SEQUOIA PROJECT

"Socio-Economic Impact Assessment for Research Projects"

Contract n° 258346

WP2: SEQUOIA Self-Assessment Methodology development

Deliverable 2.1b - Results from Call5 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°: 2.1b

Due date: 15 March 2011

Delivery Date (Resubmission): 26 September 2011

Authors: T6, LSE, EK, ENG

Partners contributed: T6, LSE, EK, ENG

Made available to: Project partners and EC

Versioning		
Version	**Date**	**Name, organization**
1	10/01/2011	Antonella Passani (T6)
2	25/01/2011	Marie Debicki (T6)
3	1/02/2011	Mauro Navarra (EK)
4	15/02/2011	Shenja Van Der Graaf (LSE)
5	16/02/2011	Martino di Maggio (ENG)
6	20/02/2011	Marie Debicki (T6)
7	1/03/2011	Antonella Passani (T6)
8	10/03/2011	Antonella Passani (T6), Shenja Van Der Graaf (LSE), Mauro Navarra (EK), Martino di Maggio (ENG), Marie Debicki (T6)
9	14/03/2011	Antonella Passani (T6), Shenja Van Der Graaf (LSE), Mauro Navarra (EK), Martino di Maggio (ENG), Marie Debicki (T6)
10	28/03/2011	Antonella Passani (T6), Shenja Van Der Graaf (LSE), Mauro Navarra (EK), Martino di Maggio (ENG), Marie Debicki (T6)
11	20/07/2011	Paolo Dini, Shenja Van Der Graaf (LSE)
12	20/09/2011	Shenja Van Der Graaf (LSE)

Quality check: Louise Newton-Clare - LSE

Internal Reviewers: Francesco Bellini - EK

TABLE OF CONTENTS

EXECUTIVE SUMMARY

This deliverable presents a complete description of the information available to date on Call 5 Internet of Services (IoS) projects and represents an extension of deliverable D2.1a. Deliverable D2.1a "Results from Call 1 projects' document inventory" was submitted to the European Commission (EC) in December 2010 and presents a description of the documents available for Call1 projects.

In this deliverable all the 23 projects funded under Call 5 were taken into consideration with the exception of the support actions. The SEQUOIA team analysed approximately 40 deliverables and 18 publications from 12 different projects. The remaining 11 projects have no deliverables yet as Call 5 projects started recently. In addition to document collection, further information was gathered through 5 focus group sessions, which provided relevant insights and information on the group of projects under analysis (please see D2.2b for the results of the focus group).

The last part of this deliverable provides a short description of all the documents analysed and the annex synthesises goals and documents available for each of the 23 Call 5 projects.

INTRODUCTION

This deliverable aims to present a complete description of the information available to date on Call 5 Internet of Services (IoS) projects. It represents an extension of deliverable D2.1a. Deliverable D2.1a "Results from Call 1 projects' document inventory" focusing on Call 1 projects was submitted to the European Commission (EC) in December 2010. The present deliverable (D2.1b) was prepared in response to the considerations made during the 1st project review, in agreement with the EC representatives, to work also with the Call 5 projects to develop the SEQUOIA methodology.

Before starting a description of the information gathered, it is important to note that Call 5 projects have all started only a few months ago (between June 2010 and November 2010). This implies that the documents available about Call 5 projects are considerably fewer than those available for Call 1 projects. A large number of projects have not yet published their first deliverable, so the analysis presented here focuses mainly on their factsheets and their websites (of course, deliverables and articles have been analysed when available). The 5 Focus Group sessions took place between 14th and 23rd of February and involved 27 projects, of which 16 were from Call 5. These sessions provided a good opportunity to collect more information about the projects, their objectives and beneficiaries/target users. The outputs of the focus group are presented in D2.2b.

In deliverable D2.1a, the consortium decided to cluster the projects under analysis in order to develop a deeper understanding from the point of view of the technology developed and potential beneficiaries. Consequently, in this deliverable we will list and present the documents gathered and used for the clustering exercise undertaken and for further developing the second questionnaire.

D2.1a provides more information on the location of the present deliverable in the general map of the SEQUOIA socio-economic impact assessment development process.

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

CHAPTER 1. CALL 5 PROJECTS: AN OVERVIEW

In this deliverable the SEQUOIA partnership took into consideration 23 projects, which represent the total number of projects funded under Call 5, with the exception of the support actions. According to the organisation of the call in sub-areas developed by the EC (Fig. 1) we can see the projects grouped in the following 3 areas: general Internet of Services, which is divided between Service Front-ends (5 projects) and Service Platforms (4 projects); Cloud Computing, which is divided between Platform as a Service (3 projects) and Infrastructure as a Service (5 projects); and finally Advanced Software Engineering (6 projects).

