoS with Dynamic and Automated SLAs Nexof-RA 2nd open call for contributions 16 March 09 NTUA and other members of the IRMOS consortium IRMOS Newsletter Issue N°2 02 March 09 IRMOS Consortium
- Panel presentation in FISE session on "Social Benefits and Business Threats in System Customisation" Future Internet Assembly Meeting, Madrid 09 December 08 Mike Boniface, IT Innovation

M:CIUDAD

Service creation on the move, with the mobile device, for mobile users.

Imagine what kind of applications become possible when our mobile devices do not only present data but provide valuable information to other users. Suppose that you are able to create instant services with information, contents and knowledge with your mobile device and in your mobile device. And suppose that this knowledge can be used remotely by other users in a simple way, with their mobile devices. And now, imagine the amount of available knowledge of those services. Millions or perhaps billions of potential sources with valuable information: constantly updated, relevant to our instant interests and still context aware.

m:Ciudad, a service infrastructure

So, what are the required tools to let each user to become a service provider with a mobile device? In which way should the mobile platform behave to make it simple to use and efficient? How to reach this type of distributed, volatile services and their associated knowledge or information? And how to exploit the business opportunities that this new scenario brings about?

m:Ciudad faces and answers these questions providing user-friendly creation tools in the mobile, an optimised execution environment, a model for knowledge warehouse, a proposed specific searching engine and a set of business models for users, for service providers and for third parties, mainly SMEs.

m:Ciudad is a service architecture, a set of mobile tools and a platform to allow users to create focused, knowledge-based mobile micro-services, which are called m:Ciudad U+ Services.

Deliverables:

- mCiudad Scenarios and Use Case Analysis
- m:Ciudad User and System Requirements
- State of the Art in Service Description
- m:Ciudad Service Specification
- mCiudad Architecture Design

Publications:

- Context-aware Recommendations on Mobile Services: The m:Ciudad Approach. Andreas Emrich, Alexandra Chapko, Dirk Werth, 4th European Conference on Smart Sensing and Context (EuroSSC), 16-18 September 2009, University of Surrey, Guildford (UK)
- Distributed RESTful Web Services for Mobile Person-to-Person Collaboration, Ville Antila, Jani Mäntyjärvi, NGMAST 2009 (Next Generation Mobile Applications, Services and Technologies) Conference, 15-18 September 2009, Cardiff, Wales.

- Towards a Semantic Infrastructure for User Generated Mobile Services, Marcin Davies. European Semantic Web Conference 2009, 31 May - 4 June 2009, Heraklion, Greece.
- m:Ciudad - Unleashing mobile user-provided services, Diego Urdiales, Rafael de las Heras, Marcin Davies, Anna V. Zhdanova, Milla Immonen, Juhani Heinilä, Maribel Narganes, Benoît Christophe, Linas Maknavi_ius, Payam Barnaghi. WWW2009 event, 20-24 April 2009, Madrid, Spain.
- m:Ciudad: An Infrastructure for Creation and Sharing of End User Generated Microservices, Davies M., Gil G., Maknavicius L., Narganes M., Urdiales D., Zhdanova A.V. In Proceedings of the Poster and Demonstration Track at the 1st Future Internet Symposium, 28-30 September 2008, Vienna, Austria.

MANCOOSI

The Mancoosi project aims at developing tools for the system administrator. It pursues two main avenues:

1. Develop mechanisms that provide for rollbacks of failed upgrade attempts, allowing the system administrator to revert the system to the state before the upgrade
2. Develop better algorithms and tools to plan upgrade paths based on various information sources about software packages and on optimization criteria

Deliverables:

- D4.2 - First version of the Mancoosi specialised CUDF solver plugin for the modular platform manager, 01/02/2010.
- D2.2 - Instantiation of the metamodel on a widely used GNU/Linux distribution, 31/01/2010.
- D7.3 - Updated Collaboration Plan, 31/01/2010 (revised 07/04/2010).
- D6.3 - Dissemination Activity Report, 31/01/2010.
- D3.2 - First version of the DSL based on the model developed in WP2, 01/11/2009.
- D2.1 - Metamodel for Describing System Structure and State, 31/01/2009.
- D6.2 - Periodic Dissemination Report, 31/01/2009.
- D7.2 - Updated Collaboration Plan, 28/01/2009.
- D3.1 - Survey of the state of the art technologies for handling versioning, rollback and state snapshot in complex systems, 30/09/2008.
- D7.1 - Collaboration activities plan, 01/08/2008.
- D1.1 - Project presentation, 01/08/2008.

Publications:

- Autocompletion for Mashups, VLDB 2009, August 24-28, 2009, Lyon, France. By Ohad Greenshpan, Tova Milo and Neoklis Polyzotis.
- Querying DAG-shaped Execution Traces Through Views, WebDB, June 28, 2009, Rhode Island, USA. By Maya Ben-Ari, Tova Milo and Elad Verbin.
- Enforcing Type-Safe Linking using Inter-Package Relationships, JFLA 2010. January 30 - February 02, 2010, La Ciotat, France. By Mehdi Dogguy, Stéphane Glondu, Sylvain Le Gall, and Stefano Zacchiroli.
- Strong Dependencies between Software Components, ESEM 2009, October 15-16, 2009, Lake Buena Vista, Florida, USA. By Pietro Abate, Jaap Boender, Roberto di Cosmo, and Stefano Zacchiroli.
- Expressing Advanced User preferences in Component Installation, IWOCE 2009 International Workshop on Open Component Ecosystem, affiliated with ESEC/FSE 2009. August 24, 2009, Amsterdam, The Netherlands. By Ralf Treinen and Stefano Zacchiroli.
- Towards maintainer script modernization in FOSS distributions, IWOCE 2009 International Workshop on Open Component Ecosystem, affiliated with

ESEC/FSE 2009. August 24, 2009, Amsterdam, The Netherlands. By Davide Di Ruscio, Patrizio Pelliccione, Alfonso Pierantonio, and Stefano Zacchiroli.

- On Solving Boolean Multilevel Optimization Problems, IJCAI 2009, July 11-17 2009, Pasadena, California, USA. By Josep Argelich, Ines Lynce, and Joao Marques-Silva.
- Towards a Model Driven Approach to Upgrade Complex Software Systems, ENASE 2009, May 9-10 2009, Milan, Italy. By Antonio Cicchetti, Davide Di Ruscio, Patrizio Pelliccione, Alfonso Pierantonio, and Stefano Zacchiroli.
- Optimized Union of Non-disjoint Distributed Data Sets, EDBT 2009, March 23-26 2009, Saint-Petersburg, Russia. By Itay Dar, Tova Milo and Elad Verbin.
- CNF instances from the software package installation problem, RCRA 2008, Udine, Italy, December 12-13 2008. By Josep Argelich and Ines Lynce.
- Package upgrades in FOSS distributions: details and challenges, HotSWUp'08, October 20, 2008, Nashville, Tennessee, USA. By Roberto Di Cosmo, Paulo Trezentos, and Stefano Zacchiroli.
- Managing model conflicts in distributed development, MODELS 2008, September 2008, Toulouse, France. By Antonio Cicchetti, Davide Di Ruscio, and Alfonso Pierantonio.
- Automating Model Co-evolution in Model-Driven Engineering, EDOC 2008, September 2008, München, Germany. By Antonio Cicchetti, Davide Di Ruscio, Romina Eramo, and Alfonso Pierantonio.
- On SAT Technologies for dependency management and beyond, ASPL 2008, Limerick, September 2008. By Daniel Le Berre and Anne Parrain.
- Enriching Topic-Based Publish-Subscribe Systems with Related Content, SIGMOD'08. By Rubi Boim and Tova Milo.

MOST

The MOST project improves software engineering by leveraging ontology and reasoning technology. It will develop a seamless integration technology for ontologies into model-driven software development (MDSD), resulting in ontology driven software development (ODSD).

One major objective of the MOST approach is to leverage MDSD and other software development processes with ontology technology. This goal will be reached by implementing integration technology for ontologies into MDSD, resulting in ontology-driven software development (ODSD). This integration technology concerns all involved artefacts (ontology and modelling languages, models, tools), as well as the development processes. The artefact integration technology will include capabilities for managing software development and ontology-based validation and knowledge management in an integrated platform. For this purpose, the MOST approach will integrate ontologies and software technologies at three different levels. First, integration will be addressed at the level of modelling languages. Second, ontologies and models will be integrated (modelling level) so that ontologies can be used in models, and vice versa. Third, ontology-based queries will be integrated on the level of MDA transformations.

As a second major objective, the MOST approach creates an integrated ontology-driven, model-aware software process (ODSD). This development process will be guided by process guidance ontologies that formalize the rules, conditions and actions a software engineer has to conduct in specific development situations. Instead of specifying a workflow operationally, the ODSD will be driven by guidance ontologies that can formalize company-specific rules, conventions, and enactments. Hence, technology from the ontology engineering domain will benefit to the software process domain. Secondly, the guidance ontologies will guide the modelling engineer to change between ontology and modelling languages, i.e., to transform models to different, but more appropriate reasoning languages. Finally, to be able to guide software engineers through the development of integrated model chains, traceability information must be provided for the entire chains.

Deliverables:

- D8.1.c Proof of Concept release, 07/20/2010
- D8.0.c Processes, Principles and Selection Criteria behind PoC, 07/20/2010
- D2.1 IT Report: Runtime Service Composition, 06/30/2010
- D3.4 Guidance tools for language trans_formations, February 2010
- D7.2 Planning and organisation of dissemination activities, February 2010
- D7.5.3 Collaboration Plan, February 2010
- D8.4 Dissemination and exploitation plan, February 2010
- D12.1.2 NEXOF Roadmap v2: Construction Process, 01/18/2010
- D1.3 Report on transformation patterns, September 2009
- D2.3 Initial prototype of ontology-driven software process guidance system, September 2009
- D2.5.1 Ontology Services for Model-Driven Software Development, September 2009

- D3.3 Refinement of language hierarchy (input from use-cases) , September 2009
- D4.2 Report on Traceability Information, September 2009
- D8.2a Report on Proof of concept, 12/09/2009
- D8.0 b Proof of Concept Overall Process, 12/09/2009
- D1.2 Report on querying the combined metamodel, February 2009
- D2.2 Modeled Software Guidance/Engineering Processes and Systems, February 2009
- D3.2 Initial Prototype of Language Transformations, February 2009
- D3.6 Successful Stories and Potential Patterns on Applying Ontologies in Software Engineering, February 2009
- D4.1 Taxonomy and Definition of the Explicit Traceability Information, February 2009
- D7.1 Market analysis and preparation of marketing and standardization materials, February 2009
- D7.5.2 Collaboration Plan, February 2009
- D8.1.b Proof of concept, 12/09/2009
- D12.1 NEXOF Context, 04/02/2009
- D5.2.1 Open Architecture Specification Process - Programme 1, 10/17/2008
- D5.1.1 Open Architecture Specification Process – Definition, 10/17/2008
- D1.1 Report on the combined metamodel, September 2008
- D2.1 Requirements definition and design of ontology-driven software process guidance system, September 2008
- D3.1 Classification of language hierarchy and complexities, September 2008
- D5.1 Definition of the case study requirements, September 2008
- D6.1 Description of functional and non-functional requirements, September 2008
- D7.5.1 Collaboration Plan, September 2008
- D8.4 Dissemination and exploitation plan, September 2008

Publications:

- Christo Büurger, Sven Karol, and Christian Wende. Applying Attribute Grammars for Metamodel Semantics. International Workshop on Formalization of Modeling Languages (FML 2010) at ECOOP 2010, 2010.
- Cristo Büurger, Sven Karol, Christian Wende, and Uwe Aßmann. Reference Attribute Grammars for Metamodel Semantics. 3rd International Conference on Software Language Engineering (SLE 2010), 2010.
- Florian Heidenreich, Jendrik Johannes, Mirko Seifert, Michael Thiele, Christian Wende, and Claas Wilke. Integrating OCL and Textual Modelling Languages. In Jordi Cabot, Tony Clark, Manuel Clavel, and Martin Gogolla, editors, Proceedings of the Workshop on OCL and Textual Modelling (co-located with MODELS 2010), 2010.
- Henri Mhle and Christian Wende. Describing Role Models in Terms of Formal Concept Analysis. 8th International Conference on Formal Concept Analysis (ICFCA10), 2010. Guilin Qi, Qiu Ji, Je. Pan, and Jianfeng Du. Possdl - a possibilistic dl reasoner for uncertainty reasoning and inconsistency handling. In Proceedings of 7th the Extended Semantic Web Conference (ESWC2010), 2010.
- Jan Reimann, Mirko Seifert, and Uwe Aßmann. Role-Based Generic Model Refactoring. In Dorina C. Petriu, Nicolas Rouquette, and Øystein Haugen, editors, MoDELS (2), volume 6395 of Lecture Notes in Computer Science, pages 78-92. Springer, 2010.

- Yuan Ren, Kees van Deemter, and Je. Z. Pan. Charting the Potential of Description Logic for the Generation of Referring Expressions. In Proc. of 6th International Conference on Natural Language Generation(INLG-2010), 2010.
- Quentin Reul and Je. Z. Pan. KOSIMap: Use of Description Logic Reasoning to Align Heterogeneous Ontologies. In the Proc. of International Description Logic Workshop (DL2010), 2010.
- Mirko Seifert, Christian Wende, and Uwe Aßmann. Anticipating Unanticipated Tool Interoperability using Role Models. In Jean Bézivin, Richard Mark Soley, and Antonio Vallecillo, editors, MDI 2010, pages 52-60, New York, NY, USA, 2010. ACM.
- Edward Thomas, Je. Z. Pan, Stuart Taylor, and Yuan Ren. Lightweight Reasoning and the Web of Data for Web Science. In the Proc. of the 2nd International Conference on Web Science (WebSci 2010), 2010.
- Claas Wilke, Michael Thiele, and Christian Wende. Extending Variability for OCL Interpretation. Proceedings of the ACM/IEEE 13th International Conference on Model Driven Engineering Languages and Systems (MoD-ELS2010), Oslo, 2010.

NEXOF-RA

The overall ambition of the NEXOF-RA project is to build a Reference Architecture for the NESSI Open Framework - ranging from the infrastructure up to the interfaces with the end users - leveraging research in the area of service-based systems to consolidate and trigger innovation in service oriented economies. In particular; the aim of NEXOF-RA is to deliver:

- The NEXOF Reference Architecture. Following an Open Architecture Specification Process which will allow contributions from many sources also outside NEXOF-RA, - focusing on the “NESSI Framework” as defined in the NESSI Holistic Model.

- A Proof-of-Concept. This will be a set of software artifacts the project team will use to validate the key architectural choices made.

- The NEXOF Roadmap. This will consist in the roadmap for the implementation and adoption of the whole of NEXOF and will indicate some major research trends to be pursued by the wider community.

Deliverables:

- D10.3.b Assessment report, 09/10/2010
- D10.2-v2 Assessment Criteria, 09/10/2010
- D10.2-v1 Assessment Criteria, 03/08/2009
- D10.1 Scenarios and Requirements for Open Framework construction, 10/13/2008
- D6.3 v1.0 The NEXOF Reference Model, 07/20/2010
- D6.1 v1.1 NEXOF Reference Architecture Model, 04/02/2009
- D6.2 v1.0 NEXOF Reference Architecture Model, 03/21/2009
- D7.5b RA Specification, 07/20/2010
- D7.2c Definition of an architectural framework and principles, 07/20/2010
- D7.1 State of the art report, 12/15/2009
- D7.5 RA Specification (Executive Summary)
- D7.5 RA Specification (Pattern Compass)
- D7.2b Definition of an architectural framework and principles, 12/02/2009
- D7.4 RA Specification Sample, 04/07/2009
- D7.2 Design Pattern example and template, 11/25/2008
- D7.2 Definition of an architectural framework principles
- D7.1 State of the art report, 07/23/2008
- D7.2 Definition of an architectural framework and principles, 07/22/2008
- D9.2 2nd Report on Standardisation, 08/29/2010
- D9.2 Report on Standardisation, 08/05/2009
- D9.1 Relevant Standardisation Bodies and Standards for NEXOF, 09/10/2008
- D4.1c Non-Functional Aspects, 08/12/2010
- D3.1 Resource Infrastructure in the NEXOF Reference Architecture, 07/20/2010

- D2.1 Service Centric System Architecture Contributions to Model and Architecture, 07/20/2010
- D1.1c Advanced User-Service Interactions, 07/20/2010
- D1.1b Advanced User Service Interactions contribution, 12/15/2009
- D2.1 Service Centric System Architecture Contributions to Model and Architecture 08/05/2009
- D4.1 Non Functional Aspects, 04/06/2009
- D3.1 Adaptive Service-Aware Infrastructure Contributions to Model and Architecture, 03/24/2009
- D1.1 Advanced User Service Interactions Contribution, 03/11/2009

Publications:

- Decentralized web service orchestration: a reflective approach, Ricardo Jimenez-Peris, Marta Patiño-Martinez, Ernestina Martel-Jordán, IEEE Int. Symp. on Applied Computing (SAC). Pp. 494-498., Fortaleza, Ceara, Brazil, March 16 - 20, 2008.
- An Autonomic Approach for Replication of Internet-based Services, Damián Serrano, Marta Patiño-Martinez, Ricardo Jimenez-Peris, Bettina Kemme, IEEE Int. Symp. on Reliable Distributed Systems (SRDS). Oct. 2008.Naples (Italy)

OPEN MEDIA PLATFORM

Open Media Platform (OMP) aims to define an open and extensible service & software architecture for media-rich end-user devices, such as mobile phones or mobile media players, that will address software productivity and optimal service delivery challenges. The OMP project brings together Europe's leading researchers, industrial partners and emerging SMEs within the semiconductor industry.
The overall objective of the OMP project is developing technologies to dynamically compose media-rich services that scale with device and network capabilities. This includes the development and definition of :

1. Tool Infrastructure: To develop and disseminate an innovative and extensible computing tool infrastructure consisting of static and dynamic composition tools and methodologies
2. Media Infrastructure: To define, implement a novel media infrastructure consisting of enhanced standard APIs to support resource and context awareness
3. Reference Platforms: To prototype a standards-based run-time environment and methodologies

Deliverables:

- D2.1 OMP Component Based SE Tool Specification, Dec. 2008
- D2.2 OMP Split Compilation Tools (static and dynamic compilers) Specification, Dec. 2008
- D4.4 OMP Prototype Evaluation report, Dec. 2009
- D5.4 ICT Collaboration Report 1, Dec. 2008
- D5.6 ICT Collaboration Report 2, Dec. 2009
- D2.4 OMP Component Based SE Tool, June 2009
- D2.5 OMP Split Compilation Tools (static and dynamic compilers), June 2009
- D2.7 OMP/OpenMAX Component Schema, June 2009
- D3.4 Bellagio OpenSource Implementation, Dec. 2008
- D3.5 Integration to Gstreamer multimedia framework, Dec. 2008
- D3.6 Bellagio OpenSource implementation, Dec. 2009
- D3.7 Integration to Gstreamer multimedia framework, Dec. 2009)

OPEN

OPEN (Open Pervasive Environments for migratory iNteractive services) aims to develop an environment, which provides people with the ability to continue to perform their tasks when they move about and change their interaction device.

The main goal of the OPEN project is to provide a general and open migratory service platform solution based on a sound and innovative scientific approach developed by a multidisciplinary consortium combining the expertise of three technological world leaders, three well known research organisations and one SME.