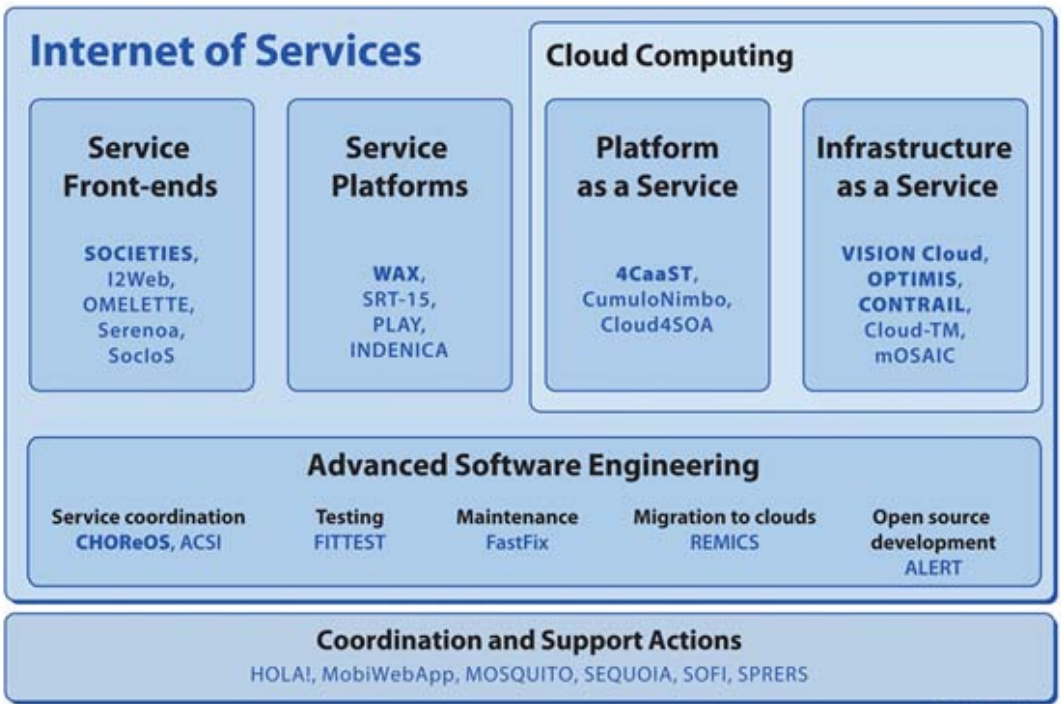

Fig. 1 - Call 5 Projects

The list of projects under analysis follows:

4CAAST
CLOUD-TM
ACSI
CONTRAIL
I2WEB
OMELETTE
SERENOA
SRT-15
ALERT
CUMULONIMBO
INDENICA
OPTIMIS
SOCIETIES
VISION CLOUD
CHOREOS
FASTFIX
PLAY
WEBINOS (was WAX before)
CLOUD4SOA
FITTEST
MOSAIC
REMICS
SocioS

As indicated above, Call 5 projects have started only some months ago and many have not published their first deliverable yet (with the exception of management or communication-related deliverables). Considering all the projects, there are approximately 40 deliverables and 18 publications available from 12 different projects. The remaining 11 projects have no deliverables yet. Moreover some of these deliverables are not available on the website. We analysed all the available deliverables and the main publications. In addition, for each project the website and the factsheet were taken in consideration.

As for Call 1 projects, the consortium paid particular attention to the concrete solutions that the projects aim to develop, including the use cases and prototypes foreseen. In fact, to have concrete realisations and case studies would make the development of variables for the impact assessment easier. The focus groups organised were very useful in this regard, as for each project the participants were asked to present their use cases. More than 26 use cases are foreseen by the projects in very different sectors. Some projects are still deciding or have to better specify the use cases that will be developed.

Table 1 includes a short description of each project and its main outputs. As for Call 1 projects, we used this preliminary description internally in order to build a common understanding about the projects under analysis. The consortium used the approach developed for Call 1 projects also for Call 5 projects, as it was useful in the analysis of the projects (with some small modifications to improve its efficiency).

Table 1 - Call 5 projects' brief description and outputs identification

Project	Brief description
4CAAST	4CaaSt (Building the PaaS Cloud of the Future) embeds features that ease programming of rich applications and enable the creation of a business ecosystem where applications from different providers can be tailored to different users, mashed up and traded together. It aims to create a PaaS Cloud Platform which supports the optimized and elastic hosting of Internet-scale multi-tier applications. **Outputs**: A PaaS Cloud Platform that will provide: • Virtualization of processing and networking end-to-end; • Higher level of abstraction through advanced service life-cycle management based on blueprint; • RT scaling and adaptation based on convergent QoS management; • Support of service platform technologies as immigrants on the cloud; • Components enabling business ecosystems through marketplace functions; • Possibility to combine with apps/services from other app providers and the Cloud Provider; • Possibility for the end users develop their own apps (DIY); • A complete set of built-in programming libraries and common facilities beyond what is offered by state-of-the-art PaaS Clouds, easing development of killer applications; • The necessary tools to monitor execution and manage lifecycle of applications.
CLOUD-TM	*Cloud-TM: A novel programming paradigm for Cloud computing* develops a Self-Optimizing Distributed Transactional Memory middleware that will spare programmers the burden of coding for distribution, persistence and fault-tolerance, thus facilitating the development and administration of cloud applications.

	Outputs: A Cloud platform that will consist of three components: • i) Distributed Software Transactional Memory: DSTMs represent the confluence be-tween Transactional Memories (TMs) and Distributed Shared Memories (DSMs). DSTMs enrich the TM model to breach the boundaries of a single machine and transparently leverage the resources of commodity, shared-nothing clusters, thus achieving higher scalability and dependability levels. In DSTMs, the transaction abstraction serves not only to simplify parallel programming. It is also a natural means to batch consistency actions in a single synchronization phase taking place at commit time, amortizing communication overheads across a (possibly) large number of memory accesses. This allows to overcome the performance limits of strongly consistent DSMs, without forcing programmers to em-brace subtle relaxed consistency models as in weakly consistent DSMs. • ii) Distributed Storage System: The DSS provides durability for both the user-level application data and long-lived middleware data. This component will implement the persistence mechanisms and expose an API capable of abstracting several concerns of persistence in the cloud, such as scalability and data partitioning. • iii) Autonomic Manager: The AM component will implement a feedback control loop, collecting and pro-cessing comprehensive workload information from each layer of the Cloud-TM architecture. The AM will be in charge of determining the adaptation criteria according to either a) on-line computed performance forecasts, or b) off-line determined adaptation policies (e.g. based on static thresholds). • The project will also define a friendly programming model for large-scale distributed applications that integrates the familiar notion of atomic transaction as a first-class programming language construct. This would spare programmers from the burden of implementing low level, error-prone mechanisms (e.g. distribution, persistence and fault-tolerance) and attaining major reductions of the time and cost for the development process.
ACSI	The project will develop scientific advances, techniques, and tools to simplify the design and deployment of infrastructure to support service collaborations, the ability of services to join such collaborations, and the evolution of such collaborations as the marketplace and competitive landscape change. Artifact-Centric Service Interoperation (ACSI) is based on two fundamental constructs: the interoperation hub and dynamic artifacts. The project will therefore contribute to the reduction of the effort and lead-time of designing, deploying, maintaining, and joining environments that support service collaborations. **Outputs**: • ACSI Hub System, which will include a comprehensive Design Workbench for designing, deploying, maintaining, and evolving interoperation hubs. An interoperation hub serves as a virtual rendezvous for multiple services that work together towards a common goal. The project research will develop a principled, easy-to-use framework for creating, deploying, maintaining, and joining ACSI interoperation hubs in essentially any application domain. Similar to EasyChair or SalesForce.com, an ACSI interoperation hub will serve as the anchor for a collaborative IT environment that supports large numbers of service collaborations that operate independently, but focus on essentially common goals. • The interoperation hubs used here will be structured around dynamic artifacts. These provide an holistic marriage of data and process, both treated as first-class citizens, as the basic building block for modeling,

	specifying, and implementing services and business processes. Artifacts can give an end-to-end view of how key conceptual business entities evolve as they move through the business operations. As a result, artifacts can substantially simplify the management of "hand-off" of data and processing between services and organizations. A key pillar of the ACSI research is to generalize the advantages of artifacts to the context of interoperation hubs and service collaborations.
CONTRAIL	CONTRAIL (Open Computing Infrastructures for Elastic Services) will vertically integrate an opensource distributed operating system for autonomous resource management in Infrastructure-as-a-Service (IaaS) environments, and high level services and runtime environments as foundations for Platform-as-a-Service (PaaS). It will help to fill the gap in today's business Clouds by providing reliable, transparent and secure Cloud infrastructures. The project will design, implement, evaluate and promote an open source system in which resources that belong to different operators are integrated into a single homogeneous Federated Cloud that users can access seamlessly. **Outputs**: A complete Cloud platform which integrates a full Infrastructure-as-a-Service and Platform-as-a-Service offer, based on the open source XtreemOS system. In particular the project will provide: • a collection of infrastructure services offering network, computation and storage as a service; • services to federate IaaS Clouds; • a collection of PaaS services to support typical Cloud applications; • a collection of run-time environments providing elasticity, scalability and performance dependability to selected classes of applications; and • a collection of applications and use cases from the domains of e-business, e-science telecommunication and media using and demonstrating the Contrail system.
I2WEB	I2WEB provides tools to develop inclusive Future Internet services that will overcome the risk of a Future Internet Community further isolating excluded groups. The main objective of the project is to enable the Future Internet to be very extensively used by people with disabilities and the elderly. I2Web particularly responds to immediate challenges of the Future Internet: ubiquitous and mobile Web, media convergence and user-generated content, in combination with cloud computing, Web 2.0 developments, Social Networking, User-Centred Design and Inclusive Design principles. **Outputs**: • i) An open, scalable, dependable service platform architecture and online space, enabling automatic service description, discovery, composition, and negotiation with a multiplicity of reusable inclusive services, which may be mobile or nomadic, multi-device, multimodal and multi-context. • ii) User models based upon existing accessibility standards combined with an analysis of user requirements for people with special needs and older people in relation to ubiquitous Web 2.0 applications, in which multimodality and delivery context are key components. • iii) Extension of existing device models from the Mobile Web arena, to cope with the needs of other devices, ranging from standard desktops to consumer electronics devices. • iv) Open information models and generic application abstractions, that can deal with information aggregation, cloud computing applications, Semantic