The service platform will be able to interoperate with existing technologies by:

- adapting and preserving the state of the software application parts dedicated to interacting with end users;
- supporting mechanisms for logic reconfiguration;
- defining suitably flexible mechanisms from the underlying layers.

The objective of OPEN is to provide users with migratory interactive services, which enable users to change interaction platform while continuing their tasks through an interface adapted to the new context in which it will be used.

Deliverables:

- D1.1 Requirements for OPEN Service Platform
- D1.2 Initial OPEN Service Platform architectural framework
- D2.2 Document about architecture for migratory interfaces
- D3.1 Detailed Network Architecture
- D4.1 Solutions for application logic reconfiguration
- D5.1 Initial application requirements and design
- D6.1 Usability criteria for project phases: use cases selection, design, development, test and deployment
- D6.2 Evaluation parameters for enabling the environment programmability
- D6.3 Indicators for technical evaluation
- D7.3 Report on OPEN technology for standardisation

Publications:

Conference Papers

- Anders Nickelsen, Rasmus L. Olsen and Carmen Santoro. Open Pervasive Environments for migratory iNteractive services (OPEN). Proceedings of FIS2008, Vienna, Austria, September 28-30, 2008.

- Fabio Paternò, Carmen Santoro and Antonio Scorcia. Preserving Rich User Interface State in Web Applications across Various Platforms. Proceedings of EIS2008, Pisa, Italy, September 25-26, 2008.
- Fabio Paternò, Carmen Santoro and Antonio Scorcia. User Interface Migration between Mobile Devices and Digital TV. Proceedings of EIS2008, Pisa, Italy, September 25-26, 2008.
- Stühmer, R., Anicic, D., Sen, S., Ma, J., Schmidt, K.-U., and Stojanovic, N. (2009b). Lifting events in rdf from interactions with annotated web pages. In A. Bernstein, D. R. Karger, T. Heath, L. Feigenbaum, D. Maynard, E. Motta, and K. Thirunarayan (Eds.) International Semantic Web Conference , vol. 5823 of Lecture Notes in Computer Science , (pp. 893-908). Springer.
- D. Niebuhr, A. Rausch. Guaranteeing Correctness of Component Bindings in Dynamic Adaptive Systems based on Runtime Testing, Proceedings of the 4th Workshop on Services Integration in Pervasive Environments (SIPE 09) at the International Conference on Pervasive Services 2009 (ICSP 2009)
- D. Niebuhr, A. Rausch. Guaranteeing Correctness of Component Bindings in Dynamic Adaptive Systems, Proceedings of the 35th EUROMICRO Conference on Software Engineering and Advanced Applications (SEAA), Track on Service and Component Based Software Engineering (SCBSE)
- D. Niebuhr, A. Rausch, C. Klein, J. Reichmann, R. Schmid. Achieving Dependable Component Bindings in Dynamic Adaptive Systems – AnRuntime Testing Approach, Proceedings of the 3rd Third IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO 2009)
- A. Rausch, D. Niebuhr, M. Schindler, D. Herrling. Emergency Management System, Proceedings of the International Conference on Pervasive Services 2009 (ICSP 2009)
- G. Ghiani, F. Paternò, C. Santoro, On-demand Cross-Device Interface Components Migration, Proceedings of Mobile HCI 2010, Lisbon
- K. Schmidt, Socially filtered web search: an approach using social bookmarking tags to personalize web search, Proceedings of the 2009 ACM symposium on Applied Computing, Hawaii, USA
- F. Paternò, G. Zichitella, Desktop-to-mobile Web adaptation through customizable two-dimensional semantic redesign, Proceedings of HCSE 2010
- G. Ghiani, F. Paternò, C. Santoro, Partial Web interface migration, Proceedings of the International Conference on Advanced Visual Interfaces, Roma, Italy
- K. Schmidt, Client-Side Event Processing for Personalized Web Advertisement, ODBASE, LNCS 5871, Springer
- M. L. Jakobsen, J. G. Rasmussen, R. L. Olsen, Online estimation of context dynamics and its impact on context sensitive applications, IEEE
- K. Schmidt, Gaining Reactivity for Rich Internet Applications by Introducing Client-side Complex Event Processing and Declarative Rules, 2009 AAAI
- A. Nickelsen, M. Martin, H. Schwefel, Service Migration Protocol for NFC Links, 16th EUNICE/IFIP WG 6.6, LNCS 6164, Norway 2010

Journals Papers (JP)

- Paternò F, Santoro C. and Spano L.D. . MARIA: A universal, declarative, multiple abstraction-level language for service-oriented applications in ubiquitous environments, ACM Transactions on Computer-Human Interaction (TOCHI), Volume 16 , Issue 4 (November 2009)
- Paternò F., Santoro C., Scorcia A. Ambient intelligence for supporting task continuity across multiple devices and implementation languages. In:

Computer Journal, Oxford University Press on behalf of The British Computer Society, 2009.

PERSIST

The objective of PERSIST is to develop Personal Smart Spaces that provide a minimum set of functionalities which can be extended and enhanced as users encounter other smart spaces during their everyday activities. They will be capable of learning and reasoning about users, their intentions, preferences and context. They will be endowed with pro-active behaviours, which enable them to share context information with neighbouring Personal Smart Spaces, resolve conflicts between the preferences of multiple users, make recommendations and act upon them, prioritise, share and balance limited resources between users, services and devices, reason about trustworthiness to protect privacy and be sufficiently fault-tolerant to guarantee their own robustness and dependability.

The vision of PERSIST is of a Personal Smart Space, which is associated with the portable devices carried by the user and which moves around with him/her, providing context-aware pervasiveness to the user at all times and places. The Personal Smart Space will cater for the needs of users, adapting to their preferences and learning new ones as these arise.

Deliverables:

- D1.1 Project Management Handbook
- D1.2 Project report (year 1)
- D1.3 Project report (year 2)
- D1.4 Project final report
- D2.1 Business modelling, cases and opportunities report
- D2.3 Federation and deployment management report
- D2.4 Initial architecture design
- D2.5 Revised architecture design
- D3.1 Detailed design for personal smart spaces
- D3.2 Core architecture dependability and privacy analysis
- D3.3 Prototype of integrated networking and mobility system
- D3.4 Prototype of context and preferences
- D3.5Prototype of service discovery and composition
- D3.6 Dependability and privacy prototype implementation
- D4.1 Prototype grouping and sharing systems
- D4.2 Prototype user intent system
- D4.3 Prototype recommender system
- D4.4 Prototype learning and reasoning systems
- D4.5 Prototype pro-active behaviour system
- D5.1 Testbed environment specification
- D5.2 Test case specification
- D5.3 Final Test report
- D5. Integrated prototype
- D5.5 Evaluation process description
- D5.6 Evaluation software
- D5.7 Evaluation report

- D6.1 Website
- D6.2 Initial exploitation report
- D6.3 Initial standardisation report
- D6.4 Exploitation report
- D6.5 Standardisation report
- D6.6 Dissemination report
- D6.7 Future evolution report
- D6.8 Collaboration Plan
- D6.9 Collaboration Report

Publications:

- Implicit Adaptation of User Preferences in Pervasive Systems, The 4th International Conference on Systems (ICONS 09), Sarah McBurney, Eliza Papadopoulou, Nick Taylor, Howard Williams
- A Model for Personalised Communications Control in Pervasive Systems, The IASTED International Conference on Advances in Computer Science and Engineering (ACSE 09), Eliza Papadopoulou, Sarah McBurney, Howard Williams
- User Preference Management in a Pervasive System should be a Trusted FunctioN, The IASTED International Conference on Advances in Computer Science and Engineering (ACSE 09), Sarah McBurney, Eliza Papadopoulou, Nick Taylor, Howard Williams
- Linking Privacy and User Preferences in the Identity Management for a Pervasive System, IEEE/WIC/ACM International Conference on Web Intelligence (WI-08), Eliza Papadopoulou, Sarah McBurney, Nick Taylor, Howard Williams
- A Dynamic Approach to Dealing with User Preferences, The 2008 International Conference on Intelligent Pervasive Computing (IPC-08), Eliza Papadopoulou, Sarah McBurney, Nick Taylor, Howard Williams
- Adapting Pervasive Environments through Machine Learning and Dynamic Personalisation, The 2008 International Conference on Intelligent Pervasive Computing (IPC-08), Sarah McBurney, Eliza Papadopoulou, Nick Taylor, Howard Williams
- Personal eSpaces and Personal Smart Spaces, 1st PerAda Workshop on Pervasive Adaptation at the 2nd IEEE International Conference on Self-Adaptive and Self-Organising Systems (SASO 08), Nick Taylor
- Novel Pervasive Computing Services experienced through Personal Smart Spaces, International Workshop on Data Management in Ad Hoc and Pervasive Computing (DMAHPC09), C. Venezia, N.K.Taylor, M.H.Williams, K. Doolin, I. Roussaki
- A Pervasive Environment Based on Personal Self-Improving Smart Spaces, Workshop on Architectures and Platforms for AmI at the European Conference on Ambient Intelligence 2008 (AmI 08), Micheal Crotty, Nick Taylor, Howard Williams, Korbinian Frank, Ioanna Roussaki and Mark Roddy
- Knowledge Representation and Inference in Context-Aware Computing Environments, The Second International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies (UBICOMM 2008), M. Rockl, K. Frank, G. Hermann, and M. Vera
- The Bayeslet Concept for modular Context Inference, The Second International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies (UBICOMM 2008), Korbinian Frank, Matthias Rockl, Patrick Robertson

- Technical patterns for long term trusted archiving, The Third International Conference on Digital Society (ICDS 2009), Jan Porekar
- How to make Personal Smart Spaces Context-aware, PERSIST Workshop on Intelligent Pervasive Environments (AISB 2009), Ioanna Roussaki, Nicolas Liampotis, Nikos Kalatzis, Korbinian Frank, Patrick Hayden
- A Hybrid Preference Learning and Context Refinement Architecture, PERSIST Workshop on Intelligent Pervasive Environments (AISB 2009), Korbinian Frank, Patrick Robertson, Sarah McBurney, Nikos Kalatzis, Ioanna Roussaki, Marco Marengo
- Smart Space a new dimension of context, PERSIST Workshop on Intelligent Pervasive Environments (AISB 2009), L. Lamorte, Claudio Venezia
- User Intent to Support Proactivity in a Pervasive System, PERSIST Workshop on Intelligent Pervasive Environments (AISB 2009), Yussuf Abu-Shaaban, Sarah McBurney, Nick Taylor, Howard Williams, Nikos Kalatzis, Ioanna Roussaki
- Giving the User Explicit Control over Implicit Personalisation, PERSIST Workshop on Intelligent Pervasive Environments (AISB 2009), Sarah McBurney, Howard Williams, Nick Taylor, Eliza Papadopoulou
- Protecting the Privacy of Personal Smart Spaces, PERSIST Workshop on Intelligent Pervasive Environments (AISB 2009), Kajetan Dolinar, Eliza Papadopoulou, Nicolas Liampotis, Yussuf Abu-Shaaban, Ioanna Roussaki
- Challenges for Context Management Systems imposed by Context Inference, 6th International Workshop on Managing Ubiquitous Communications and Services (MUCS 2009), Korbinian Frank, Nikos Kalatzis, Ioanna Roussaki, Nicolas Liampotis
- Simultaneous Localization and Mapping for Pedestrians using only Foot-Mounted Inertial Sensors, 11th International Conference on Ubiquitous Computing (Ubicomp 09), Patrick Robertson, Michael Angermann, Bernhard Krach
- Open Vehicular Data Interfaces for In-Car Context Inference, 1st International Conference on Automotive User Interfaces and Interactive Vehicular Applications (AutomotiveUI 2009), Matthias Kranz, Korbinian Frank, Daniel Hermosilla Galcern, Eduard Weber
- Inertial Systems Based Joint Mapping and Positioning for Pedestrian Navigation, ION GNSS 2009, Patrick Robertson, Michael Angermann, Bernhard Krach, Mohammed Khider,
- Design Patterns for a Systemic Privacy Protection, International Journal on Advances in Security, issn 1942-2636. vol. 2, no. 2&3, year 2009, pg. 267 – 287, Kajetan Dolinar, Jan Porekar, Aljo_a Jerman - Bla_i_
- Faster Bayesian Context Inference by Using Dynamic Value Ranges, Proceedings of the Second International Workshop on Information Quality and Quality of Service for Pervasive Computing in Conjunction with IEEE PERCOM 2100, Korbinian Frank, Sergio Fortes Rodriguez, Patrick Robertson, Raquel Barco Moreno
- Comparison of Exact Static and Dynamic Bayesian Context Inference Methods for Activity Recognition, Proceedings of MUCS 2010, 7th International Workshop on
- Managing Ubiquitous Communications and Services, part of PerCom 2010, Korbinian Frank, Matthias Röckl, Maria José Vera Nadales, Patrick Robertson, Tom Pfeifer

PROTEST

The project will deliver methods and tools to support property-based development of systems. Property-driven development is a powerful new mechanism for gaining assurance of system reliability and functionality. However, in order to deliver its full benefits we need tools to integrate property-based testing into the development life cycle.

Project Goals:

- develop software engineering approaches to improve reliability of service-oriented networks
- support fault-finding and diagnosis based on specified properties of the system
- build automated tools that will generate and run tests, monitor execution at run-time, and log events for analysis.

Property-based testing will deliver more effective tests, more efficiently:

- Property discovery
- Test and property evolution
- Property monitoring
- Analysing concurrent systems

Testing with properties as objects improves the competitiveness of software developers, as they can deliver higher quality software for a lower price allows collaborating companies to improve the definition of their software interfaces and so improve the compatibility between their services.

Publications:

- Exago, Property monitoring via log file analysis, Atilla Erd_di, Erlang User Group Meeting, London, 13/5/2010
- Trexplorer, a trace exploration tool for Erlang, Judit Kõszegi, Erlang User Group Meeting, London, 17/2/2010
- QuickChecking Refactoring Tools, Dániel Horpácsi and Dániel Drienyovszky, Erlang User Group Meeting, London, 17/2/2010
- Similar Code Detection and Elimination for Erlang Programs, Huiqing Li and Simon Thompson, Twelfth International Symposium on Practical Aspects of Declarative Languages (PADL'10), 18/1/2010
- Improving your test code with Wrangler, Huiqing Li, Adam Lindberg, Andreas Schumacher and Simon Thompson, Erlang User Conference 2009, 12/11/2009
- Iterative Refinement of Reverse-Engineered Models by Model-Based Testing, N. Walkinshaw, J. Derrick, Q. Guo, 16th International Symposium on Formal Methods (FM 2009), Eindhoven, Netherlands, 2/11/2009

- Applying Testability Transformations to Achieve Structural Coverage of Erlang Programs, Q. Guo, J. Derrick, N. Walkinshaw, 21st IFIP International Conference on Testing of Communicating Systems (TestCom 2009) in Eindhoven, Netherlands, 2/11/2009
- Testing-framework-aware Refactoring, Huiqing Li and Simon Thompson, Third ACM Workshop on Refactoring Tools conference in conjunction with the OOPSLA. Pages 1-4, Orlando, Florida, October 2009. ACM Digital Library, 25/10/2009
- Computing the Structural Difference between State-Based Models, K. Bogdanov, N. Walkinshaw, 16th IEEE Working Conference on Reverse Engineering (WCRE 2009), Lille, France, 13/10/2009
- Implementing an {LTL-to-Büchi} Translator in {Erlang}, Hans Svensson, 8th ACM Sigplan Erlang Workshop, 5/9/2009
- Finding Race Conditions in Erlang with QuickCheck and PULSE, Koen Claessen, Micha_ Pa_ka, Nicholas Smallbone, John Hughes, Hans Svensson, Thomas Arts and Ulf Wiger, ICFP '09: Proceedings of the fifth ACM SIGPLAN international conference on Functional programming, 31/8/2009
- Erlang Programming, Francesco Cesarini, Simon Thompson, O'Reilly Book, 25/6/2009
- Debugging and verification of multi-agent systems, Clara Benac Earle, Eurocast 2009, Canary Islands, Spain, 15/2/2009
- Clone Detection and Removal for Erlang/OTP within a Refactoring Environment, Huiqing Li and Simon Thompson, PEPM 2009 - Partial Evaluation and Program Manipulation Conference in Savannah, Georgia, 20/1/2009
- McErlang: Verifying Distributed Systems Written in Erlang, Lars-Ake Fredlund, Clara Benac Earle and Hans Svensson, COST - European Cooperation in the Field of Scientific and Technical Research, 1/12/2008
- Formal Specification for Free!, John Hughes, 14th International Erlang/OTP User Conference, 13/11/2008
- Update on the EU project ProTest, Thomas Arts, Simon Thompson and Francesco Cesarini, 14th International Erlang/OTP User Conference, 13/11/2008
- Francesco Cesarini, Viviana Pappalardo, Corrado Santoro, Seventh ACM SIGPLAN Erlang Workshop, 1/9/2008
- Inferring Finite-State Models with Temporal Constraints, Neil Walkinshaw, 23rd IEEE/ACM International Conference on Automated Sofware Engineering (ASE'08) L'Aquila, Italy, 1/9/2008
- Results for: Erlang testing tools survey for the ProTest project, Tamas Nagy and Aniko Nagyne Vig, Seventh ACM SIGPLAN Erlang Workshop, 1/9/2008
- Testing Erlang Data Types with Quviq QuickCheck, Thomas Arts, Laura M. Castro and John Hughes, Seventh ACM SIGPLAN Erlang Workshop, 1/9/2008
- Refactoring with Wrangler, updated, Huiqing Li, Simon Thompson, Gyorgy Orosz and Melinda Toth, Seventh ACM SIGPLAN Erlang Workshop, 1/9/2008
- Erlang Testing and Tools Survey, Tamas Nagy and Aniko Nagyne Vig, Seventh ACM SIGPLAN Erlang Workshop,
- 1/9/2008
- Erlang fault detection with model-based testing, Jonas Boberg, Seventh ACM SIGPLAN Erlang Workshop, 1/9/2008
- Verifying Robocup Teams, Clara Benac Earle and Lars-Ake Fredlund, Fifth International Workshop on Model Checking and Artificial Intelligence, 1/7/2008

- ProTest Presentation slides, John Hughes, 13th International Erlang/OTP User Conference, 14/11/2007

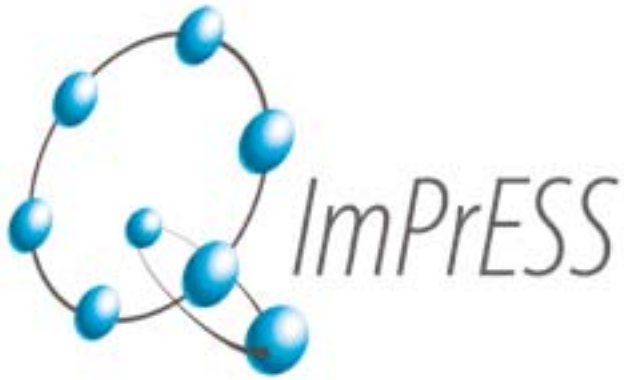

Q-IMPRESS

The Q-ImPrESS project aims at bringing service orientation to critical application domains, such as industrial production control, telecommunication and critical enterprise applications, where guaranteed end-to-end quality of service is particularly important. The main challenge here is to create a method for quality-driven software development and evolution, where the consequences of design decisions and system resource changes on performance, reliability and maintainability can be foreseen through quality impact analysis and simulation.
The Q-ImPrESS project aims at developing a tool-supported method for design-time quality prediction in service-oriented software systems. A prototypic integrated development environment, the Eclipse-based Q-ImPrESS IDE, will assist software engineers during the iterative development and evolution both in existing as well as in newly started software development projects.
The applicability of the proposed method will be validated on two real industrial case studies taken from the domains of production control systems and telecommunications.