	Web and mobile/ubiquitous Web 2.0 systems.
OMELETTE	OMELETTE (Open Mashup Enterprise service platform for LinkEd data in The TElco domain) will create a sound model of mashups that follows the REST architectural style (also supported by standard widget technology), as well as a standard specification of a mashup-containing platform that may guarantee portability and interoperability among different vendors and versions. These concepts will be based on a solid theoretical model of mashup foundations and the specific requirements gained from the telco domain. **Outputs**: • An open interoperable service platform for converged mashup services. In particular the project will define: o The specification of a mashup reference architecture; o The definition of a description language for mashup services and properties: the language will cover both the sources to be mashed up and the final mashups; o The specification of a standard proposal that will aim at interoperability and portability of mashups and mashup sources (e.g., widgets) among multiple platforms. • Research on innovative approaches on converged service engineering integrating telecommunication capabilities in multimodal mashup-based interaction. • Research on the relationship of widgets and mashup technologies, taking into account governance and security issue.
SERENOA	will develop an open platform that enables the creation of context-sensitive SFE. These will provide a user interface (UI) that exhibits some capability to be aware of the context and to react to changes of this context in a continuous way. As a result such a UI will be adapted to a person's devices, tasks, preferences, and abilities, thus improving people's satisfaction and performance compared to traditional SFEs based on manually designed UIs. **Outputs**: Concepts, languages, (intelligent) runtimes and tools are needed to support multi-dimensional context-aware adaptation of Service Front-Ends (SFEs). These artefacts will enable SFE engineers to concentrate on the functionality rather than on the implementation details concerning the adaptation to the multiple dimensions of the context of use.
SRT-15	SRT-15 (Subscription Routing Technology for 2015) will allow for dependable and scalable cloud-based processing of data coming to and from a variety of heterogeneous enterprise services spread across multiple distributed locations. In order to be able to embrace the change in the enterprise information processing landscape, SRT-15 relies on technologies that support rapid change: cloud computing, content-based routing, and complex event processing. **Outputs**: A scalable platform for connecting business applications and services by enabling the discovery and integration of dynamic enterprise services on the Internet. This platform will be the following features: • Scales Across Public and Private Clouds, • Uses Content to Route Information, • Provides Flexible and Timely Service Interaction, • Processes Data in Distributed and Parallel Fashion,

	• Is Dependable and Protects Data Privacy.
ALERT	ALERT will develop methods and tools that improve FLOSS (Free/Libre Open Source Software) coordination by maintaining awareness of community activities through real-time, personalized, context-aware notification. It will create an active collaboration platform, i.e. a virtual actor that would interact with other developers, process and recognize various kinds of interactions, suggest actions on the basis of these and remember and bring past interactions to the developers' attention, thus enabling developers to work better together. **Outputs**: • Set of ontologies that model the conceptual dependencies between community, content and interactions. • Conceptual model of an (re)active collaboration platform. • Model for declarative specification of alerts. • Intelligent, pattern-based engine that can correlate, aggregate, and filter events in order to provide triggers for relevant notifications. • Module for integration of information from structured and unstructured sources. • Interaction highway set of services (e.g. search, recommendation, visualisation, etc.). • Module for capturing explicitly or implicitly the interest/expertise of the developer. • Module for capturing relevant information for the task at hand.
CUMULO-NIMBO	CumuloNimbo will deliver a new generation of cloud computing PaaS that will exhibit high scalability while guaranteeing full consistency, and provide consistency, availability, and simpler programming abstractions, such as transactions. It will follow a radically different approach from current state of the art approaches in cluster and cloud environments. Cluster environments have been based on replication approaches that provide strong consistency but at the cost of limited scalability. For this reason cloud environments rely exclusively on data partitioning and therefore do not provide data consistency across data partitions. **Outputs**: CumuloNimbo Platform: • i) will be highly elastic with each subsystem scaling independently to optimize the use of resources and minimize the cost incurred to process the actual incoming load. • ii) will also be highly available with self-healing sub-systems that automatically repair themselves in the advent of failures without disrupting service provisioning and without introducing data inconsistencies during failures and recovery. • iii) will provide full transactional guarantees, that is, without requiring any data partitioning.
INDENICA	INDENICA will support the development of virtual domain-specific platforms and provide Virtual Service Platforms as a novel domain-specific approach to create service-based applications. In particular, its approach will ease the development of domain-specific, service-based applications by resolving the quality, interface, and technology fragmentation that can be observed in today's service platforms. It aims to provide methods, architectures, components, tools and assets for reuse-based creation of the adapted platforms. Effectively, INDENICA will research on a product line for the development of domain-specific platforms that