Expected benefits:
The methods and tools developed in Q-ImPrESS will allow for cost-effective development and evolution of service-oriented software. Developers will be able to try out different design scenarios and choose the best possible alternative with respect to the impact of these decisions on the quality of the software, before ever writing any single line of code. Design is always a compromise, a balance of tradeoffs between the different quality attributes considered; both attributes of external quality (quality of service) such as performance and reliability, but also attributes of internal quality such as maintainability. Choosing the right design from the beginning and thus avoiding later refactorings, tweaking and project delays to get certain quality attributes such as performance under control can easily mean savings of millions of euros in large software projects.

Deliverables:

- D1.1: Requirements Document
- D2.1: Service Architecture Meta-Model (SAMM)
- D2.2: Service Architecture Extraction Tool
- D3.1: Prediction Model Specification
- D3.2: Prediction Model Generation Tools
- D3.3: Resource Usage Modeling
- D4.1: Quality Prediction Tools
- D5.1: Consistency Checking

- D5.2: Experiments with Impact Analysis
- D5.3: Trade-off Analysis of QoS Attributes
- D6.1: Method and Abstract Workflow
- D6.1: Annex – Guidelines and Tool Manuals
- D7.1: Demonstrator Description
- D7.2: Demonstrator Implementation
- D8.1: Q-ImPrESS Project Website
- D8.2: Organized workshop at ESEC/FSE
- D8.3: Organized Session at Euromicro
- D8.4: Industrial Seminars
- D8.6: Enterprise SOA Showcase
- D8.7: Collaboration Plan
- D8.8: First Collaboration Report
- D8.9: Second Collaboration Report

RESERVOIR

Web 2.0 is rapidly taking hold, offering "the web as a platform". In parallel, traditional client-server computing is starting to lose ground as a new paradigm emerges - the Cloud Computing paradigm. In RESERVOIR, the partners are developing breakthrough system and service technologies that will serve as the infrastructure for Cloud Computing. They aim to achieve this goal by creative coupling of virtualization, grid computing, and business service management techniques.

Essentially, the RESERVOIR project aims to support the emergence of Service-Oriented Computing as a new computing paradigm. In this paradigm, services are software components exposed through network-accessible, platform and language independent interfaces, which enable the composition of complex distributed applications out of loosely coupled components.

Service-Oriented Computing (SOC) carries the visionary promise of reducing software complexity and costs, expediting time-to-market, improving reliability and enhancing accessibility of consumers to government and business services. However, conditional to the wide-scale penetration of SOC to the economic landscape, the ICT industry needs to solve several well-recognized technical challenges. One such key challenge is the development of a scalable and effective service-oriented infrastructure. This is the challenge addressed by RESERVOIR.

Deliverables:

- Market analysis and exploitation plan, August 2010
- Project Progress Report, August 2010
- Training Material, February 2010
- Standardization, February 2010
- D2.1.1 Virtual Machine Technologies
- D2.2.1 Virtual Java Service container
- D2.3.1 Relocation Enablement (M12
- Scientific Report, January 2010
- High Level Architectural Specification, Dec 2009
- Virtual Execution Environment (VEEH) Design and Open Specification, Nov 2009
- Project Progress Report, September 2009
- Integration Plan Final, February 2009
- Telco Scenario: Compute Cloud Capabilities for Telef´onica, February 2009
- Virtual Execution Environment Manager (VEEM) Scientific Report, Feb 2009
- Jan 29, 2009 that includes the deliverables D4.1.1, D4.1.2 and
- D4.1.3)
- Scientific Report, January 2009 that includes the scientific deliverables of Work Packages WP3.1: VEE Provisioning and Supervision, WP3.2: Allocation Policy Management, and WP3.3: Federation of Management Domains.
- Software Prototype Report, January 2009

- Service Manager Scientific Report, Nov 2008
- Service Manager Design and Open Specification, Nov 2008
- RESERVOIR- High Level Architecture, June 2008

Publications:

- Principles, Methodology and Tools for Engineering Cloud Computing Systems, Luis M. Vaquero, Luis Rodero-Merino, Juan Cceres, Clovis Chapman, Maik Lindner and Fermn Galn, chapter in book Advanced Design Approaches to Emerging Software Systems: Principles, Methodology and Tools, IGI Global, 2010
- The Cloud Supply Chain: A Framework for Information, Monitoring, Accounting and Billing, M. Lindner, F.G. Marquez, C. Chapman, S. Clayman, D. Henriksson, and E. Elmroth, 2nd International ICST Conference on Cloud Computing (CloudComp 2010), in press, Springer Verlag, 2010.
- RESERVOIR - A European Cloud Computing Project, Naqvi, Philippe Massonet, ERCIM News 83, October 2010.
- From Infrastructure Delivery to Service Management in Clouds, L. Rodero-Merino, L. M. Vaquero, V. Gil, F. Galn, J. Fontn, R. S. Montero, I. M. Llorente - Future Generation Computer Systems journal
- An Architecture For Federated Cloud Computing, B. Rochwerger, D. Breitgand, D. Hadas, J. Caceres, A. Galis, E. Levy, P. Massonet, M. Villari, E. Elmroth, T. Vazquez, R. Montero, J. Tordsson; Chapter 15 of ”Cloud Computing: Principles and Paradigms” book - Rajkumar Buyya (Ed.), James Broberg (Ed.) and Andrzej Goscinski (Ed.), Wiley 2010
- Improving Virtual Machine Migration In Federated Cloud Environments, A. Celesti, F. Tusa, M. Villari, A. Puliafito - Published on Proceedings of The Second International Conference on Evolving Internet (INTERNET 2010); 20-25 September 2010
- A Naming System Applied To A Reservoir Cloud, A. Celesti, M. Vil- lari, A. Puliafito -Published on The 6th IEEE International Conference on Information Assurance and Security (IAS 2010) -The 1st International Workshop on Security and Performance in Emerging Distributed Architectures (SPEDA2010) August 2010
- How To Enhance Cloud Architectures To Enable Cross-federation, A. Celesti, F. Tusa, M. Villari, A. Puliafito - Published on Proceedings of The 3rd IEEE International Conference on Cloud Computing (IEEE Cloud 2010), Miami, Florida, USA July 2010.
- An Aspect-Oriented Approach to Consistency-Preserving Caching and Compression of Web Service Response Messages, W. Li, J. Tordsson, and E. Elmroth, The 2010 IEEE International Conference on Web Services (ICWS 2010), DOI 10.1109/ICWS.2010.83, pp. 526-533, 2010.
- Three-phase Cross-cloud Federation Model: The Cloud Sso Authentication, A. Celesti, F. Tusa, M. Villari, A. Puliafito - Published on Proceed- ings of The 2nd IEEE International Conference on Advances in Future Internet (AFIN 2010), Venice, Italy July 2010
- Ecosystem Of Cloud Naming Systems: An Approach For The Management And Integration Of Independent Cloud Name Spaces, A. Celesti, M. Villari, A. Puliafito - Published on Proceedings of The 9th IEEE In- ternational Symposium on Network Computing and Applications (IEEE NCA10), Boston, USA July 2010

- Engineering Autonomic Controllers For Virtualized Web Applications, G. Toffetti, A. Gambi, M. Pezzé, C. Pautasso, International Conference of Web Engineering (ICWE), Vienna, July 2010
- Software Architecture Definition for On-demand Cloud Provisioning, C. Chapman, W. Emmerich, F. Galn, S. Clayman, A. Galis - ACM International Symposium on High Performance Distributed Computing (HPDC 2010), 21-25 June 2010, Chicago hpdc2010.eecs.northwestern.edu/
- Security And Cloud Computing: Intercloud Identity Management Infrastructure, A. Celesti, F. Tusa, M. Villari, A. Puliafito - Published on Proceedings of The 19th IEEE International Workshops on Enabling Technologies: Infrastructures for Collaborative Enterprises (WETICE 2010) - ETNGRID, Tei of Larissa, Greece June 2010
- Scalable Cloud Computing Defenses for Detection and Analysis of DDOS attacks, M. Villari, B. Rochwerger, P. Massonet, J. Latanicki - Networking Conference on Trans-European Research and Education Networking Association (TERENA 2010) Vilnius, Lithuania, from 31 May to 3 June 2010
- Cloud economics: Dynamic Business models for the Cloud Market, Maik Lindner, Luis M. Vaquero, Luis Rodero-Merino, Juan Cáceres, International Journal of Business Information Systems (IJBIS), Issue 3/4, Volume 5, 2010.
- Protecting SLAs with surrogate models, A. Gambi, G. Toffetti, M. Pezz, International Workshop on Principles of Engineering Service-Oriented Systems (PESOS) @ ICSE, Cape Town, May 2010
- Design Of A Cloud Naming Framework, A. Celesti, M. Villari, A. Puliafito - Published on Proceedings of The ACM International Conference on Computing Frontiers (ACM CF10), Bertinoro, Italy, ACM pp.105-106, May 2010; ISBN:978-1-4503-0044-5.
- Monitoring Future Internet Service Clouds, S. Clayman, A. Galis, C. Chapman, L.R, Merino, .L M. Vaquero, K. Nagin, B. Rochwerger, G. Toffetti - ”Towards the Future Internet - A European Research Perspective book, April 2010 IOS Press, http://www.iospress.nl/
- Monitoring Virtual Networks, S. Clayman, A. Galis, L. Mamatas; 12th IEEE/IFIP Network Operations and Management Symposium (NOMS 2010) - International on Management of the Future Internet (ManFI 2010) 19-23 April 2010, Osaka http://www.manfi.org/2010/
- Elastic Service Management in Computational Clouds, C. Chapman, W. Emmerich, F. Galn, S. Clayman, A. Galis 12th IEEE/IFIP Network Operations and Management Symposium (NOMS2010) / International Workshop on Cloud Management (CloudMan 2010) 19-23 April 2010, Os- aka http://cloudman2010.lncc.br/
- Scalable Cloud Defenses for Detection, Analysis and Mitigation of DDoS Attacks, J. Latanicki, P. Massonet, S. Naqvi, B. Rochwerger and M. Villari - Future Internet Assembly (FIA) - 14-16 April 2010 Valencia Conference; FIA book chapter IOSPress http://www.iospress.nl
- Distributed Usage Logging for Federated Grids, E. Elmroth and D. Henriksson, Future Generation Computer Systems., Vol 26, No. 8, pp. 1215 - 1225, 2010
- Facing Services in Computational Clouds, T. Metsch, L. Vaquero, L. Rodero-Merino, M. Lindner, P. Massonet; in the book: ”High Speed and Large Scale Computing” IOSPress ISBN: 978-1-60750-073-5; January 2010
- Reservoir - Cloud Computing: by the Reservoir project, S. Beco, A. Maraschini, F. Pacini, O. Biran, D. Breitgand, K. Meth, B. Rochwerger, E. Salant, E. Silvera, S. Tal, Y. Wolfsthal, M. B. Yehuda, J. Caceres, J. Hierro, W. Emmerich, A. Galis, L. Edblom, E. Elmroth, D. Henriksson, F. Hernandez, J. Tordsson, A. Hohl, E. Levy, A. Sampaio, B. Scheuermann, M. Wusthoff, J.

Latanicki, G. Lopez, J. Marin-Frisonroche, A. Dorr, F. Ferstl, E. Huedo, I. Llorente, R. Montero, P. Massonet, S. Naqvi, G. Dal- Ions, M. Pezz, A. Puliafito, C. Ragusa, M. Scarpa and S. Muscella; Nuovo Cimento Journal C, Vol. 32, No. 2, pp. 99 - 103, 2009.

- Virtual Infrastructure Management in Private and Hybrid Clouds, B. Sotomayor, R. S. Montero, I. M. Llorente and I. Foster, (IEEE Internet Computing, in press, ISSN: 1089-7801)
- Management and Service-aware Networking Architectures for Future Internet Position Paper: System Functions, Capabilities and Requirements, A. Galis, H. Abramowicz, M. Brunner, D. Raz, P. Chemouil, J. Butler, C. Polychronopoulos, S. Clayman, H. de Meer, T. Coupaye, A. Pras, K. Sabnani, P. Massonet, S. Naqvi - Invited paper IEEE 2009 Fourth International Conference on Networking in China (ChinaCom09) 26-28 August 2009, Xi'an, China
- Interfaces for Placement, Migration, and Monitoring of Virtual Machines in Federated Clouds, E. Elmroth and L. Larsson. In Juan. E. Guerrero (ed), Proceedings of the Eighth International Conference on Grid and Cooperative Computing (GCC 2009), IEEE Computer Society Press, pp. 253 - 260, 2009
- Accounting and Billing Architecture for Federated Cloud Infrastructures, E. Elmroth, F. Galn, D. Henriksson and D. Perales. In Juan. E. Guerrero (ed), Proceedings of the Eighth International Conference on Grid and Cooperative Computing (GCC 2009), IEEE Computer Society Press, pp. 268 - 275, 2009.
- Mitigating Security Threats To Largescale Cross Border Virtualization Infrastructures, M. Villari, F. Tusa, P. Massonet, S. Naqvi, J. Latanicki - International Conference on Cloud Computing, Munich, Germany, 19-21 October 2009
- From Grids To Clouds Shift In Security Services Architecture, J. Latan- icki, M. Villari, P. Massonet, S. Naqvi CGW'09 - Cracow Grid Workshop, Krakow, Poland 12-14 October 2009
- New Cloud Computing Challenges: The Security Impact In The Social World, M. Villari - ZISC Workshop on Security in Virtualized Environments and Cloud Computing, Invited Speaker ETH - Zurich, Switzerland, 10-12 September 2009.
- Cloud Computing for On-Demand Grid Resource Provisioning, I. M. Llorente, R. Moreno-Vozmediano, and R. S. Montero, (Advances in Parallel Computing, Volume 18 (2009): "High Speed and Large Scale Scientific Computing", pp. 177 - 191. IOS Press, 2009, ISBN 978-1-60750-073-5)
- RESERVOIR: Management Technologies and Requirements for Next Generation Service Oriented Infrastructures, B. Rochwerger, A. Galis, E. Levy, J. A. Cáceres, D. Breitgand, Y. Wolfsthal, I. M. Llorente, M. Wusthoff, R. S. Montero, E. Elmroth, (accepted as a short paper at The 11th IFIP/IEEE International Symposium on Integrated Management 1-5 June 2009, New York, USA)
- Dynamic Provision of Computing Resources from Grid Infrastructures and Cloud Providers,. Vazquez, E. Huedo, R. S. Montero, y I. M. Llorente, Workshop on Grids, Clouds and Virtualization (WGCV), held in conjunction with the 4th International Conference on Grid and Pervasive Computing, GPC 2009 Geneva, Proceedings, IEEE Computer Society, in press
- Resource Leasing and the Art of Suspending Virtual Machines, B. Sotomayor, R. S. Montero, I. M. Llorente and I. Foster, IEEE International Conference on High Performance Computing and Communications (HPCC-09), Seoul, Korea, Proceedings, IEEE Computer Society, in press
- Elastic Management of Cluster-based Services in the Cloud, R. Moreno, R. S. Montero, e I. M. Llorente, First Workshop on Automated Control for

Datacenters and Clouds (ACDC09), held in conjunction with the 6th International Conference on Autonomic Computing and Communications, Barcelona, Proceedings, ACM, in press

- Cloud Computing for on-Demand Grid Resource Provisioning, I.M. Llorente, R. Moreno-Vozmediano, and R.S. Montero; In Advances in Parallel Computing, IOS Press (paper in press), 09
- Service Specification in Cloud Environments Based on Extensions to Open Standards, Fermín Galán, Americo Sampaio, Luis Rodero-Merino, Irit Loy, Victor Gil, Luis M. Vaquero, Mark Wusthoff, Fourth International Conference on COMmunication System softWAre and middlewaRE (COMSWARE 2009), June 2009, Dublin (Ireland)
- NAP: a Building Block for Remediating Performance Bottlenecks via Black Box Network Analysis, Muli Ben-Yehuda, D. Breitgand, Michael Factor, Hillel Kiolodner, Valentin Kravtsov, Dan Pelleg; 6th IEEE Interna- tional Conference on Autonomic Computing, 15-19 June, 2009, Barcelona, Spain
- An Open Source Solution for Virtual Infrastructure Management in Private Clouds, B. Sotomayor, R. S. Montero, I. M. Llorente, and I. Foster. - Internet Computing, 2009
- Over the Clouds: Holistic Service Management and Custom Automated Scalability, Luis Rodero-Merino, et.al.; Special Issue on Cloud Comput- ing of journal 'IEEE Internet Computing', 2009
- Design for Future Internet Service Infrastructures, B. Rochwerger, A. Galis, D. Breitgand, E. Levy, J. A. Caceres, I. M. Llorente, Y. Wolfsthal, M. Wusthoff, S. Clayman, W. Emmerich, E. Elmroth, R. S. Montero- part of "Towards Future Internet" book published by IOS Press, ISBN 978-1- 60750-007-0, pp350, April 2009, IOS Press
- The RESERVOIR Model And Architecture for Open Federated Cloud Computing, B. Rochwerger, J. Caceres, R.S. Montero, D. Breitgand, E. Elmroth, A. Galis, E. Levy, I.M. Llorente, K. Nagin, Y. Wolfstha, IBM Systems Journal, in press
- Capacity Leasing in Cloud System using the OpenNebula Engine, B. Sotomayor, R. S. Montero, I. M. Llorente and I. Foster, Workshop on Cloud Computing and its Applications 2008 (CCA08), Chicago, Illinois, USA, October, 2008, Cornell University Library, in press.
- A Break in the Clouds: Towards a Cloud Definition, L.M. Vaquero, L. Rodero-Merino, J. Cáceres, M. Lindner. ACM Computer Communication Reviews. 2009;39(1):50-55. 2008
- Dynamic Deployment of Custom Execution Environments in Grids, R. S. Montero, E. Huedo and I. M. Llorente; was accepted for publication in the The Second International Conference on Advanced Engineering Computing and Applications in Sciences (ADVCOMP 2008), Valencia Spain, October 2008.
- Capacity Leasing in Cloud Systems using the OpenNebula Engine. B. Sotomayor, R. S. Montero, I. M. Llorente and I. Foster Workshop on Cloud Computing and its Applications 2008 (CCA08), Chicago, Illinois, USA, October, 2008, Cornell University Library, in press
- Design for Future Internet Service Infrastructures, B. Rochwerger, A. Galis, D. Breitgand, E. Levy, J. A. Cáceres, I. M. Llorente, Y. Wolfsthal, M. Wusthoff, S. Clayman, W. Emmerich, E. Elmroth , R. S. Montero- submitted on 30th Oct 2008 to Future Internet Assembly (FIA) - March 09 Pregue Conference & FIA book (i.e. to be published by IOSPress http://www.iospress.nl by March 09) http://www.future-internet.eu/publications/fi-conference-publication/call-for-papers.html
- Dynamic Provisioning of Virtual Clusters for Grid Computing, M. Rodriguez, D. Tapiador, J. Fontan, E. Huedo, R. S. Montero and I. M. Llorente; was

accepted for publication in the 3rd Workshop on Virtualization in High-Performance Cluster and Grid Computing (VHPC 08) as part of Euro-Par 2008, Canary Island, Spain, August 2008.
- Automatic Configuration of an Autonomic Controller: an Experimental Study with Zero-Configuration Policies, T. Heinis, C. Pautasso. Accepted to the 5th International Conference on Autonomic Computing (ICAC 2008), Chicago, USA, June 2008.
- Reservoir Whitepaper for OGF23, Barcelona, June 2008
- RESTful Web Services vs. Big Web Services: Making the Right Architectural Decision, C. Pautasso, O. Zimmermann, F. Leymann. Accepted to 17th International World Wide Web Conference (WWW2008), Bejing, China, April 2008.
- RESERVOIR - An ICT Infrastructure for Reliable and Effective Delivery of Services as Utilities, IBM Technical Library H-0262; Feb 2008

ROMULUS

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

The main concept of ROMULUS is researching on novel methods for increasing productivity and reliability of web software development, in particularly, focused on Java web development.

ROMULUS proposal is based on recognising some of the deficiencies of standard Java Enterprise Edition, and proposing a new paradigm for developing web applications taking advantage of new trends in software engineering, such as domain driven design combined with agile development methodologies, and some of the principles from Ruby on Rails.

In order to have a serious impact, the project does not start from scratch, it is based on two mature open source projects, Roma and LIFERAY, which will be extended according to this proposal needs and following an open source project development methodology, in order to disseminate and exploit the results of the project. The idea of using these projects is to solve the former problems, with Roma will get to make Java application development easy, integrate frameworks and tools using a metaframework and decreasing the implementation time. While with LIFERAY, we are integrating the leading enterprise open source portal framework, with relevant industrial references.

Deliverables:

- Romulus Requirements
- Evaluation Plan
- First Evaluation and FP7 Liasons report
- Final Evaluation and FP7 Liasons report
- Open Source Community set up
- Roma MetaFramework Early Report
- Roma MetaFramework Final Report
- Romulus RAD early
- Romulus RAD final
- Semantic Powered Development Tools Report
- Annual Report on Mashup Integration
- Final Report on mashup integration
- Report on Emerging Portal Technologies
- Report on Emerging Portal Technologies integration
- Report of Integration of Testing and Security

- Annual Dissemination Report, Explotation Plan and liasons with other projects
- Final Dissemination Report, Explotation Plan and liaisons with other projects
- Report on Rationale and Requirements of the demonstrator
- Report on Early Demonstrator
- Report on the Final Demonstrator

Publications:

- Michael Hausenblas and Marcel Karnstedt, Understanding Linked Open Data as a Web-Scale Database. In Second International Conference on Advances in Databases, Knowledge, and Data Applications 2010.
- O. Ureche, A. Iqbal, M. Hausenblas Accessing Site-Specific APIs Through Write-Wrappers From The Web of Data. In Semantics for the Rest of Us - Variants of Semantic Web Languages in the Real World Workshop at 8th International Semantic Web Conference (ISWC 2009), Washington DC, USA, 2009.
- A. Iqbal, O. Ureche, M. Hausenblas. Integrating Linked Data Driven Software Development Interaction into an IDE. In 5th International Workshop on Semantic Web Enabled Software Engineering Workshop at 8th International Semantic Web Conference (ISWC 2009), Washington DC, USA, 2009.
- José Fernández-Villamor, Mercedes Garijo Ayestarán, Carlos Ángel Iglesias. Descripción semántica de aplicaciones web mediante microservicios (Semantic description of Web applications through microservices). Proceedings of the Symposium on Telematics Engineering, 8th edition, Universidad Politécnica de Cartagena, Spain. September 15-17, 2009
- José Fernández-Villamor, Carlos Ángel Iglesias, Mercedes Garijo Ayestarán, Microservices: lightweight service descriptions for REST architectural style. In Proceedings of the Second International Conference on Agents and Artificial Intelligence, 2010.
- Boni García, Juan C. Dueñas, Hugo A. Parada. Functional Testing based on Web Navigation with Contracts. IADIS International Conference (WWW/INTERNET09). Rome, Italy. November 2009.
- Boni García, Juan C. Dueñas, Jose I. Fernández-Villamor, Adam Westerski, Mercedes Garijo, Carlos A. Iglesias. ROMULUS: Domain Driven Design and Mashup Oriented Development based on Open Source Java Metaframework for Pragmatic, Reliable and Secure Web Development. 14th European Conference on Software Maintenance and Reengineering (CMSR10). Madrid, Spain. March 2010.
- Pablo Martín, Guillermo Hernández, Carlos Ángel Iglesias, Luca Garulli, Giordano Maestro. Janus: un Generador de la Vista de Roma Framework basado en Plantillas (Janus, a new View Generator of Roma Framework based on templates). Proceedings of the Symposium on Telematics Engineering, 8th edition, Universidad Politécnica de Cartagena, Spain. September 15-17, 2009
- Boni Garcia, Juan C. Dueñas, Hugo A. Parada. Automatic Functional and Structural Test Case Generation for Web Applications based on Agile Frameworks. 5th Int'l Workshop on Automated Specification and Verification of Web Systems (WWV09). Castle of Hagenberg, Austria. 17th July 2009
- José Fernández-Villamor, Laura Díaz-Casillas, Carlos Ángel Iglesias. A comparison model for Agile Web Frameworks. Proceedings of the 3rd Euro American Conference on Telematics and Information Systems, 2008

- Boni García, Juan C. Dueñas, Hugo A. Parada. A Template-Based Method for Automatic Unit Test Case Generation of Web Applications. 6th Workshop on System Testing and Validation (STV08). Madrid, Spain. 13rd December 2008
- A. Iqbal, O. Ureche, M. Hausenblas, G. Tummarello. LD2SD: Linked Data Driven Software Development. 21st International Conference on Software Engineering and Knowledge Engineering (SEKE 09). Jul 2009
- K. Alexander, R. Cyganiak, M. Hausenblas, J. Zhao. Describing Linked Datasets - On the Design and Usage of voiD, the 'Vocabulary of Interlinked Datasets'. Linked Data on the Web Workshop (LDOW 09) at WWW09. Apr 2009
- M. Hausenblas, W. Halb, Y. Raimond, L. Feigenbaum, D. Ayers. SCOVO: Using Statistics on the Web of Data. ESWC2009 Semantic Web In-Use Track. Feb 2009
- A Westerski, A. Iqbal, G. Tummarello, S. Decker Sindice Widgets: Lightweight embedding of Semantic Web capabilities into existing user applications. In Proceedings of the 4th International Workshop on Semantic Web Enabled Software Engineering 2008

S-Cube

S-Cube, the European Network of Excellence in Software Services and Systems, will establish an integrated, multidisciplinary, vibrant research community. This will enable Europe to lead the software-services revolution, thereby helping shape the software-service based Internet which is the backbone of our future interactive society.

S-Cube's Major Research Objectives are:
- Europe as a leader in the software services revolution
- Incubator for the next wave of service technologies
- Agile & holistic service engineering & adaptation principles, techniques & methods to foster innovation

S-Cube's Major Integration Objectives are to establish a unified, multidisciplinary, vibrant research community, to inaugurate a Europe-wide education and training programme for researchers and industry and to establish a trust relationship with industry.

Deliverables:

- CD JRA 1.1.5 Analysis on how to exploit codified HCI and codified context knowledge for SBA engineering, May 17, 2010
- CD-IA-1.1.1 Comprehensive Overview of the State of the Art on Service-based Systems, May 19, 2009
- CD-IA-1.1.2 Separate Knowledge Models for Functional Layers, May 19, 2009
- CD-IA-1.1.3 Integrated Knowledge Model, May 17, 2010
- CD-IA-2.1.2 Identification of Scientific Subjects and Partners for Mobility, Jun 16, 2009
- CD-IA-2.1.3 Initial assessment of results of a separate mobility program for researchers and students, May 17, 2010
- CD-IA-2.2.2 Collection of Industrial Best Practices, Scenarios and Business Cases, Jun 16, 2009
- CD-IA-2.2.4 Report on common Pilot Cases, May 17, 2010
- CD-IA-3.1.1 Integration Framework Baseline, May 19, 2009
- CD-IA-3.1.3 First Version of Integration Framework, May 17, 2010
- CD-IA-3.2.2 Results of the First Validation, May 17, 2010
- CD-JRA-1.1.2 Separate Design Knowledge Models for Software Engineering and Service Based Computing, Jun 16, 2009
- CD-JRA-1.1.4 Coordinated design knowledge models for software engineering and service-based computing, May 17, 2010
- CD-JRA-1.2.2 Taxonomy of Adaptation Principles and Mechanisms, Jun 16, 2009
- CD-JRA-1.2.4 Integrated adaptation and monitoring principles, techniques and methodologies across functional SBA layers, May 17, 2010
- CD-JRA-1.3.2 Quality Reference Model for SBA, May 19, 2009

- CD-JRA-1.3.3 Initial Concepts for Specifying End-to-End Quality Characteristics and Negotiating SLAs, May 17, 2010
- CD-JRA-1.3.4 Initial Set of Principles, Techniques and Methodologies for Assuring End-toend Quality and Monitoring of SLAs, May 17, 2010
- CD-JRA-2.1.2 Initial Models and Mechanisms for Quantitative Analysis of Correlations Between KPIs, SLAs and Underlying Business Processes, May 19, 2009
- CD-JRA-2.1.3 Business Transaction Language, May 17, 2010
- CD-JRA-2.2.2 Models and Mechanisms for Coordinated Service Compositions, May 19, 2009
- CD-JRA-2.2.3 Algorithms and Techniques for Splitting and Merging Service Compositions, May 17, 2010
- CD-JRA-2.3.3 Requirements for Service Registries in Dynamic Environments and Evaluation of Existing Service Registries, May 19, 2009
- CD-JRA-2.3.2 Basic Requirements for Self-Healing Services and Decision Support for Local Adaptation, May 19, 2009
- CD-JRA-2.3.4 Decision Support for Local Adaptation, May 17, 2010
- CD-SoE-1.1.2 Organizational Structure for Virtual Campus, May 18, 2009
- CD-SoE-1.1.3 Initial Learning Programme and Module Collection, May 17, 2010
- CD-SoE-1.2.2 Collaboration Plan for Joint Activities with ICT SSAI&E Projects, Nov 21, 2008
- CD-SoE-1.2.3 Definition of Strategy for Community Outreach, May 17, 2010
- CD-SoE-1.2.4a Report on Dissemination of Network Results and Collaboration with ICT SSAI&E Projects, 14, 2010
- CD-SoE-1.2.4b Report on dissemination of network results and collaboration with ICT SSAI&E Projects, May 17, 2010
- CD-SoE-1.2.5a Plan for Join Publications, May 17, 2010
- CD-SoE-1.2.6 Creation of Course Registries, May 17, 2010
- PO-IA-2.1.1 Definition and Set-up of Mobility Program Policies, Nov 21, 2008
- PO-IA-2.2.1 Identification of Potential Industrial Collaborators, May 19, 2009
- PO-IA-3.2.1 Initial Definition of Validation Scenarios, May 17, 2010
- PO-IA-3.2.6 Results of the Initial Empirical Evaluation, May 17, 2010
- PO-JRA-1.1.1 State of the art report on software engineering design knowledge and Survey of HCI and contextual knowledge, May 19, 2009
- PO-JRA-1.1.3 Codified Human-Computer Interaction (HCI) Knowledge and Context Factors, May 19, 2009
- PO-JRA-1.2.1 State of the Art Report, Gap Analysis of Knowledge on Principles, Techniques and Methodologies for Monitoring and Adaptation of SBAs, Aug 06, 2009
- PO-JRA-1.2.3-Baseline of Adaptation and Monitoring Principles, Techniques, and Methodologies across Functional SBA Layers, Jun 08, 2010
- PO-JRA-1.3.1 Survey of quality related aspects relevant for SBAs, May 19, 2009
- PO-JRA-2.1.1 Survey on business process modelling and Management, Nov 24, 2010
- PO-JRA-2.2.1 Overview of the state of the art in composition and coordination of services, May 19, 2009
- PO-JRA-2.3.1 Use Case Description and State-of-the-Art, May 19, 2009
- PO-SoE-1.1.1 Analysis of Current Programmes, Nov 21, 2008
- PO-SoE-1.2.9 Updated Project Web Portal, May 18, 2009
- PO-SoE-1.2.5b Plan for Joint Publications, May 17, 2010
- PO-SoE-1.2.7 Thematic Workshops Targeting Industry and Specialised Communities, May 17, 2010

- PO-SoE-1.2.9b Updated Project Web Portal, May 17, 2010
- PO-SoE-1.2.9c Updated Project Web Portal, May 17, 2010

Publications:

- Andrikopoulos, V., & Parkin, M. (2009). Integrated Knowledge Model. S-Cube Consortium.
- Andrikopoulos, V., van den Heuvel, W. - J., & Papazoglou, M. P. (2008). Comprehensive overview of the state of the art on service-based systems. S-Cube Consortium.
- Di Nitto, E., Karastoyanova, D., Metzger, A., Parkin, M., Pistore, M., Pohl, K., et al. (2009). S-Cube: Addressing Multidisciplinary Research Challenges for the Internet of Services. In Towards the Future Internet: A European Research Perspective. Amsterdam: IOS Press.
- Gehlert, A. (2008). Identification of Potential Industrial Collaborators. S-Cube Consortium.
- Gehlert, A. (2010). Analysis on how to exploit codified HCI and codified context knowledge for SBA engineering. S-Cube Consortium.
- Gehlert, A., & Metzger, A. (2008). Quality Reference Model for SBA. S-Cube Consortium.
- Gehlert, A., Pistore, M., Plebani, P., & Versienti, L. (2009). First Version of Integration Framework. S-Cube Consortium.
- Hacid, M. - S. (2008). Definition and Set-up of Mobility Program Policies.
- Hashmi, S. I., Lane, S., Karastoyanova, D., & Richardson, I. (2010). A CMMI Based Configuration Management Framework to Manage the Quality of Service Based Applications. In EuroSPI: Industrial Proceedings of the 17th European conference on Software Process Improvement.
- Kazhamiakin, R. (2009). Initial Definition of Validation Scenarios. S-Cube Consortium.
- Kazhamiakin, R. (2010). Results of the First Validation. S-Cube Consortium.
- Kounkou, A. (2010). Creation of course registries. S-Cube Consortium.
- Kounkou, A. (2010). Initial learning programme and module collection. S-Cube Consortium.
- Kounkou, A. (2010). Results of the initial Empirical Evaluation. S-Cube Consortium.
- Kounkou, A., Cullinane, A., & Maiden, N. (2008). Using HCI Knowledge in Service-Centric Applications: the S-Cube Network of Excellence.
- Laforenza, D. (2008). Definition of the Collaboration Infrastructure and Identification of Appropriate Tools for Exchange of Knowledge. S-Cube Consortium.
- Leitner, P. (2009). Requirements for Service Infrastructures in Dynamic Environments and Evaluation of Existing Service Registries. S-Cube Consortium.
- Maiden, N., & Kounkou, A. (2010). Plan for Joint Publications. S-Cube Consortium.
- Mancioppi, M. (2009). Detailed requirements analysis for a Distributed Services Laboratory Network. S-Cube Consortium.
- Mancioppi, M. (2010). Report on dissemination of network results and collaboration with ICT SSAI&E Projects. S-Cube Consortium.
- Mazza, V., & Pernici, B. (2009). Report on common Pilot Cases. S-Cube Consortium.
- Metzger, A. (2008). Updated Project Web Portal. S-Cube Consortium.
- Metzger, A. (2009). Updated Project Web Portal. S-Cube Consortium.

- Metzger, A., & Pohl, K. (2009). S-Cube: Enabling the Next Generation of Software Services. In J. Filipe, J. Cordeiro, & J. Cardoso (Eds.), Proceedings of the 5th Intl. Conference on Web Information Systems and Technologies (WEBIST 2008), ENASE 2008 Best Papers. (pp. 40–47). Lecture Notes in Business Information Processing. Berlin Heidelberg: Springer-Verlag.
- Metzger, A., & Pohl, K. (2009). Towards the Next Generation of Service-Based Systems: The S-Cube Research Framework. In J. P. van Eck, J. Gordijin, & R. Wieringa (Eds.), CAiSE 2009 (pp. 11–16). LNCS. Berlin Heidelberg: Springer-Verlag.
- Metzger, A., Schmieders, E., Cappiello, C., Di Nitto, E., Kazhamiakin, R., Pernici, B., et al. (2010). Towards Proactive Adaptation: A Journey along the S-Cube Service Life-Cycle. In MESOA: 4th International Workshop on Maintenance and Evolution of Service-Oriented Systems.
- Nikolaou, C., & Bitsaki, M. (2009). Definition of Strategy for Community Outreach. S-Cube Consortium.
- Nikolaou, C., & Bitsaki, M. (2009). Plan for Join Publications. S-Cube Consortium.
- Nitto, E. D. (2008). State of the art report on software engineering design knowledge and Survey of HCI and contextual Knowledge. S-Cube Consortium.
- Di Nitto, E., Karastoyanova, D., Metzger, A., Parkin, M., Pistore, M., Pohl, K., et al. (2009). S-Cube: Addressing Multidisciplinary Research Challenges for the Internet of Services. In Towards the Future Internet: A European Research Perspective. Amsterdam: IOS Press.
- Kounkou, A., Cullinane, A., & Maiden, N. (2008). Using HCI Knowledge in Service-Centric Applications: the S-Cube Network of Excellence.
- Metzger, A., & Pohl, K. (2009). S-Cube: Enabling the Next Generation of Software Services. In J. Filipe, J. Cordeiro, & J. Cardoso (Eds.), Proceedings of the 5th Intl. Conference on Web Information Systems and Technologies (WEBIST 2008), ENASE 2008 Best Papers. (pp. 40–47). Lecture Notes in Business Information Processing. Berlin Heidelberg: Springer-Verlag.
- Metzger, A., & Pohl, K. (2009). Towards the Next Generation of Service-Based Systems: The S-Cube Research Framework. In J. P. van Eck, J. Gordijin, & R. Wieringa (Eds.), CAiSE 2009 (pp. 11–16). LNCS. Berlin Heidelberg: Springer-Verlag.
- Metzger, A., Schmieders, E., Cappiello, C., Di Nitto, E., Kazhamiakin, R., Pernici, B., et al. (2010). Towards Proactive Adaptation: A Journey along the S-Cube Service Life-Cycle. In MESOA: 4th International Workshop on Maintenance and Evolution of Service-Oriented Systems.
- Papazoglou, M., & Pohl, K. (2009). S-CUBE: The Network of Excellence on Software Services and Systems. In E. Di Nitto, A. Sassen, P. Traverso, & A. Zwegers (Eds.), At Your Service: An Overview of Results of Projects in the Field of Service Engineering of the IST Programme. MIT Press Series on Information Systems.

ServFace (service annotations for userface composition)

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 technology. 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 existing methodologies and tools. Instead, user interfaces are designed manually for service interfaces and business processes.