	are interoperable by design. Members of the product line (i.e., platforms) can be modified after deployment, thus providing capabilities to make decisions that are currently made at design time at runtime. **Outputs**: • i) Methods, architectures, components, tools and assets for reuse-based creation that support the development of context and domain-specific platforms along with the services. • ii) A common basis for families of platforms which integrate system management and interoperability capabilities right from the start. • iii) A reusable infrastructure for platform development and support the automatic deployment and the monitoring, governance, and adaptation of services in a Virtual Service Platform.
OPTIMIS	The motivation for OPTIMIS is the vision that hybrid clouds will become commonplace, realized by private clouds interacting with a rich ecosystem of public and other cloud providers. OPTIMIS is aimed at enabling organizations to automatically externalize services and applications to trustworthy and auditable cloud providers in the hybrid model. Consequently, OPTIMIS believes its activities will support and facilitate an ecosystem of providers and consumers that will benefit from the optimal operation of services and infrastructures. The optimization covers the full lifecycle of services and their interactions. **Outputs**: i) Open specification and a toolkit that supports the construction of next generation cloud architectures. ii) The OPTIMIS toolkit will provide a set of independent components that can be adopted, either in full or in part, by Infrastructure Providers that offer the capacity required by services, and by Service Providers that use this capacity to deliver services. In particular the toolkit will be composed by the following components: • Monitoring Manager: gathers all monitoring data from the different monitoring collectors and store it in the Monitoring Database. • SLA/SLO Monitoring/Evaluation Tool: communicates with the Monitoring Manager in order to receive the monitoring data of the SLOs and SLAs during the execution. For the effective evaluation of the QoS requested, the component must retrieve the agreed SLA/SLO from the respective repository. • Trust Monitoring/Assessment Tool: takes as input Metering Data from the Monitoring Database in order to assess the trust and the reputation of the providers through the previous executions. • Risk Monitoring/Assessment Tool: takes as input Metering Data from the Monitoring Database in order to assess the risk and the expected integrity of a provider. • Cost Monitoring/Assessment Tool: has as input the SLA/SLO from the respective Tool/Service as well as metering information from the Monitoring Manager, it assesses the cost of the execution. The outcome of the assessment is forwarded to the Monitoring Manager in order to be stored to the Monitoring Database for future use. • High level application/services Monitoring Collector: collects high level monitoring information from the applications/services that are being executed within the virtual environment. That Service Runtime metering data are being pushed to the Monitoring Manager for aggregation and storage. • Virtual IT-infrastructure Monitoring Collector: collects virtual IT-infrastructure monitoring data (e.g. VM parameters). • Physical IT-infrastructure Monitoring Collector: collects physical IT-infrastructure monitoring data (e.g. network link supervision, computing

	resources data, storage data, etc.). • Energy efficiency Monitoring Collector: collects energy efficiency-related monitoring data (e.g. energy consumption). • Data Mining System: receives/collects monitoring data from the Monitoring Database and performs the necessary actions to store the data in an appropriate format and level of granularity, in the Data Mining Database (selection of relevant data, classification, aggregation, etc.). Furthermore, the Data Mining System is in charge of performing prediction of future. • Monitoring Database: is a database storing raw monitoring data (aggregated to certain extent, but not modified). • Data Mining Database: it is a database storing monitoring information previously processed by the Data Mining System. • Monitoring Infrastructure Toolset GUIs: Graphical User Interfaces which allow different kind of users to access the monitoring information, both "real-time" evolution and historical information.
SOCIETIES	The vision of SOCIETIES is to develop a complete, integrated Community Smart Space (CSS), which extends pervasive systems beyond the individual to dynamic communities of users. CSSs will embrace on-line community services, such as Social Networking in order to offer new and powerful ways of working, communicating and socialising. CSSs provide the following benefits: Shared resources with other community members in a seamless unobtrusive manner. CSSs will provide the means by which users of pervasive systems can come together to co-operate, socialise and share their communal resources. **Outputs**: • The resource sharing facilities will address cases of multiple requests for the same resource at the same time. CSSs will support sharing of resources owned by community members as well as by publicly available resources. • The support of multiple techniques for the discovery of relationships and behaviours within communities. • The orchestration of multiple communities to which an owner belongs, maintaining a registry of super- and sub- communities in community hierarchies along with policies on information disclosure and service access to members of other related communities. • The proactive exchange of information on the situation, interests and resources of community members. • The intelligent use of information learnt from monitoring communities and the exchanges between their members. This can be used to support the proactive discovery, configuration, adaptation, control and sharing of services and resources. • Intelligent conflict resolution among the members of a community based on mediation and negotiation. The support of ad-hoc communication at both intra- and inter- community levels, across heterogeneous networks and device platforms.
VISION CLOUD	The goal of VISION Cloud is to introduce a powerful ICT infrastructure for reliable and effective delivery of data-intensive storage services, facilitating the convergence of ICT, media and telecommunications. This infrastructure will support the setup and deployment of data and storage services on demand, at competitive costs, across disparate administrative domains, while providing QoS and security guarantees. **Outputs**: An architecture and a reference implementation of a cloud-based infrastructure,