The ServFace project aims at creating a model-driven service engineering methodology for an integrated development process for service-based applications. ServFace looks at this process from two different perspectives: First, the annotation of services with corresponding user interface annotations that are 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. Together with technical service descriptions like WSDL, 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 is being developed in ServFace will provide tool support for defining service annotations. Service annotations will be published in an annotation repository, from where they are globally available.

Deliverables:

- D 1.1: User and Technological Requirements and Evaluation Plan (Initial Version), 11/16/2009
- D 1.2: Service engineering methodology (initial version), 11/16/2009
- D 1.4: User and Technological Requirements and Evaluation Plan (Final Version), 11/16/2009
- D 2. 1: XML Languages for User Interface Models, 11/16/2009
- D 2. 2: Methods and Solutions for User Interface Composition, 11/16/2009
- D 2. 3: Guidelines for Multimodal User Interface Adaptation, 11/16/2009
- D 3.1: Web Service technology extension and Platform specification (initial version) ,11/16/2009
- D 3.2: Spec. of platform mapping and adaptation, service creation and management tools (initial v.), 11/16/2009
- D 3.3: Runtime Implementation and Web Service Extensions (initial version),11/16/2009

- D 3.5: Presentation-Oriented Service Architecture Meta-Models (final version), 11/16/2009
- D 4.1: Use Case and Scenario Definition, 11/16/2009
- D 4.2: Demonstrator Implementation (initial version), 11/16/2009
- D 5.1: Project Website (launch), 11/16/2009
- D 5.3: Plan of dissemination activities, 11/16/2009
- D 4.3: Report on Evaluation Results (initial version), 11/25/2009
- D 1.3: Design Patterns and Guidelines (initial version), 11/25/2009
- D 2. 4: Authoring Tools for User Interface Models and Generation (initial version), 11/25/2009
- D 2. 5: Models for Service Annotations, User Interfaces, and Servicebased Interactive Applications (interims version), 11/25/2009
- D 3.4: Implementation of Tool Environment for Platform Mapping and Service Creation (initial version), 11/25/2009
- D 2. 6: Methods and Solutions for User Interface Composition (interims version), 01/29/2010
- D 2. 8: Authoring Tools for user interface models and generation (interims version), 01/29/2010
- D 1.5: Methodology for Service-Based Interactive Application Development, 02/26/2010
- D 1.6: Annotation-related Guidelines for UI Design, 02/26/2010
- D 2. 7: Logical Description of Multimodal User Interfaces, 02/26/2010
- D 3.7: Discussion of approaches for platform mappings and runtime support (final version), 02/26/2010
- D 2. 9 Models for Service Annotations, User Interfaces, and Service-based Interactive Applications (final version), 07/13/2010
- D 3.9: Specification of the ServFace Platform Mapping, 09/27/2010
- D 3.8: Automatic Model-Based Generation of User Interfaces from WSDL Descriptions, 09/29/2010
- D 2.10: Methods and Solutions for User Interface Composition (final version), 10/15/2010
- D 4.5: Report on Evaluation Results (final version), 10/15/2010
- D 2.11: Logical Description of Multimodal User Interfaces, 11/08/2010
- D 2.12: Authoring Tools for user interface models and generation (final version), 11/08/2010
- D 5.6: Project Website (continuous updates), 11/10/2010
- D 4.6: Demonstrator Implementation (final version), 11/10/2010

Publications:

- Service discovery supported by task models, Kritikos K., Paternò F. - In: EICS'10 - 2nd ACM SIGCHI Symposium on Engineering Interactive Computing Systems (Berlin, Germany, 19-23 June 2010). Proceedings, pp. 261 - 266. ACM, 2010
- User task-based development of multi-device service-oriented applications, Paternò F., Santoro C., Spano L. D. - AVI 2010 - International Conference on Advanced Visual Interfaces. Workshop session (Rome, Italy, 25-29 May 2010). Proceedings, article n. 9. ACM
- Supporting multimodality in service-oriented model-based development environments, Manca M., Paternò F. - HCSE 2010 - 3rd Conference on Human-Centred Software Engineering (Reykjavik, Iceland, 14-15 October 2010). Proceedings, pp. 135 - 148. (Lecture Notes in Computer Science, vol. 6409). Springer

- Task-Driven Service Discovery and Selection, Kritikos K., Paternò F. - Proceedings of the International Conference on Advanced Visual Interfaces, pp 89-92, Roma, Italy, acm
- Exploiting Web service annotations in model-based user interface development, Paternò F., Santoro C., Spano L. D. - In: EICS'10 - 2nd ACM SIGCHI Symposium on Engineering Interactive Computing Systems (Berlin, Germany, 19-23 June 2010). Proceedings, pp. 219 - 224. ACM
- Deriving vocal interfaces from logical descriptions in multi-device authoring environments, Paternò F., Sisti C., In: ICWE 2010 - 10th International Conference on Web Engineering (Vienna, Austria, 5-9 July 2010). Proceedings, pp. 204 - 217. (Lecture Notes in Computer Science, vol.
- Rapid UI Development for Enterprise Applications: Combining Manual and Model-Driven Techniques, Arne Schramm, Andre Preußner, Matthias Heinrich, Lars Vogel, In Proceedings of the Model Driven Engineering Languages and Systems Conference, Springer,
- Rapid Development of Composite Applications Using Annotated Web Services, Lars Dannecker, Marius Feldmann, Tobias Nestler, Gerald Hübsch, Uwe Jugel, In Proceedings of the 6th Model-driven Web Engineering Workshop (at ICWE), 2010
- Conceptual and Usability Issues in the Composable Web of Software Services, Abdallah Namoun, Tobias Nestler, Antonella De Angeli, In Proceedings of the 2nd International Workshop on Lightweight Composition on the Web (at ICWE), Springer, 2010
- The ServFace Builder - A WYSIWYG Approach for Building Service-Based Applications, Tobias Nestler, Marius Feldmann, Gerald Hübsch, Andre Preussner, Uwe Jugel, In Proceedings of the 10th Interantional Conference on Web Engineering, Springer, 2010
- Generating Smart Wrapper Libraries for Arbitrary APIs, Uwe Jugel, Paper on a technique to generate refined APIs from legacy APIs using Eclipse Xpand/Xtend
- A Case Study on API Generation, Uwe Jugel, Andre Preußner, Presentation of the paper "A Case Study on API Generation" at SAM 2010 in Oslo, Paper yet to be published.
- Visual, In-place Data Flow Modeling, Uwe Jugel, Paper on a visual modeling technique for graphical UI builders, presented at AVI 2010 in Roma.
- An Integrated Approach for Creating Service-Based Interactive Applications, Marius Feldmann, Jordan Janeiro, Tobias Nestler, Gerald Hübsch, Uwe Jugel, Andre Preußner, Alexander Schill, An Integrated Approach for Creating Service-Based Interactive Applications; Conference in Human-Computer Interaction (Interact); 2009
- MARIA: A Universal, Declarative, Multiple Abstraction Level Language for Service-Oriented Applications in Ubiquitous Environments, Fabio Paterno', Carmen Santoro, Lucio Davide Spano, ACM Transactions on Computer-Human Interaction (TOCHI) archive Volume 16 , Issue 4 (November 2009)
- Model-Based Design of Multi-device Interactive Applications Based on Web Service, Fabio Paternò, Carmen Santoro, Lucio Davide Spano, Model-Based Design of Multi-device Interactive Applications Based on Web Service; LNCS 5726; 892 - 905; Springer Berlin/Heidelberg; 2009
- Design-Time Support to Create User Interfaces for Service Based Applications, Tobias Nestler, Marius Feldmann, Alexander Schill, Design-time support to create user interfaces for service based applications; IADIS, WWW/Internet; 2008
- Exploiting Web Services and Model-Based User Interfaces for Multi-device Access to Home Applications, Giulio Mori, Fabio Paternò, Lucio Davide

Spano, Exploiting Web Services and Model-Based User Interfaces for Multi-device Access to Home Applications; LNCS 5136; 181 - 193; Springer Berlin/Heidelberg; 2008
Designing Usable Applications Based on Web Services, Fabio Paternò, Carmen Santoro, Lucio Davide Spano, Designing Usable Applications based on Web Services; Proceedings of the International Workshop on: Interplay between Usability Evaluation and Software Development (I-USED 2008); Pisa, Italy; 2008

- End User Development of Service-based Applications, Abdallah Namoune, Tobias Nestler, Antonella De Angeli, Proceedings of the 2nd Workshop on HCI and Services (at HCI); Special Issue for BIT; 2009
- Improving task-driven software development approaches for creating service-based interactive applications by using annotated Web services, Marius Feldmann, Gerald Hübsch, Thomas Springer, Alexander Schill, Improving task-driven software development approaches for creating service-based interactive applications by using annotated Web services; 5th International Conference on Next Generation Web Services Practices; 2009
- Overview of an End User enabled Model-driven Development Approach for Interactive Applications based on Annotated Services, Marius Feldmann, Tobias Nestler, Klemens Muthmann, Uwe Jugel, Gerald Hübsch, Alexander Schill, Overview of an End User enabled Model-driven Development Approach for Interactive Applications based on Annotated Services; 4th Workshop on Emerging Web Services Technology, ACM International Conference Proceeding Series; ACM; 2009
- Overview of a Model-to-Code Transformation Approach for Generating Service-based Interactive Applications for Google Android, Marius Feldmann, Vladyslav Oleniuk, Larissa Globa, Alexander Schill, Overview of a Model-to-Code Transformation Approach for Generating Service-based Interactive Applications for Google Android; Proceedings of the 19th International Crimean Conference on Microwave and Telecommunication Technology (CriMiCo); Krim; 2009
- An Annotation Tool for Enhancing the User Interface Generation Process for Services, Paoli Izquierdo, Jordan Janeiro, Gerald Hübsch, Thomas Springer, Alexander Schill, An Annotation Tool for Enhancing the User Interface Generation Process for Services; International Crimean Conference (Crimico); 2009
- User-centric Composition of Service Front-ends at the Presentation Layer, Tobias Nestler, Lars Dannecker, Andreas Pursche, User-centric Composition of Service Front-ends at the Presentation Layer; Proceedings of the 1st International Workshop on User-generated Services (at ICSOC 2009); Springer, LNCS; 2009
- Service Composition at the Presentation Layer using Web Service Annotations, Tobias Nestler, Marius Feldmann, Andre Preußner, Alexander Schill, Service Composition at the Presentation Layer using Web Service Annotations; Proceedings of the First International Workshop on Lightweight Integration on the Web (at ICWE 2009); CEUR workshop proceedings; 2009
- Towards a Mashup-driven End-User Programming of SOA-based Applications, Tobias Nestler, Towards a Mashup-driven End-User Programming of SOA-based Applications; Proceedings of the 10th International Conference on Information Integration and Web-based Applications & Services (iiWAS); ACM and Austrian Computer Society; 11/2008
- Improving the Development of Service-Based Applications through Service Annotations, Jordan Janeiro, André Preussner, Thomas Springer, Alexander Schill, Matthias Wauer, Improving the Development of Service-Based

Applications through Service Annotations; IADIS WWW/Internet; Rome, Italy; 2009

- Enhancing User Interface Design Patterns with Design Rationale Structures,Jordan Janeiro, Simone DJ Barbosa, Thomas Springer, Alexander Schill, 2009
- Retrieval of User Interface Templates Based on Tasks, Jordan Janeiro, Thomas Springer, Simone DJ Barbosa, Alexander Schill, Retrieval of User Interface Templates Based on Tasks; Conference in Human Computer Interaction (Interact); Uppsala, Sweden; 2009
- Flexible Human Service Interfaces, Josef Spillner and Iris Braun and Alexander Schill, Flexible Human Service Interfaces; Proceedings of 9th International Conference on Enterprise Information Systems (ICEIS 2007); June 2007; Portugal
- MARIA: A Universal, Declarative, Multiple Abstraction-Level Language for Service-Oriented Applications in Ubiquitous Environments, 11/02/2009

SHAPE

The SHAPE project aims to support the development and realization of enterprise systems based on a Semantically-enabled Heterogeneous service Architecture (SHA). SHA extends Service Oriented Architectures (SOA) with semantics and heterogeneous infrastructures (Web services, Agents, Semantic Web Services, P2P and Grid) under a unified service oriented approach. To achieve this the consortium of the SHAPE project will develop a Model-Driven Engineering (MDE) tool-supported methodology and will take an active role in the standardization of metamodels and languages for SHA.

The main challenges addressed by the SHAPE project are:

- How to map the flow of business logic and data to services and functionality in a platform independent way?
- How to integrate the various models of processes, requirements, services and functions in a common model that can also be adopted by individual projects and their implementation environment?
- How to manage such models and provide links between them that can be used for service composition or managing changes?

The SHAPE project will promote a new development paradigm with a higher degree of involvement of joint users 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.

Deliverables:

- D1.1 - Case Study Scenario Descriptions and Requirements Specification
- D1.2 - Case Study Execution and Validation
- D1.3 - Case Study Execution and Validation
- D1.4 - Case Study Execution and Validation
- D2.1 - Model-Driven Methodology and Architecture Specification
- D2.2 - SHAPE Integrated and Tool Supported Methodology
- D2.3 - SHAPE Integrated and Tool Supported Methodology
- D2.4 - SHAPE Integrated and Tool Supported Methodology
- D3.1 - Consolidation of Common Basis for Metamodels and Languages for SHA
- D3.2 - Metamodel and Language for service modelling (UPMS), Extensions for Flexible Business Models and Service Variability
- D3.3 - Metamodel and Language Extension for Semantic Web Services, Agents and P2P and Grid
- D3.4 - Metamodel and Language for service modelling (UPMS), Extensions for Flexible Business Models and Service Variability
- D3.5 - Metamodel and Language Extension for Semantic Web Services, Agents and P2P and Grid
- D3.6 - Integrated Set of Models for SHA
- D4.1 - MDE Toolset and architecture specification
- D4.2 - SHAPE MDE Toolset

- D4.3 - SHAPE MDE Toolset
- D4.4 - SHAPE MDE Toolset
- D5.1 - Model Transformation and Deployment Architecture Description
- D5.2 - Model Transformations from Business Models to UMPSHA
- D5.3 - Model Transformations and Deployment - from UPMSHA to WSA, Agents, P2P, Grid, and SWS platforms
- D5.4 - Model Transformations from Business Models to UMPSHA
- D5.5 - Model Transformations and Deployment - from UPMSHA to WSA, Agents, P2P, Grid, and SWS platforms
- D6.1 - Public Website
- D6.2 - Standardisation and Dissemination Plan
- D6.3 - Plan for Collaboration with ICT SSAI&E Projects
- D6.4 - Standardisation and Dissemination Plan
- D6.5 - Plan on Use and Dissemination of Foreground (Confidential)
- D6.6 - Standardisation and Dissemination Plan
- D6.7 - Plan on Use and Dissemination of Foreground (Confidential)
- D6.8 - Standardisation and Dissemination Plan

Publications:

- Hahn, C., Nesbigall, S., Warwas, S., Zinnikus, I., Fischer K., Klusch, M. (2008): Model-driven Approach to the Integration of Multiagent Systems and Semantic Web Services. Proceedings of MDA4SOA workshop, EDOC 2008. (Details - Link)
- Hahn, C., Slomic, I, (2008): Agent-based Extensions for the UML Profile and Metamodel for Service-oriented architectures. Proceedings of MDA4SOA workshop, EDOC 2008. (Details - Link)
- Facca, F.M., Komazec, S., Zaremba, M. (2008): Towards a Semantic Enabled Middleware for Publish/Subscribe Applications. Proceedings of the International Conference on Semantic Computing, pp. 498-503, 2008 IEEE International Conference on Semantic Computing, 2008. (Details - Link)
- Roman, D., Kifer, M. (2008): Semantic Web Service Choreography: Contracting and Enactment. Proceedings of the 7th international Conference on the Semantic Web, Karlsruhe, Germany, October 26 - 30, 2008. (Details - Link)
- Hahn, C., Nesbigall, S., Warwas, S., Zinnikus, I., Fischer K., Klusch, M. (2008): Integration of Multiagent Systems and Semantic Web Services on a Platform Independent Level. Proceedings of International IEEE/ACM International Conference on Intelligent Agent Technology (IAT), Sydney, Australia, IEEE CS Press. (Details - Link)
- Fischer, K., Hahn, C., Warwas, S. (2008): Modeling TeleTruck: A Case Study. Proceedings of Engineering Societies in the Agents World IX, 9th International Workshop, ESAW 2008, Saint-Etienne, France, September 24-26, 2008, Revised Selected Papers, Lecture Notes in Computer Science, Springer Verlag, Berlin et al., 1-26. (Details - Link)
- Jacobi, S., Hahn, C., Raber, D. (2009): Integration of agent-based Manufacturing Execution Systems into a Service Oriented Architecture. Proceedings of the International Workshop on Agent-based Technologies and applications for enterprise interOPerability (ATOP 2009), Budapest, Hungary. (Details - Link)
- Jacobi, S., Hahn, C., Raber, D. (2009): MasDISPO_xt - Process optimisation by use of integrated, agent-based Manufacturing Execution Systems inside the Supply Chain of Steel Production. Proceedings of the eleventh

International Conference on Enterprise Information Systems (ICEIS 2009), Milan, Italy. (Details - Link)
- Warwas, S., Hahn, C. (2009): The DSML4MAS development environment. Proceedings of AAMAS (2) 2009: 1379-1380. (Details - Link)
- Hahn, C., Zinnikus, I., Warwas, S., Fischer, K. (2009): From agent interaction protocols to executable code: a model-driven approach. Proceedings of AAMAS (2) 2009: 1199-1200. (Details - Link)
- Warwas, S., Hahn, C., Fischer, K. (2009): A visual development environment for Jade. Proceedings of AAMAS (2) 2009: 1349-1350. (Details - Link)
- Sadovykh, A., Hahn, C., Panfilenko, D., Shafiq, O., Limyr, A. (2009): SOA and SHA Tools Developed in SHAPE Project. Proceedings of the Tools and Consultancy Track. European Conference on Model-Driven Architecture (ECMDA-09), in Conjunction with Fifth European Conference on Model-Driven Architecture Foundations and Applications, June 23-26, Enschede, Netherlands, Page 113, CTIT Proceedings Series (WP), Vol. 09-12, 6/2009. (Details - Link)
- Stollberg, M., Muth, M. (2009): Service Customization by Variability Modeling. Proceedings of the 5th International Workshop on Engineering Service-Oriented Applications (WESOA'09) co-located with the ICSOC-ServiceWave 2009, Stockholm, Sweden, November 23rd, 2009. (Details - Link)
- Elvesæter, B., Berre, A.-J. (2010): OMG Specifications for Enterprise Interoperability. IESA 2010 workshop "Standardisation - a foundation for Interoperability" Coventry, United Kingdom, April 13, 2010.
- Elvesæter, B., Landre, E., Berre, A.-J. (2010): Application of SHAPE Technologies in Production and Process Optimization. IESA 2010 Workshop "Use of MDI/SOA concepts in industry", Coventry, United Kingdom, April 13, 2010.
- Hahn, C., Dmytro, P., Fischer, K. (2010): A Model-Driven Approach to Close the Gap between Business Requirements and Agent-Based Execution. In Proc. of the 4th Workshop on Agent-based Technologies and applications for enterprise interoperability (ATOP 2010) held in conjunction with AAMAS 2010, Toronto, Canada, June 10, 2010.
- Dmytro, P., Emrich, A., Weber, S. (2010): MDA Organization Platform: A Holistic Approach for the Management of Model-Driven Architectures. In Proc. of the 4th Workshop on Modeling, Design, and Analysis for the Service Cloud (MDA4ServiceCloud'10), co-located with the 6th European Conference on Modelling Foundations and Applications (ECMFA 2010), Paris, France, June 15, 2010.
- Grammel, B., Kastenholz, S. (2010): A Generic Traceability Framework for Facet-based Traceability Data Extraction in Model-driven Software Development. In Proc. of the 6th ECMFA Traceability Workshop held in conjunction with the ECMFA 2010, Paris, France, June 15, 2010.
- Mohagheghi, P., Berre, A.-J., Sadovykh, A., Barbier, F., Benguria, G. (2010): Reuse and Migration of Legacy Systems to Interoperable Cloud Services - The REMICS project. In Proc. of the 4th Workshop on Modeling, Design, and Analysis for the Service Cloud (MDA4ServiceCloud'10), co-located with the 6th European Conference on Modelling Foundations and Applications (ECMFA 2010), Paris, France, June 15, 2010.
- Stollberg, M., Elvesæter, B., Shafran, V., Magarshak, R. (2010): A Customizable Methodology for the Model-driven Engineering of Service-based System Landscapes. In Proc. of the 4th Workshop on Modeling, Design, and Analysis for the Service Cloud (MDA4ServiceCloud'10), co-located with the 6th European Conference on Modelling Foundations and Applications (ECMFA 2010), Paris, France, June 15, 2010. (Link)

- Hahn, C., Jacobi, S., Raber, D. (2010): Enhancing the Interoperability between Multiagent Systems and Service-Oriented Architectures through a Model-Driven Approach. In Proc. of the 2010 IEEE/WIC/ACM/ International Conferences of Web Intelligence and Intelligent Agent Technology. August 31 – September 3, 2010.
- Jacobi, S., Hahn, C., Raber, D. (2010): Integration of Multiagent Systems and Service Oriented Architectures in the Steel Industry. In Proc. of the 2010 IEEE/WIC/ACM/ International Conferences of Web Intelligence and Intelligent Agent Technology. August 31 – September 3, 2010.
- Hahn, C., Jacobi, S., Raber, D. (2010): Enhancing the Interoperability between Multiagent Systems and Service-Oriented Architectures through a Model-Driven Approach. In Proc. of the 8th German Conference on Multi-Agent System Technologies (MATES 2010), Leipzig, Germany, September 27 – 29, 2010.
- Elvesæter, B., Carrez, C., Mohagheghi, P., Berre, A.-J., Johnsen, S. G., Solberg, A. (2010): Model-Driven Service Engineering with SoaML. in “Service Engineering, European Research Results”, Schahram Dustdar and Fei Li (editors), Springer Verlag, book to be published in 2010.

SLA@SOI

Launched on June 2nd, 2008, the project is committed to research, engineer and demonstrate technologies that can embed SLA-aware infrastructures into the service economy.

The Vision adopted is a “business-ready service-oriented infrastructure empowering the service economy in a flexible and dependable way.”

Deliverables:

- D.A1a Framework Architecture, October 2009
- D.A2a Business SLA Management, October 2009
- D.A3a SLA-aware Service Management, October 2009
- D.A4a SLA-Aware Infrastructure Management, October 2009
- D.A5a-M12-SLA Foundations and Management, October 2009
- D.A6a Predictable Service Engineering Methodology, October 2009
- D.B1a Scientific Evaluation Report, October 2009
- D.B2a Adhoc Demonstrator, October 2009
- D.B3a ERP Hosting, October 2009
- D.B4a Use Case Specification Enterprise IT, October 2009
- D.B5a Service Aggregator, October 2009
- D.B6a Use Case Specification eGovernment, October 2009
- D.B7a Use Case Specification Financial Grids, October 2009
- D.B8c Annual Review of External Collaborations, October 2009
- D.B9b Dissemination Plan, May 2009
- D.B9a Project Website, January 2009

Publications:

- Enterprise IT Use Case (Public release), SLA@SOI Consortium
- NEXOF-RA RESERVOIR & SLA@SOI Collaboration, RESERVOIR & SLA@SOI Consortium, NEXOF-RA
- Using Cloud Standards for Interoperability of Cloud Frameworks, Thijs Metsch, Andy Edmonds, Victor Bayon
- Statistical Inference of Software Performance Models for Parametric Performance Completions, Jens Happe, Dennis Westermann, Kai Sachs, and Lucia Kapova, Lecture Notes in Computer Science, 2010, Volume 6093/2010, pp 20-35
- Research into Practice – Reality and Gaps, Franz Brosch, Heiko Koziolek, Barbora Buhnova and Ralf Reussner, QoSA 2010, Lecture Notes in Computer Science, 2010, Volume 6093/2010, pp 36-51
- XMPP in the SLA@SOI project, Primoz Hadalin, RealTime WebCamp Ljubljana

Dynamic set-up of Monitoring Infrastructures for SLA Management, George Spanoudakis, SLA@SOI Technical Paper
- SAMI: The SLA Management Instance, Costas Kotsokalis, Ramin Yahyapour, Miguel Angel Rojas Gonzalez,ICIW 2010, Proceedings of Fifth International Conference on Internet and Web Applications and Services, pp. 303-308

- Dynamic set-up of monitoring infrastructures for service based systems, Marco Comuzzi, George Spanoudakis, ACM SAC 2010, Proceedings of the 2010 ACM Symposium on Applied Computing, pp. 2414-2421
- A Framework for Multi-level SLA Management, Marco Comuzzi, Wolfgang Theilmann, Gabriele Zacco, Christoph Rathfelder, Constantinos Kotsokalis, and Ulrich Winkler, NFPSLAM 2009, Lecture Notes in Computer Science, 2010, Volume 6275/2010, pp. 187-196
- Translation of Service Level Agreements: A Generic Problem Definition, Costas Kotsokalis, Ulrich Winkler, NFPSLAM 2009, Lecture Notes in Computer Science, 2010, Volume 6275/2010, pp. 248-257
- Performance Evaluation of Scheduling Policies in Symmetric Multiprocessing Environments, Jens Happe, Henning Groenda, and Ralf H. Reussner, MASCOTS 2009
- SLA-aware Resource Management, Yih Leong Sun, Andy Edmonds, John Kennedy, Victor Bayon, Primoz Hadalin,
- Parametric Performance Completions for Model-Driven Performance Prediction, Jens Happe, Steffen Becker, Christoph Rathfelder, Holger Friedrich, and Ralf H. Reussner, Performance Evaluation Journal, Volume 67, Issue 8, August 2010, pp. 694-716
- Black-Box Performance Models: Prediction based on Observation, Jens Happe, Hui Li, and Wolfgang Theilmann,
- Modeling Service Level Agreements with Binary Decision Diagrams, Costas Kotsokalis, Ramin Yahyapour, Miguel Angel Rojas Gonzalez, ICSOC / ServiceWave 2009, Lecture Notes in Computer Science, 2009, Volume 5900/2009, pp. 190-204
- A Model-Driven Approach for Monitoring Business Performance in Web Service Compositions, Christof Momm, Michael Gebhart, Sebastian Abeck,
- Modeling Event-Driven Service-Oriented Systems using the Palladio Component Model, Christoph Rathfelder, Samuel Kounev, QUASOSS
- Model-based Performance Prediction for Event-driven Systems, Christoph Rathfelder, Samuel Kounev, DEBS 2009
- Challenges in SLA Translation, Hui Li, Wolfgang Theilmann, Jens Happe, SLA@SOI Technical Paper
- Establishing and Monitoring SLAs in Complex Service-Based Systems, Marco Comuzzi, George Spanoudakis, Costas Kotsokalis, Ramin Yahyapour, ICWS 2009, Proceedings of 2009 IEEE International Conference on Web Services, pp. 783-790
- Web Service Trust: Towards a Dynamic Assessment Framework, George Spanoudakis, Stephane LoPresti, ARES 2009, Proceedings of Intl. Conference on Availability, Reliability and Security, 2009, pp. 33-40
- Parameter Dependencies for Component Reliability Specifications, Heiko Koziolek, Franz Brosch, FESCA 2009, ETAPS 2009
- Describing and Verifying Monitoring Capabilities for SLA-Driven Service Based Systems, Marco Comuzzi, George Spanoudakis, CAiSE 2009
- A Framework for Hierarchical and Recursive Monitoring of Service Based Systems, Marco Comuzzi, George Spanoudakis, ICIW 2009
- Multi-level SLAs for Harmonized Management in the Future Internet, Wolfgang Theilmann, Luciano Baresi, Towards the Future Internet – A European Research Perspective

- A Model-Driven Approach for Monitoring Business Performance in Web Service Compositions, Christof Momm, Thomas Detsch, Sebastan Abeck, Advances in Quality of Service Management Workshop, EDOC 2008
- Multi-level SLA Management for Service-Oriented Infrastructures, Wolfgang Theilmann, Ramin Yahyapour, Joe Butler, Servicewave 2008
- Quality Considerations in SAP Architectures, Wolfgang Theilmann, Industrial Experience Report Track, CompArch 2008

SmartLM

Current software licensing practices are limiting the acceleration of Grid adoption. Indeed, the rapid emergence of service and virtualization environments requires a rapid evolution in licensing models. *SmartLM* will provide a generic and flexible licensing virtualization technology for new service-oriented business models across organization boundaries.

Deliverables:

- D1.1 Functional requirements of the new licensing architecture for Grids
- D2.1 Software Licensing Panorama Today
- D2.2 Report on business impact in adoption of Grid licensing mechanisms and suggested Business Models
- D3.1 Design of the license service, specification of WS-Negotiation protocol
- D4.1 Process-Model for Accounting
- D6.1 Scenarios, criteria and methodology for evaluation
- D7.3 Product description and Market context analysis for exploitation
- D7.4 Preliminary exploitation plan

Publications:

Whitepapers

- The Business Side of Software Licensing
- Architectural overview of the SmartLM innovative license management system

Booklets

- SmartLM – Grid-friendly software licensing for location independent application execution

Articles

- Cacciari, Claudio; D'andria, Francesco; Hagemeier, Björn; Mallmann, Daniel; Martrat, Josep; Rumpl, Angela; Ziegler, Wolfgang; Zsigri, Csilla: elasticLM – Software License Management for Distributed Computing Infrastructures. In: ERCIM News, volume 83.
- Cacciari, Claudio; D'andria, Francesco; Hagemeier, Björn; Mallmann, Daniel; Martrat, Josep; Peréz, David; Rumpl, Angela; Ziegler, Wolfgang; Zsigri, Csilla: elasticLM - A novel approach for software licensing in distributed computing infrastructures. In: Proceedings of the 2nd IEEE International Conference on Cloud Computing Technology and Science, Indianapolis, USA, Nov. 30-Dec. 3, 2010.

- Battré, Dominic; Brazier, Frances M.T.; Clark, Kassidy P.; Oey, Michael; Papaspyrou, Alexander; Wäldrich, Oliver; Wieder, Philipp; Ziegler, Wolfgang: A proposal for WS-Agreement Negotiation. In: Association for Computing Machinery u.a.: Grid 2010 : 11th ACM/IEEE International Conference on Grid Computing, 25-29 October 2010 in Brussels, Belgium. 2010
- Cacciari, Claudio; D'andria, Francesco; Hagemeier, Björn; Mallmann, Daniel; Martrat, Josep; Rumpl, Angela; Ziegler, Wolfgang; Zsigri, Csilla: SLA-based management of software licenses as web service resources in distributed environments. In: Euro-Par 2010 : Ischia, Italien, 31.8.-3.9. 2010. Berlin [u.a.] Springer, 2010
- Cacciari, Claudio; D'andria, Francesco; Hagemeier, Björn; Mallmann, Daniel; Martrat, Josep; García Peréz, David; Rumpl, Angela; Ziegler, Wolfgang; Zsigri, Csilla: Software licenses as mobile objects in distributed computing environments. In: Euro-Par 2010 : Ischia, Italien, 31.8.-3.9. 2010. Berlin [u.a.] : Springer, 2010
- Rumpl, Angela; Wäldrich, Oliver; Ziegler, Wolfgang: Extending WS-AGREEMENT with multi-round negotiation capability. In: Institute of Electrical and Electronics Engineers u.a.: Grid 2009 : the 10th IEEE / ACM International Conference on Grid Computing, October 12-16, 2009 in Banff, Alberta, Canada. New York, NY [u.a.] : Springer, 2010 (CoreGrid series).
- Battré. Domonic (Ed.); Wieder, Philipp (Ed.); Ziegler, Wolfgang (Ed.); Open Grid Forum: WS-Agreement Specification Version 1.0 Experience Document. Open Grid Forum, 2010 (OGF document series)
- Rana, Omar; Ziegler, Wolfgang: Research challenges in managing and using Service Level Agreements. In: Desprez, F. (Hrsg.) u.a.; European Research Consortium for Informatics and Mathematics: Grids, P2P and services computing : based on the CoreGrid ERCIM Working Group Workshop on Grids, P2P and Service Computing in Conjunction with EuroPar 2009, August 24th, 2009 in Delft, The Netherlands. New York, NY [u.a.] : Springer, 2010 (CoreGrid series 12).
- Francesco D'andria, Josep Martrat, Csilla Zsigri, Daniel Mallmann, Björn Hagemeier, Claudio Cacciari: An Enhanced Strategy for License Management in a Dynamic Grid Environment. eChallenges e-2009 Conference Proceedings ISBN: 978-1-905824-13-7
- Hagemeier, Björn; Mallmann, Daniel; Ziegler, Wolfgang: Securing software licenses in a location independent manner. In: Bubak, Marian (Hrsg.) u.a.: Cracow Grid Workshop '09 : October 12 - 14, 2009, Cracow, Poland ; proceedings. Cracow : Academic Computer Centre CYFRONET AGH, 2010, S. 112-119
- Csilla Zsigri : SmartLM's flexible licensing virtualization technology makes licenses mobile objects The 451 Group, July 2009, Link, PDF
- Daniel Mallmann, Josep Martrat, Wolfgang Ziegler: Smart LM: Grid-friendly Software Licensing for Location independent Application Execution. inSiDE, Vol. 6 No. 1, Spring 2008, Link
- Jiadao Li; Oliver Wäldrich; Wolfgang Ziegler: Towards SLA-based Software Licenses and License Management in Grid Computing. CoreGRID Symposium, Las Palmas de Gran Canaria, Gran Canaria, Spain, August 2008, Springer, 2008, CoreGRID Series, p. 139 - 152.
- M. Riedel, D. Mallmann, A. Streit: Usage of the UNICORE Grid Technology in Scientific und Economic Domains, Proceedings of the Workshop on Grid and Scientific and Engineering Applications (Grid&SEA) at Engineering and Economics (AMEE'08), June 2008, Sozopol, Bulgaria, AIP Conference Proceedings 1067, pp. 559-566, Link
- Csilla Zsigri: SmartLM. 451 MIS 2008 Review EU Research, 5 December 2008

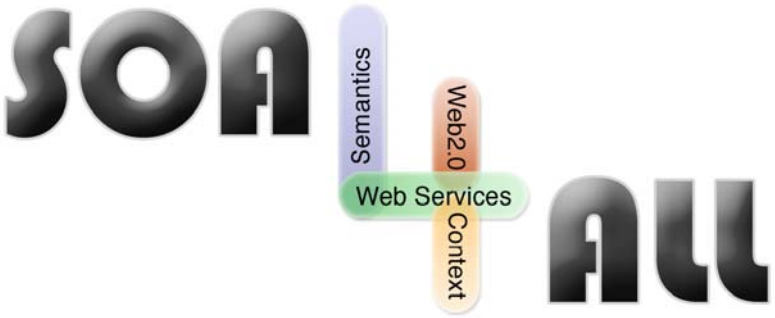

SOA4ALL

Service Oriented Architectures for All (SOA4All) aims at realizing a world where billions of parties are exposing and consuming services via advanced Web technology: the main objective of the project is to provide a comprehensive framework that integrates complementary and evolutionary technical advances (i.e., SOA, context management, Web principles, Web 2.0 and semantic technologies) into a coherent and domain-independent service delivery platform.

Deliverables:

- D0.3 SOA4All: Moving Beyond Related EU Projects
- D1.1.1 Design Principles for a Service Web v1
- D1.2.1 WSMO grounding in SAWSDL
- D1.3.1 Semantic Spaces: A Unified Semantic Data Coordination Infrastructure
- D1.3.2A Distributed Semantic Spaces: A Scalable Approach To Coordination
- D1.3.2B Distributed Semantic Spaces: A First Implementation
- D1.3.3A A Distributed Semantic Marketplace
- D1.4.1A SOA4All Reference Architecture Specification
- D1.4.1B SOA4All Runtime
- D1.4.2A Final SOA4All Reference Architecture Specification
- D1.5.1 SOA4All Testbeds Specification And Methodology
- D1.5.2 Setup SOA4All Testbeds
- D2.1.1 Service Provisioning Platform Design
- D2.1.2 Service Modelling Tools Design
- D2.1.3 Service Provisioning Platform First Prototype
- D2.2.1 Service Consumption Platform Design
- D2.2.2 Service Consumption Platform First Prototype
- D2.3.1 Service Monitoring and Management Tool Suite Design
- D2.3.2 Service Monitoring and Management Tool Suite First Prototype
- D2.4.1 SOA4All Studio First Demonstrator & Interface Specification
- D2.4.2 First Demonstrator & Interface Specification
- D2.6.1 Specification of the SOA4All Process Editor
- D2.7.1 Recommender System First Prototype
- D3.1.1 Defining the features of the WSML-Quark language
- D3.1.2 Defining the features of the WSML-Core v2.0 language
- D3.1.3 Defining the features of the WSML-DL v2.0 language
- D3.1.4 Defining the features of the WSML-Rule v2.0 language
- D3.2.1 Framework and APIs for integrated reasoning support
- D3.2.2 First Prototype Reasoner for WSML-Core v2.0
- D3.2.3 First Prototype Rule Reasoner for WSML-Rule v2.0
- D3.2.4 First Prototype Description Logic Reasoner for WSML-DLv2.0
- D3.2.5 Second Prototype Repository Reasoner for WSML-Core v2.0