	built on open standards and new technologies, to provide a scalable, flexible and dependable framework for optimised delivery of data-intensive storage services. In particular the cloud platform will be based on the following concepts: • Raising the Abstraction Level of Storage: encapsulating storage into objects with user-defined and system-defined attributes (a.k.a. "meta-data"). This metadata supports effective access, management and manipulation of the storage and stored data and derivation of new information. With that, we enable high scalability and simplification of all storage and data functions. • Data Mobility and Federation without Boundaries: substantially extending the limited data migration capabilities available in contemporary infrastructures to migrate and federate data across geographically distributed administrative domains, ensuring comprehensive and transparent data interoperability and overcoming the problem of data lock-in. • Computational Storage: providing an architected solution for bringing computing to storage, enabling secure execution of computational tasks near their data, applicable both to user-driven computational tasks as well as autonomous data derivation and transformation. • Content Centric Storage: facilitating access to data by content and its relationships, rather than details of the underlying storage containers, further supported by mechanisms to define domain-specific c storage optimisations, and altogether leading to highly simplified and efficient storage access. • Advanced Capabilities for Cloud-Based Storage: combining the base concepts of cloud computing with additional capabilities of storage virtualisation, advancing the state of the art by creating a platform where data-intensive services can be provided securely, at the desired level of service, and at competitive costs.
CHOREOS	The CHOReOS objective is to sustain decentralized service choreographies in the Future Internet. It revisits the concept of choreography-centric service-oriented systems to introduce a dynamic development process and associated methods, tools, and middleware – referred to as CHOReOS Integrated Development and Runtime Environment (IDRE) – for the software systems that implement and coordinate the services in the Ultra Large Scale Future Internet. **Outputs**: A framework for scalable choreography development, specifically enabling the development of decentralized, compositional, adaptable, QoS-aware choreographies by domain-experts. In particular: • Reference software architecture for the target choreography-based software systems, introducing abstractions for services, interaction protocols, and coordination (choreography patterns); • A model-driven process and tools enabling rigorous and systematized development, as well as requirements specifications and choreography modelling by domain experts (as opposed to IT professionals); • Service middleware support, effectively enabling the deployment of adaptable, QoS-aware choreographies in the ULS Future Internet; • Definition of concepts of governance of services and choreographies and development of mechanisms for Verification and Validation (V&V) of adaptable, QoS-aware choreographies in the context of the ULS Future Internet.

FASTFIX	The overall goal of the FastFix project is to provide software developers with a maintenance environment that combines time efficiency with low cost and high precision. FastFix will develop a platform and a set of tools that will continuously monitor customer environments, while collecting information on application execution and user interaction. The overall objective is to identify symptoms of execution errors, performance degradation, or changes in user behaviour. By using correlation techniques, the platform will also support failure replication in order to identify incorrect execution patterns and, in particular cases, automatically generate and deploy remedial patches. **Outputs**: A platform and a set of open source tools to on-line monitoring of execution environments, gathering semantic information on application and user behaviour. This information is sent in real time to a support centre, taking special care on privacy and security issues. Using event correlation techniques, FastFix identifies failure symptoms, performance degradation or changes in user behaviour and allow for failure replication, patch generation and patch deployment, resulting in a self-healing software application. In particular will be developed: • tools to gather context information on user and application, • a run-time with minimum impact on application performance, • a secure method to send this information to a centralized fault analysis platform, • a tool to detect software failures, undesirable execution trends and performance degradation, • a platform to replicate failure conditions within a virtual machine, • a tool to generate change strategies and necessary patches.
PLAY	The PLAY project will provide an open highly distributed Platform for event-driven interaction between services that scales at the Internet level based on a federated architecture able to address the multiplicity and the heterogeneity of service networks and address Quality of Service (QoS) requirements (such as dependability, security and scalability) in very large scale Internet based networks. The project will develop and validate an elastic and reliable architecture for dynamic and complex, event-driven interaction in large highly distributed and heterogeneous service systems. **Outputs**: A platform which collects information in real time (events) from many, heterogeneous, distributed event sources, process those events in a complex way, and, after discovering something relevant, forward such a situation (combination of events) to the parties interested in that issue. In particular the project will provide: • Federated middleware layer: a peer-to-peer overlay network combined with a publish/subscribe mechanism, that has the task of collecting events coming from the heterogeneous and distributed services; • Distributed complex event processor: an elastic, distributed computing cloud based engine for complex processing/ combination of events coming from different services in order to detect interesting situations that a service should react on; • Situational-aware business adapter: a recommender engine for proposing adaptation and changes in running business processes and services in a non-predefined (ad-hoc) way, by ensuring the consistency of the system.
WEBINOS	Promoting a “single service for every device” vision, Webinos will move the

	existing baseline of web development from installed applications to services, running consistently across a wide range of connected devices, ensuring that the technologies for describing, negotiating, securing, utilizing device functionalities and adapting to context are fit for purpose. The project will in particular deliver an OS platform for web applications across mobile, PC, home media (TV) and in-car devices. **Outputs**: An Open Source Platform and software components for the Future Internet in the form of web runtime extensions, and complementary infrastructure components, to enable web applications and services to be used and shared consistently and securely over a broad spectrum of converged and connected devices, including mobile, PC, home media (TV) and in-car units. In particular webinos: • bases on the achievements of the Web community and extends an open source web runtime environment. • is a federated web runtime that offers a common set of APIs to allow apps easy access to cross-user cross-service cross-device functionality in an open yet secure manner. • aims at easy programming of applications by offering a single virtual device that can consist of all devices owned by a user. This device inherently represents the changing circumstances of context and devices. • creates open specifications and open source reference implementations that not only show the feasibility of the specifications but also simplify their adaptation by the industry.
CLOUD4SOA	CLOUD4SOA aims at resolving the semantic interoperability issues that exist in current Clouds infrastructures and on introducing a user-centric approach for applications which are built upon and deployed using Cloud resources. The project will identify and analyse the semantic interoperability problems within the Cloud by carrying out a comprehensive analysis on both IaaS and PaaS layers, while including stakeholder experience, requirements and feedback. The Cloud Semantic Interoperability Framework will be the result of this exercise, acting as a roadmap for creating such user-oriented platforms. It will also implement, deploy and evaluate the usability, applicability and adaptability of the Cloud4SOA Reference Architecture by developing a proof-of-concept platform which will be applied in different scenarios and setups across 3 European industrial and business showcases. **Outputs**: • Cloud Semantic Interoperability Framework that represents a solution for the semantic interoperability problems within the Cloud. • An open, generic architecture for a semantically interoperable Cloud, capitalising on Service Oriented Architecture, lightweight semantics and user-centric design and development principles. • Implementation and deployment of the Cloud4SOA Reference Architecture by developing a proof of concept platform which will be applied in different scenarios and setups across 3 European industrial and business showcases.
FITTEST	The Future Internet will be a complex interconnection of services, applications, content and media, on which our society will become increasingly dependent for critical activities such as social services, learning, finance, business, as well as entertainment. This presents challenging problems during testing; challenges that simply cannot be avoided since testing is a vital part of the quality assurance process. FITTEST aims to attack the problems of testing the Future

	Internet with Search Based Testing. **Outputs**: Search Based Software Testing (SBST) based on search algorithms that are related to an objective function, that measures the degree to which the current solution (e.g. test case) achieves the testing goal (e.g. coverage, concurrency problems, etc.). SBST represents the unifying conceptual framework for the various testing techniques that will be adapted for FI testing in the project. Implementations of such techniques will be integrated into a unified environment for continuous, automated testing of FI applications.
MOSAIC	mOSAIC aims to develop an open-source platform that enables applications to negotiate Cloud services as requested by their users. Using mOSAIC solutions, Cloud-based application developers, maintainers and users will be able to: • specify the service requirements in terms of a Cloud ontology and communicate them to the platform via the innovative API; • postpone their decision on the procurement of Cloud services until run-time; • find best-fitting Cloud services to their actual needs and efficiently outsource computations. The platform will implement a multi-agent brokering mechanism that will search for services matching the applications' request, and possibly compose the requested service if no direct hit is found. It will also facilitate competition between Cloud providers, who, in return, will be able to reach customers they could not reach before. **Outputs**: • The mOSAIC API that will be a Cloud-based language - and platform - independent application programming interface that extends the existing language - or platform-dependent API capabilities with composite features based on usage patterns. • The mOSAIC framework that will include the complete set of scientific and technical solutions provided by mOSAIC. In particular, behind the solutions related to the API definition and instantiation, the framework will include semantic-oriented proposals like a Cloud ontology, which will allow for the semantic representation of Cloud resources and a direct application of semantically driven information processing. • The open-source platform that will be a proof-of concept prototype ready to be tested, exploited or extended by its users. It will include instances of the APIs for several programming languages and application tools (e.g., a workflow editor). Its semantic engine that, based on the Cloud ontology, expresses the application's needs for Cloud resources in terms of SLAs and QoS requirements.
REMICS	REMICS proposes a leap progress in legacy systems migration to Service Clouds by providing a model driven methodology and tools, which significantly improve the baseline ADM concept. This will be achieved through driving the standardization work in OMG including PIM4Cloud specification, and support and increase impact further by providing project results under open source licences. **Outputs**: • Extension of the KDM metamodel. On Platform Independent Model level, the components and services are defined using SoaML (SOA Modelling Language). The REMICS project will extend this language to address the

	specific architectural patterns and model driven methods for architecture migration, and to cover specificities of service clouds development paradigm. In particular, the PIM4 Cloud Computing, model driven Service Interoperability and Models@Runtime extensions are intended to support the REMICS methodology for service cloud architecture modelling. • EMF Eclipse based versions of the models representing the standards being worked on and provided in REMICS (Open source Metamodels and Models with an emphasis on Open Models for standards based on Open Interfaces)
SocioS	SocIoS views the democratization of the service lifecycle management as a major step towards the Future Internet of Services. The focal point of the SocIoS concept is the provision of the necessary tools to common Internet users that allow them to compose, provide and consume application services over abstract resources on the Web. The project will build qualitative, functional and usable business applications exploiting the User Created Content (UCC) and the Social Graph of users in Social Networks. **Outputs**: • A SOA infrastructure that will act as a virtualization layer on top of social networking containers; • A toolset for supporting the business extensions of the developed applications; mainly through SLA support; • An API that will grant end users with single point access to the underlying functionality, including a set of core wrapper services; • A usable user interface that will simplify the use of the SocIoS API for easier application development and deployment from the home user; • A toolset for third-party services support that perhaps can “blend” with those developed and deployed using SocIoS for more effectively delivering applications that exploit the UCC and social graph.