- D3.3.1 Ontology Instantiation State of the Art Report
- D. D3.4.2 WSMO-Lite: Lightweight Semantic Descriptions for Services on the Web
- D3.4.3 MicroWSMO and hRESTS
- D3.4.4 WSMO Data Grounding Component
- D3.4.5 Semantic Annotations For WS-Policy (SA-Policy)
- D3.4.6 MicroWSMO v2 – Defining the second version of MicroWSMO as a systematic approach for rich tagging
- D3.4.7 Defining extensions to WSMO for capturing contextual information
- D5.1.1 State of the Art Report On Service Description and Existing Discovery Techniques
- D5.1.3 Second Crawler Prototype
- D5.3.1 First Service Discovery Prototype
- D5.4.1 First Service Ranking Prototype
- D6.1.1 State of the Art Report and Requirements for Service Construction
- D6.3.1 First Specification of Lightweight Process Modelling Language
- D6.3.2 Advanced Specification Of Lightweight, Context-aware Process Modelling Language
- D6.4.1 Specification and First Prototype of Service Composition and Adaptation Environment
- D6.4.2 Advanced Prototype For Service Composition and Adaptation Environment
- D6.5.1 Specification and first prototype of the composition framework
- D6.5.2 Advanced Prototype For Adaptive Service Composition Execution
- D7.3 End User Service Design
- D7.6 End User Service Annotation and Context Descriptions
- D8.1 Web21C Future Requirements
- D8.2 Semantic Telco Analysis
- D8.3 Web21c Futures Design
- D8.4 Web 21c Prototype v1
- D8.5 Telco 2.0 Recommendations
- D9.1.1 Future C2C eCommerce Requirements and Scenario Descriptions
- D9.2.1 eCommerce Framework Infrastructure Design
- D9.3.1 C2C e-Commerce Prototype v1
- D11.1.1 OASIS Specifications for SEE
- D11.1.2 Initialise OASIS WG on SEE With SOA4All Extensions
- D11.2.1 Initialise working group on Aligning WSMO With Other Approaches To W3C
- D12.1.1 Initial Dissemination Strategy
- D12.2.1 SOA4All Branding
- D12.2.2 First Generation of SOA4All Publicity Material
- D12.2.3 Second Generation of SOA4All Publicity Material
- D12.3.1 Initial SOA4All Website
- D12.3.2 First Update of the SOA4All Website with SOA4All Technology
- D12.6.1 Definition of Collaboration Activities
- D12.6.2 First Collaboration Activity Report
- D12.6.3 Second Collaboration Activity Report
- D13.1.1 Baseline Training Tutorial
- D13.2.1 First External SOA4All Tutorial
- D13.3.1 First SOA4All Online Webcast

Publications:

Book chapters

- Pedrinaci C., Lambert D., Maleshkova M., Liu D., Domingue J., Krummenacher R. Adaptive Service Binding with Lightweight Semantic Web Services In Service Engineering: European Research Results (S. Dustdar and F. Li eds.), Springer, 2010
- Pedrinaci, C., Grenon, P., Galizia, S., Gugliotta, A., Domingue, J. Enabling Context-Aware Web Services: Methods, Architectures, and Technologies eds. Michael Sheng, Jian Yu, and Schahram Dustdar, Chapman and Hall/CRC, 2010
- Krummenacher R., Domingue, J., Pedrinaci, C., Simperl E. SOA4All: Towards a Global Service Delivery Platform Towards the Future Internet - Emerging Trends from European Research, Eds. Georgios Tselentis, Alex Galis, Anastasius Gavras, Srdjan Krco, Volkmar Lotz, Elena Simperl, Burkhard Stiller, Theodore Zahariadis, IOS Press, 2010
- Domingue, J., Fensel, D., Davies, J., González-Cabero, R., Pedrinaci, C. The Service Web: a Web of Billions of Services In Towards a Future Internet: a European Research Perspective - Eds. Tselentis, G., Domingue, J., Galis, A., Gavras, A., Hausheer, D., Krco, S., Lotz, V. and Zehariadis, T., IOS Press, May, 2009

Conferences and symposia

- Namoune, A., Wajid, U., Mehandjiev, N. A Comparative Study: Application Development by Ordinary End Users and IT Professionals ServiceWave 2010
- Speiser, S., Studer, R. A Self-Policing Policy Language The Semantic Web (ISWC 2010)
- Mehandjiev, N., Lecue, F., Wajid, U., Namoune, A. Assisted Service Composition for End-Users 8th IEEE European Conference on Web Services (ECOWS 2010)
- Lecue, F. Combining Collaborative Filtering and Semantic Content-based Approaches to Recommend Web Services IEEE International Conference on Semantic Computing (ICSC 2010)
- Mehandjiev, N., Namoune, A., Wajid, U., Macaulay, L., Sutcliffe, Al End User Service Composition – Perceptions and Requirements 8th IEEE European Conference on Web Services (ECOWS 2010)
- Baude, F., Filali, I., Huet, F., Legrand, V., Mathias, E., Merle, P., Krummenacher, R., Simperl, E., ESB Federation for Large-Scale SOA 25th ACM Symposium On Applied Computing, March 2010
- Lecue, F., Delteil, A., Leger, A. Extending Web Service Composition Languages with Semantic Data Flow Proceedings of the 3rd IEEE International Conference on Semantic Computing (ICSC09), Berkeley, CA, USA, September 14-16, 2009
- Hausheer, D., Ristol Jorba, S., et al. Future Internet Socio-Economics – Challenges and Perspectives FIA event, Madrid, Spain, December 2008
- Lecue, F., Wajid, U., Mehandjiev, N. Negotiating Robustness in Semantic Web Service Composition In Proceedings of the 7th IEEE European Conference on Web Services (ECOWS09), Eindhoven, The Netherlands, November 9-11, 2009
- Lecue, F. Optimizing QoS-Aware Semantic Web Service Composition In Proceedings of the 8th International Semantic Web Conference (ISWC 2009), Westfields Conference Center near Washington, DC, 25-29 October, 2009
- Mehandjiev, N., Lecue, F., Wajid, U. Provider-Composer Negotiations for Semantic Robustness in Service Compositions In Proceedings of the seventh International Joint Conference on Service Oriented Computing (ICSOC 2009), November 24-27, 2009, Stockholm, Sweden

- Heymans, S., Toma, I. Ranking Services using Fuzzy HEX Programs In Proceedings of the 2nd International Conference on Web Reasoning and Rules Systems (RR 2008), 31 October - 2 November 2008, Karlsruhe, Germany. LNCS 5341, pp. 181-196
- Speiser, S. Semantic Annotations for WS-Policy IEEE International Conference on Web Services (ICWS 2010)
- Agarwal, S., Rudolph, S., Abecker, A. Semantic Description of Distributed Business Processes In Knut Hinkelmann, Andreas Abecker, Harold Boley, John Hall, Martin Hepp, Amit Sheth, Barbara Thönssen, AAAI Spring Symposium - AI Meets Business Rules and Process Management. Stanford, USA, March 2008
- Kopeck_, J., Simperl, E. Semantic Web Service Offer Discovery for E-commerce In Proceedings of the 10th International Conference on Electronic Commerce (ICEC 2008), Innsbruck, Austria, August 2008
- Hoxha, J., Agarwal, S. Semi-automatic Acquisition of Semantic Descriptions of Processes in the Web Proceedings of The International Conference on Web Intelligence, IEEE, September, 2010
- Agarwal, S. Semi-Automatic Acquisition of Semantic Descriptions of Web Sites In Proceedings of the Third International Conference on Advances in Semantic Processing, IEEE, Sliema, Malta, October, 2009
- Domingue, J., Fensel D., González-Cabero, R. SOA4All, Enabling the SOA Revolution on a World Wide Scale In Proceedings of the 2nd IEEE International Conference on Semantic Computing (ICSC 2008), Santa Clara, CA, USA, August 2008
- Lecue, F., Gonzalez, R., Gorronogoitia, Y., Radzimski, M., Villa, M. SOA4All: An Innovative Integrated Approach to Services Composition In Proceedings of the IEEE International Conference on Web Services (ICWS 2010), July 5-10, 2010, Miami, FL, USA. IEEE Computer Society 2010
- Lecue, F., Gonzalez, R., Gorronogoitia, Y., Radzimski, M., Villa, M. SOA4All: An Innovative Integrated Approach to Services Composition VI Jornadas Científico-Técnicas en Servicios Web y SOA (JSWEB2010)
- Krummenacher, R., Norton, B., Simperl, E., Pedrinaci, C. SOA4All: Enabling Web-scale Service Economies 3rd IEEE International Conference on Semantic Computing (ICSC) 2009, IEEE Computer Society
- Agarwal, S., Junghans, M. Swapping Out Coordination of Web Processes to the Web Browser The 8th IEEE European Conference on Web Services, ECOWS 2010
- Duke, A., Stincic, S., Davies, J., Lecue, F., Mehandjiev, N., Pedrinaci, C., Maleshkova, M., Domingu Telecommunication mashups using RESTful services ServiceWave 2010
- Alvaro, G., Martinez, I., Ruiz, C. Towards a lighter way of consuming semantically enriched services V edición de las Jornadas Científico-técnicas en Servicios Web y SOA (JSWEB09), Madrid, 30 September - 1 October, 2009 (also available the SOA4All Studio demo explanation paper)
- Un, P., Vogel, J. Towards Lighweight Semantic Metadata Annotation and Management for Enterprise Services I-KNOW 2010, Graz
- Krummenacher, R., Norton, B., Marte, A. Towards Linked Open Services and Processes 3rd Future Internet Symposium, Berlin, September 2010
- Junghans, M., Agarwal, S., Studer, R. Towards Practical Semantic Web Service Discovery 7th Extended Semantic Web Conference (to appear), 2010, Heraklion, Crete, Greece
- Lecue, F., Mehandjiev, N. Towards Scalability of Quality Driven Semantic Web Service Composition In Proceedings of the 7th IEEE International Conference on Web Services (ICWS09), July 6-10, 2009, Los Angeles, CA, USA

- Yu, H., Benn, N., Dietze, S., Ronald Siebes, Pedrinaci, C., Liu, D., Lambert, D., Domingue, J. Two-staged approach for semantically annotating and brokering TV-related services The IEEE International Conference on Web Services (ICWS), Miami, Florida, 2010
  Speiser, S., Studer, R. Usage Policies for Document Compositions Extended Semantic Web Conference, (ESWC 2010)
- Liu, L., Lecue, F., Mehandjiev, N., Dong-Ling Xu Using Context Similarity for Service Recommendation IEEE International Conference on Semantic Computing (ICSC 2010)
- Toma, I., Ding, Y., Chalermsook, K., Simperl, E., Fensel, D. Utilizing Web 2.0 in Web Service Ranking In Proceedings of the 3rd International Conference on Digital Society (ICDS 2009), Cancun, Mexico, February 1-7, 2009
- Baader, F., Liu, H., ul Mehdi, A. Verifying Properties of Infinite Sequences of Description Logic Actions In Proceedings of the European Conference on Artificial Intelligence (ECAI), Lisbon, Portugal, 2010 (to appear)
- Steinmetz, N., Lausen, H., Brunner, M. Web Service Search on Large Scale In Proceedings of the International Conference on Service Oriented Computing (ICSOC) 2009, Joint with the ServiceWave Conference 2009, Stockholm, Sweden, November 2009
- Vitvar, T., Kopeck_, J., Viskova, J., Fensel, D. WSMO-Lite Annotations for Web Services In The Semantic Web: Research and Applications, 5th European Semantic Web Conferences (ESWC 2008), pp. 674-689, Springer LNCS 5021. Tenerife, Spain, June 2008
- Kopecky, J., Vitvar, T. WSMO-Lite: Lowering the Semantic Web Services Barrier with Modular and Light-Weight Annotations In Proceedings of the 2nd IEEE International Conference on Semantic Computing (ICSC 2008), Santa Clara, CA, USA, 4-7 August, 2008
- Akhtar, W., Kopeck_, J., Krennwallner, T., Polleres, A. XSPARQL: Traveling between the XML and RDF worlds - and avoiding the XSLT pilgrimage In The Semantic Web: Research and Applications, 5th European Semantic Web Conference (ESWC 2008), pp. 432-447, Springer LNCS 5021, June 2008, Tenerife, Spain

Publications in journal

- Lecue, F., Mehandjiev, N. Computing Scalable Semantic Service Composition in the Web with Appropriate Quality International Journal of Web Services Research (JWSR)
- Benjamins, V., Davies, J., Baeza-Yates, R., Mika, P., Zaragoza, H., Greaves, M., et al Near-Term Prospects for Semantic Technologies IEEE Intelligent Systems, January/February 2008, vol. 23, no. 1, pp. 76-88
- Domingue, J., Fensel, D. PSMs in a Global Networked Age AIEDAM Special Issue Problem Solving Methods: Past, Present and Future, 23, 3, Cambridge University Press, 2009
- Lecue, F., Mehandjiev, N. Seeking Quality of Web Service Composition in a Semantic Dimension IEEE Transactions on Knowledge and Data Engineering (TKDE)
- Abels, S. SOA4All: Enabling a Web of Billions of Services In International Journal of Interoperability in Business Information Systems (IBIS), Issue 3 (1), 2009, pp. 35-38
- Pedrinaci, C., Domingue, J. Toward the Next Wave of Services: Linked Services for the Web of Data Journal of Universal Computer Science, 2010 (to appear)
- Lecue, F., Delteil, A., Leger, A. Towards a Semantic State Transition System for Automated Generation of Data Flow in Web Service Composition In

International Journal of Semantic Computing (IJSC) 2009, Vol 3 No 4 p-499-526

- Davies, J., Domingue, J., Pedrinaci, C., Fensel, D., Gonzalez-Cabero, R., Potter, M., Richardson, M. Towards the Open Service Web BT Technology Journal, Vol. 26, No. 2, February 2009. Springer
- Pedrinaci, C., Simperl, E., Krummenacher, R., Norton, B. Towards Web-scale SOA The Architecture Journal, 21, Microsoft, 2009
- Ding, Y., Jacob, E.K., Fried, M., Toma, I., Yan, E., Foo, S., Milojevic, S. Upper tag ontology for integrating social tagging data In Journal of the American Society for Information Science and Technology (JASIST), Vol 61, No. 3, March 2010

Papers and other publications

- Baude, F., Legrand-Cortes, V. Exécution distribuée et agile de compositions de services Distributed and agile execution of service compositions, a contribution for supporting millions of Web 2.0 services over Internet. Journées du GDR CNRS Génie de la Programmation et du Logiciel, Groupe de Travail COSMAL (Composants Objets Services : Modèles, Architectures et Langages). Toulouse, France, January 28-30, 2009
- Kopecky, J., Pedrinaci, C. Lightweight Semantic Extensions for SOA SemanticUniverse.com
- Pedrinaci, C., Domingue, J. Web Services Are Dead. Long Live Internet Services SOA4All White Paper, 2010
- Fensel, D., Fischer, F., Kopeck_, J., Krummenacher, R., Lambert, D., Vitvar, T. WSMO-Lite: Lightweight Semantic Descriptions for Services on the Web W3C Member Submission, 23 August 2010

Publications in workshops

- Wajid, U., Namoune, A., Mehandjiev, N. A Comparison of Three Service Composition Approaches for End Users EUD4Services Workshop (co-located with AVI 2010)
- Liu, D., Pedrinaci, C., Domingue, J. A Framework for Feeding Linked Data to Complex Event Processing Engines The 1st International Workshop on Consuming Linked Data (COLD 2010) at the 9th International Semantic Web Conference (ISWC 2010), Shangai, China
- Maleshkova, M., Kopecky, J., Pedrinaci, C. Adapting SAWSDL for Semantic Annotations of RESTful Services Workshop: Beyond SAWSDL at OnTheMove Federated Conferences & Workshops, Vilamoura, Portugal, 2009
- Cavallaro, L., Ripa, G., Zuccalà, M. Adapting Service Requests to Actual Service Interfaces through Semantic Annotations In Proceedings of Principles of Engineering Service Oriented Systems (PESOS 2009), ICSE 2009 Workshop, May 18-19 2009, Vancouver, Canada
- De Giorgio, T., Ripa, G., Zuccalà, M. An Approach to Enable Replacement of SOAP Services and REST Services in Lightweight Processes 2nd International Workshop on Lightweight Composition on the Web (ComposableWeb 2010), of the 10th International Conference on Web Engineering (ICWE 2010), Vienna, Austria, 7-9 July 2010
- Turati, T., Cerizza, D., Celino, I., Della Valle, E. Analyzing User Actions within a Web 2.0 Portal to Improve a Collaborative Filtering Recommendation System WPRRS (Personalization, Reputation and Recommender Systems) 2009, The 2009 IEEE / WIC / ACM International Conferences on Web Intelligence (WI09), September 2009
- Dell`Aglio, D., Celino, I., Cerizza, D. Anatomy of a Semantic Web-enabled Knowledge-based Recommender System In Proceedings of the 4th

international workshop Semantic Matchmaking and Resource Retrieval in the Semantic Web, at the 9th International Semantic Web Conference 2010, Shangai, China, November 2010
- Krummenacher, R., Norton, B. Consuming Dynamic Linked Data The 1st International Workshop on Consuming Linked Data (COLD 2010) at The 9th International Semantic Web Conference (ISWC 2010), Shanghai, China
- Born, M., Kirchner, J., Mueller, J.P. Context-driven Business Process Modeling 4th International Workshop on Technologies for Context-Aware Business Process Management (TCoB 2009), Milan, Italy, May 6, 2009 - In conjunction with the 11th International Conference on Enterprise Information Systems (ICEIS 2009)
- Michels, C., Agarwal, S. Elicitation of Preferences for Web Service Compositions 5th International Workshop on Applications of Semantic Technologies
- Schnabel, F., Gorronogoitia, Y., Radzimski, M., Lecue, F., Mehandjiev, N., Ripa, G., Abels, S., Bloo Empowering Business Users to Model and Execute Business Processes The Third Workshop on Business Process Management and Social Software, BPMS2\'10, 13 September, 2010
- Krummenacher, R., Norton, B. Geospatial Linked Open Services Proceedings of the Workshop “Towards Digital Earth: Search, Discover and Share Geospatial Data 2010_ at Future Internet Symposium (FIS), Berlin, Germany, September 20, 2010
- Pedrinaci, C., Liu, D., Maleshkova, M., Lambert, D., Kopecky, J., Domingue, J. iServe: a Linked Services Publishing Platform Workshop: Ontology Repositories and Editors for the Semantic Web at the 7th Extended Semantic Web Conference (ESWC) 2010, Heraklion, Greece, 30 May - 3 June 2010
- Mathias, E., Baude, F. Multi-domain Grid/Cloud Computing Through a Hierarchical Component-Based Middleware In 8th International Workshop on Middleware for Grids, Clouds and e-Science (MGC\'2010), in conjunction with 11th Int. Middleware Conference November 2010
- Mos, A., Pedrinaci, C., Alvaro Rey, G., Gomez, J.M., Liu, D., Vaudaux-Ruth, G., Quaireau, S. Multi-Level Monitoring and Analysis of Web-Scale Service Based Applications In Proceedings of the 2nd Workshop on Monitoring, Adaptation and Beyond (MONA+), ICSOC/ServiceWave 2009, Stockholm, Sweden, November 2009
- Steinmetz, N., Lausen, H. Ontology-based Feature Aggregation for Multi-Valued Ranking In Proceedings of the 3rd Workshop on Non-Functional Properties and SLA Management in Service-Oriented Computing, November 2009, Stockholm, Sweden. Springer
- Pedrinaci, C., Domingue, J. Ontology-based Metrics Computation for Business Process Analysis In Proceedings of the 4th International Workshop on Semantic Business Process Management (SBPM2009), Jun 1, 2009 - colocated with ESWC 2009, 31 May - 4 June 2009, Heraklion, Greece
- De Giorgio, T., Ripa, G., Zuccalà, M. SAWSDL for Self-Adaptive Service Composition Beyond SAWSDL 2009 Workshop, at the OnTheMove Conference, Vilamoura, Portugal, November 2009
- Di Nitto, E., Ripa, G. Self-Organization and SOA for Networked Enterprises 1st International Future Enterprise Systems Workshop, co-located with the 3rd Future Internet Symposium Berlin, Germany, 20 September, 2010
- Maleshkova, M., Pedrinaci, C., Domingue, J. Semantic Annotation of Web APIs with SWEET 6th Workshop on Scripting and Development for the Semantic Web at Extended Semantic Web Conference (ESWC2010), Heraklion, Greece, 30 May - 3 June, 2010
- Liu, D., Pedrinaci, C., Domingue, J. Semantic Enabled Complex Event Language for Business Process Monitoring In Proceedings of the 4th

International Workshop on Semantic Business Process Management (SBPM2009), Jun 1, 2009 - colocated with ESWC 2009, 31 May - 4 June 2009, Heraklion, Greece

- Pedrinaci, C., Lambert, D., Wetzstein, B., Lessen, T., Cekov, L., Dimitrov, M. SENTINEL: A Semantic Business Process Monitoring Tool In Proceedings of the 1st international workshop on Ontology-supported business intelligence (OBI 2008) at the 7th International Semantic Web Conference (ISWC 2008), Karlsruhe, Germany, October 2008
- Pedrinaci, C., Domingue, J., Krummenacher R. Services and the Web of Data: An Unexploited Symbiosis Workshop: Linked AI: AAAI Spring Symposium - Linked Data Meets Artificial Intelligence, Stanford, CA, USA, 22-24 March 2010
- Maleshkova, M., Pedrinaci, C., Domingue, J. Supporting the Creation of Semantic RESTful Service Descriptions Workshop: Service Matchmaking and Resource Retrieval in the Semantic Web (SMR2) at 8th International Semantic Web Conference, Proceedings of ISWC09, Washington D.C., USA, 2009
- Duke, A., Stincic, S., Davies, J., Lecue, F., Mehandjiev, N., Pedrinaci, C., Maleshkova, M., Domingu Telecommunication mashups using RESTful services IEEE 2010 First International Workshop on Service-Oriented Computing and Multi-Agent Systems (SOCMAS 2010) Miami, USA, in conjunction with IEEE 6th World Congress on Services (SERVICES 2010)
- Lombardi, J.P., Vogel, J., Duebon, K. Wizard-based Process Modeling for Business Users EUD4Services Workshop, Rome, Italy, 25 May, 2010, in conjunction with AVI 2010

STREAM

STREAM Project aims at producing a highly scalable middleware platform able to process in real time massive data streams such as the IP traffic of an organization, the output of a large sensor network, the e-mail processed by an ISP, the market feeds from stock exchange and financial markets, the calls in a telco operator, credit card payments, etc.
This will enable a myriad of new services and applications in the upcoming Internet of Services.
A growing number of applications requires the ability to analyze massive amounts of streaming data in real-time. Examples of such applications are: market data feed processing of the output of large scale ad-hoc networks, etc. The project aims at providing a highly scalable cloud computing platform to enable a new breed of services. The core is the data streaming platform, StreamCloud, that will be able to parallelize the processing of information flows in large clusters of 100s sites. Current approaches fail to scale for massive information flows. Stream aims at boosting the scalability of current approaches in 1 to 2 orders of magnitude. Stream platform will provide elastic computing, so the computing resources as used as required by the incoming load. Below the core, there is a high performance communication layer that enables an efficient interaction among sites with access between node memories of tens of microseconds in contrast with tens of milliseconds using current technology. Additionally, this layer will provide parallel IO and low cost storage for huge amounts of information. Above the core, there is a data mining layer offering higher level services to ease the development of applications processing the information flows. On top of the data streaming platform there is the application layer in which user applications & services will run.

Publications:

- Reducing Disk I/O Performance Sensitivity for Large Numbers of Sequential Streams, George Panagiotakis, Michail D. Flouris, and Angelos Bilas. Reducing Disk I/O Performance Sensitivity for Large Numbers of Sequential Streams. Proceedings of The 29th Int'l Conference on Distributed Computing Systems (ICDCS'09).June 22-26, 2009. Montreal, Canada.
- Minimizing Latency in Fault-Tolerant Distributed Stream Processing Systems, Andrey Brito, Christof Fetzer, "Minimizing Latency in Fault-Tolerant Distributed Stream Processing Systems". Proceedings of The 29th Int'l Conference on Distributed Computing Systems (ICDCS'09). June 22-26, 2009. Montreal, Canada.
- "StreamMine", Book chapter of Handbook of Research on Advanced Distributed Event-Based Systems, Publish/Subscribe and Message Filtering Technologies, Christof Fetzer, Andrey Brito, Robert Fach, Zbigniew Jerzak, "StreamMine", Book chapter of Handbook of Research on Advanced Distributed Event-Based Systems, Publish/Subscribe and Message Filtering Technologies.
- A Large Scale Data Streaming System, V. Gulisano, R. Jimenez-Peris, M. Patiño-Martínez, P. Valduriez. A Large Scale Data Streaming System. 30th IEEE Int. Conf. on Distributed Systems (ICDCS), Genoa, Italy, 2010.

- Multithreading-Enabled Active Replication for Event Stream Processing Operators, Andrey Brito, Christof Fetzer, and Pascal Felber, Multithreading-Enabled Active Replication for Event Stream Processing Operators, The 28th Int. Symp. on Reliable Distributed Systems (SRDS). Niagara, USA. Sept. 27-30, 2009.
- “Soft State in Publish/Subscribe”, Zbigniew Jerzak and Christof Fetzer, “Soft State in Publish/Subscribe”, 3rd ACM International Conference on Distributed Event-Based Systems (DEBS) Nashville, TN, USA, July 6-9, 2009.

ANNEX B – DEVELOPING A COMMON VOCABULARY

SaaS-relevant paradigms and technologies

Software as a service (SaaS)

SaaS lets companies support key applications (CRM, ERP, etc.) by licensing worker "seats" on application software that's deployed by a provider. The provider maintains the hardware and software, upgrades facilities and fixes problems, thus eliminating the cost of deploying the application on a company data center. SaaS applications are accessed via a network connection -- normally over the Internet.

Because the SaaS model eliminates self-hosting of the software, it's most effective where self-hosting an application would be difficult -- for example, when there are few workers who need it (and the cost per worker would be very high) or because the skills needed to install and maintain the application aren't available in the local labor pool. Small to midsized businesses are very likely to find SaaS a better option than deploying their own applications.

Virtualization

If SaaS is an application model, virtualization is a technology. As it's popularly used today, virtualization is any technology that allows resources to be viewed as "logical" and "physical." Applications are run on logical resources, and these are assigned to an optimum physical resource based on a variety of cost, performance and availability policies.

Virtualization is a technology that allows a physical server to appear to applications as multiple "logical servers." Many companies have purchased new servers for each new application they deploy, in some cases resulting in thousands of underutilized servers that have to be individually powered, cooled and maintained. With virtualization, a much smaller number of servers housed in concentrated data centers can replace the inefficient distributed mass of servers, often reducing costs by more than 50%.

The challenge of virtualization lies in the issue of utilization. If a server truly has excess memory, disk and CPU resources available, then virtualization can provide for server consolidation and savings. If resources are limited, then virtualization will affect application performance. Some users have also found it difficult to manage virtual server farms, though tools from vendors like Cisco, HP, IBM, Microsoft, Oracle, and EMC/VMware are making that easier. Most report that virtualization does require more professional skill in house to manage. Companies that can benefit broadly from SaaS are less likely to benefit broadly from virtualization.

Cloud computing

Cloud computing is in many ways a harmonization of SaaS and virtualization into a much broader and more flexible model. As a model for IT infrastructure, cloud computing -- or the private cloud -- provides a way for enterprises to structure their data centers to efficiently use server, storage, and network resources. Cloud computing can extend virtualization across a wide area network (WAN) to build a single virtual cloud data center. Because a cloud data center can include the servers and storage of multiple physical data centers, it provides a larger pool of resources for applications to share than would be provided by simple server virtualization -- further improving cost efficiency.

The public cloud can offer companies a way to host applications off-site, for backup or as an overflow capacity resource in periods of peak demand. The public cloud can host SaaS applications more cost effectively, too. The cloud computing architecture provides a way to link the public resources with private cloud resources to create a hybrid cloud, enabling cloud computing to build a seamless application fabric across virtualized servers, SaaS services, and public/private cloud computing facilities.

SOA

An older and in many ways more confusing technology cuts across all three of these modern software innovations -- service-oriented architectures (SOA). SOA is a software design and development methodology that componentizes applications into modular services that are then assembled in various ways to promote customization to worker needs and reuse of common software elements. SOA facilitates cloud computing by making software easier to distribute in the cloud. It can be an alternative to virtualization in server-sharing, or it can use virtualization to improve performance and reliability. Finally, SOA makes everything a service and thus supports the SaaS model, not only for complete applications but also for components of applications. As SOA principles remake applications, SOA will transform all three technologies.

SOA may be transforming, but from the user's perspective, cloud computing is harmonizing. Many would argue that cloud computing is the unification of virtualization and SaaS, but it's more complicated than that. There will be applications of virtualization that are not "cloud computing" for years to come, and there will also be SaaS applications that do not use virtualization in application hosting. All three technologies have their place in managing capital and operations costs for IT. But it does seem likely that cloud computing will become the overarching framework in which most virtualization is applied and that the majority of SaaS providers will employ cloud computing to serve their users most economically and reliably.

(Reference http://searchcio.techtarget.com.au/news/2240020675/Cloud-computing-vs-virtualization-vs-SaaS-vs-SOA)

SaaS-Relevant Applications

SaaS stands for “Software as a Service”, and it essentially refers to software that is hosted on servers and is provided as a service. Some initial uses for SaaS included customer relationship management offerings, content management systems, video conferencing, and e-mail communication systems. SaaS applications are provided over the web, which means they can be accessed from any computer without any special software installed. In fact, many applications are designed to run through a standard web browser. When updates to a SaaS application need to be installed, they are simply installed on the server, which immediately ensures that all users are running the latest version. Unlike traditional software applications that require an upfront purchase, SaaS applications typically offer subscription-based pricing and are usually licensed on a per-user basis.

Many of the early adopters of SaaS were small businesses, primarily due to the low upfront costs and simplistic integration. Larger enterprises, however, have taken a somewhat more cautious approach to implementing SaaS solutions within their organization, particularly for mission-critical applications. Other solutions based on SaaS are provided directly for end users and often include a free access for the use of basic services.

In the following section are presented some relevant Saas vendors with relative applications:

IBM

IBM in last years has continued its push into SaaS. IBM clearly wants to be a major player into the domain of Cloud computing and cloud services. The most noticeable incremental development since 2007 is the move to operate in the application layer through the recent launch of Lotus applications in SaaS mode. IBM is positioning LotusLive as a SaaS collaboration suite, including features such e-mail, instant messaging, shared workspaces, conferencing, social networking and team-based document workflow management,

emphasizing underlying enterprise-class security. Collaboration applications have become a leading SaaS workload, as vendors such as Google, Microsoft and Zoho have all brought offerings to the market in recent years.

In particular IBM's overall cloud strategy breaks down into three major components:

- Smart business on the IBM cloud: this encompasses LotusLive and Smart BusinessDesktop, plus its Compute, Storage and Backup as a Service offerings. These are allcommercial cloud services.
- Smart business cloud: these are private cloud services that reside behind customers' firewalls, built by and potentially managed by IBM. This includes private cloud development and testing environments.
- Preintegrated workload-optimized solutions : the last element is bundled solutions(i.e., hardware, software and services) for certain workloads, such as testing, whichare offered in private cloud deployments.

Microsoft

Microsoft launched its first enterprise-focused SaaS product, Microsoft Dynamics CRM Online, in the first half of 2008. Dynamics CRM Online was joined in late 2008 by Microsoft Dynamics Online services, such as Payments and Commerce, which connect with Microsoft Dynamics solutions. Microsoft reports less customer adoption for this offering.. These Dynamics-branded offerings have been joined in the past year by Business Productivity Online Suite (BPOS), which consists of the following offerings in SaaS mode:

- SharePoint Online — collaboration tools
- Exchange Online — e-mail and calendaring
- Office Communications Online — instant messaging, voice and video
- Office Live Meeting — online meeting platform

Dynamics and BPOS make up two legs of Microsoft's "Software + Services" strategy, which is its positioning of offerings across the continuum from traditional/on-premises deployments through incremental hosting capabilities to fully multitenant-based commercial services. Microsoft has had significant market presence in the emerging ideas of SaaS and online software through its MSN services such as Hotmail, which has been used by hundreds of thousands of consumers for e-mail, and some business-oriented software, such as Messenger, Live Meeting and Office Live.

The other major component of Microsoft's Software + Services vision is Azure. Azure is Microsoft's re-envisioning of its core technology products for a world of cloud computing Azure consists of the following elements, all offered online and managed by Microsoft:

- Windows Azure — compute and storage, and management functionality
- SQL Azure — database and management functionality
- .NET — connectivity and access control

Oracle

Oracle has two real SaaS offerings: Oracle CRM On Demand and Oracle Behive On Demand. The most successful of these offerings is Oracle CRM On Demand. Oracle has made significant investments in this product in the past two years, to the point where it has become a leader on Gartner's "Magic Quadrant for Sales Force Automation."

- Oracle CRM On Demand provides the following features: Sales, Marketing, Service, Call Center, Analytics, Mobile, Integration.
- Oracle Behive On Demand is designed to help organizations secure communications and add collaboration into business processes; it integrates team workspaces,

calendars, instant messaging, and e-mail into a unified object model.

NetSuite

NetSuite was one of the earliest SaaS offerings on the market, and its primary purpose is to provide integrated business management software to midsize organizations. Their NetSuite application includes a comprehensive set of features, including customer relationship management (CRM), order fulfillment, inventory, accounting and finance, product assembly, ecommerce, website management, and employee productivity

SAP

SAP provides SaaS solution:

- SAP Business ByDesign is a fully integrated business management software solution designed for midsize companies or small businesses that want the benefits of large-scale business applications without the need for a large IT infrastructure. The software enables preconfigured process best practices for managing financials, customer relationships, human resources, projects, procurement, and the supply chain. SAP takes care of installation, maintenance, and upgrades
- SAP CRM on demand: It's a configurable, Web-based application that's available on a subscription basis over the Internet as a service. The on-demand solution provides your organization with immediate access to the sales, customer service, and marketing functionality

SalesForce.com

Initially launched in 1999, SalesForce.com has become a major player in the market for customer relationship management (CRM) software services. In conjunction with their force.com platform, the company also launched an AppExchange. The AppExchange serves as an online marketplace where developers can sell their plug-ins for use in other CRM applications.

Amazon EC2

Amazon launched a cloud computing initiative in 2006 called EC2, which allows developers to build scalable applications that run on their cloud. To use the cloud, developers only pay for the computing power that they actually use.

Other SaaS implementation consists in products aimed also for end users:

MobileMe

MobileMe is Apple's suite of online applications for iPhone, iPod touch, Mac and PC users. Customers can sync their e-mail, contacts, calendars, photos, etc. with the MobileMe online service using their computer or mobile device

Google Apps

Google apps is an SaaS platform that provides several services to users: Calendar, Gmail, Google Docs etc. These services are provided for free to anyone with a Google Account but also for the organizations through a customized domain

Microsoft Office Web Apps

Office Web Apps are free web-based version of Microsoft's Office productivity suite. It includes the web-based versions of Microsoft Word, Microsoft Excel, Microsoft PowerPoint, and Microsoft OneNote. The web apps allow users to access their documents directly from anywhere within a web browser as well as share files and collaborate with other users online.

Photoshop.com

Like many other photo sharing websites, Photoshop.com offers a limited amount of free photo storage, and several upgrade options for purchasing more space. To add more value

to their service, Adobe has also pushed many of their photo editing tools from their Photoshop application to their online service. Users can login to Photoshop.com and have access to many photo editing tools absolutely free.

SaaS growing

According to Gartner, Inc. the SaaS market will show consistent growth through 2013 when worldwide SaaS revenue will total over $14 billion for the enterprise application markets.

"The adoption of SaaS continues to grow and evolve within the enterprise application markets," said Sharon Mertz, research director at Gartner. "The composition of the worldwide SaaS landscape is evolving as vendors continue to extend regionally, increase penetration within existing accounts and 'greenfield' opportunities, and offer more-vertical-specific solutions as part of their service portfolio or through partners." "Adoption of the on-demand deployment model has continued to grow as on-demand vendors have extended their services through alliances, partner offerings, and more recently, by offering and promoting user application development through platform as a service (PaaS) capabilities," Ms. Mertz said. "Although usage and adoption is still evolving, deployment of SaaS still varies between the enterprise application markets and within specific market segments because of buyer demand and applicability of the solution."

The content, communications and collaboration (CCC) market and the customer relationship management (CRM) market continue to have the largest amount of SaaS revenue across market segments, with the CCC market generating $2.6 billion in 2009, up from $2.14 billion in 2008 and the CRM segment generating $2.3 billion in 2009, up from $1.9 billion in 2008. (see Table).

	2009	**2008**
Content, Communications and Collaboration (CCC)	2,573	2,143
Office Suites	68	56
Digital Content Creation (DCC)	62	44
Customer Relationship Management (CRM)	2,281	1,872
Enterprise Resource Planning (ERP)	1,239	1,176
Supply Chain Management (SCM)	826	710
Other Application Software	472	387
Total Enterprise Software	**7,521**	**6,388**

Worldwide Software Revenue for SaaS Delivery Within the Enterprise Application Software Markets
(Millions of Dollars) Source: Gartner (November 2009)

In the 2010, Gartner predicts worldwide SaaS sales will eclipse $8.5 billion, up 14.1 percent from 2009 sales of $7.5 billion. More telling, this rapid increase in both customers and SaaS vendors means that on-demand applications will make up a larger percentage of total enterprise software sales this year and for the foreseeable future.

SaaS apps, which accounted for a little more than 10 percent of the total enterprise software market last year, are expected to represent at least 16 percent of worldwide software sales by 2014.

For the report, Gartner defined SaaS as software that is owned, delivered and managed remotely by one or more providers.

Adoption varies between and within markets, and although use is expanding to a wider range of applications and solutions, the most widespread use is still characterized by horizontal applications with common processes, among distributed virtual workforce teams and within Web 2.0 initiatives," Mertz added.

Gartner estimates that 75 percent of the current SaaS delivery revenue could be considered a cloud service, and that could exceed 90 percent by 2014 as the SaaS model matures and converges with cloud services models.

“The market landscape for on-demand CRM continues to evolve as the availability and usage of SaaS solutions becomes more pervasive,” said Ms. Mertz. “The rapid adoption of SaaS and the marketplace success of salesforce.com have compelled vendors without an on-demand solution to either acquire smaller niche SaaS providers or develop the solution internally in response to increasing buyer demand.”

"The popularity of SaaS has increased significantly within the past five years and initial concerns about security, response time and service availability have diminished for many organizations as SaaS business and computing models have matured and adoption has become more widespread," Mertz said.

While pure SaaS vendors such as Salesforce and NetSuite are benefitting most from this early adoption wave, traditional on-premises vendors such as Oracle, Microsoft, SAP and IBM (NYSE: IBM) are rapidly overhauling their business models to accommodate this shift in customer expectations and preferences.

Gartner's report found this to be especially true in customer relationship management (CRM) space, where SaaS applications represented more than 24 percent of total CRM sales last year, and will account for at least 26 percent of CRM revenue in 2010.

"The market landscape for on-demand CRM continues to evolve and mature as the availability and use of SaaS solutions becomes more pervasive," Mertz added. "Greater market competition and increased focus by the mega-vendors reinforces the legitimacy of on-demand, mitigating initial objections about security and availability for many, as acceptance of SaaS as a viable model for enterprise computing services grows."