In terms of the life cycle of the projects, we have said that most of the projects started between June 2010 and November 2010, the majority will last until the first half of 2013. The figure below shows the duration of the projects. More information about the projects can be found in Annex A.

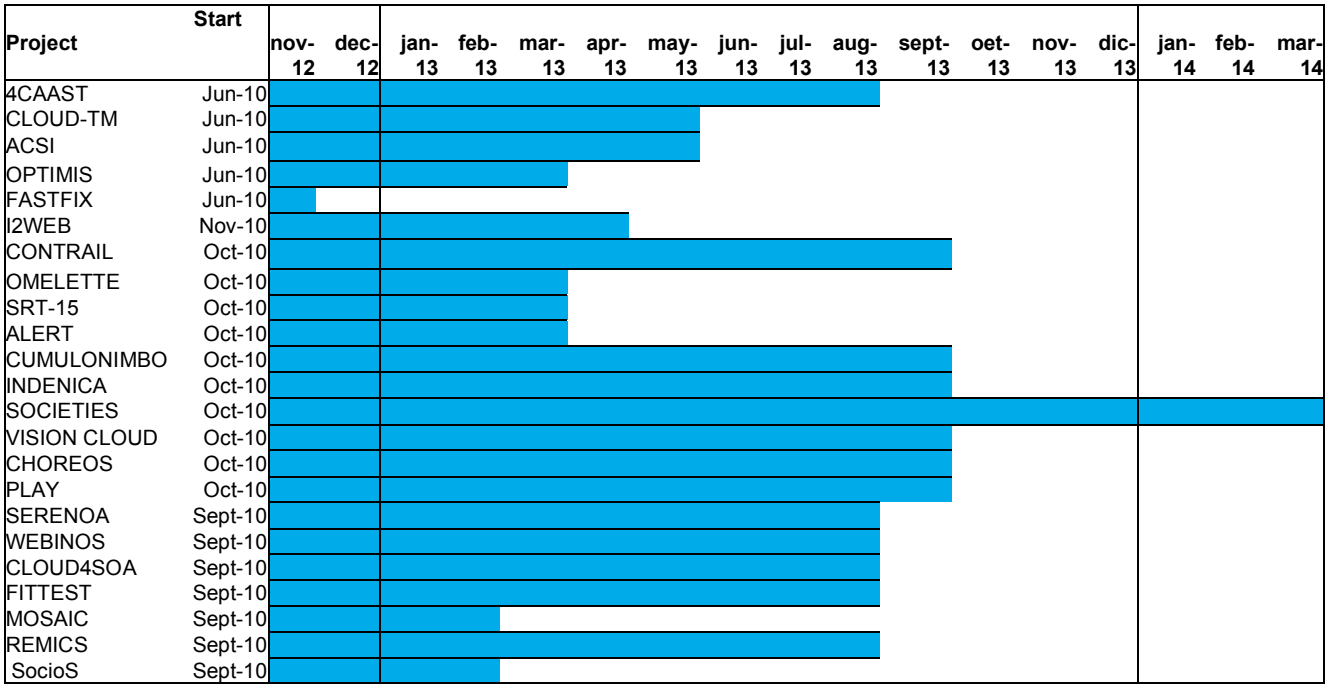

Fig. 2 – Duration of Call 5 Projects

CHAPTER 2. CLUSTERING THE PROJECTS UNDER ANALYSIS

In the previous chapter, each project under analysis was considered individually and its main outputs were briefly described. In this chapter, we aim to discover a certain 'likeness' among the projects according to their users/beneficiaries and in technological terms. In particular, this chapter seeks to provide the relational links between the 23 projects under investigation. This is for the SEQUOIA consortium's internal usage. The reason for this exercise is to develop an understanding of the nature of the projects facilitating our understanding of the centrality of the various beneficiary categories and the relevance of different technological solutions underpinning Call 5 Projects. It is important to emphasise that this exercise does not impact directly the SEQUOIA methodology developed for this project; 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.

For Call 1 projects a relational exercise was employed building on the factsheets. A similar approach was used for the Call 5 Projects so leaving room for comparison between Call 1 and Call 5 if needed (for example, at a later stage in the SEQUOIA project for the definition of variables). Again, all factsheets were entered in a content analysis package used for Call 1 projects, 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 4CAAST factsheet including content analysis (yellow) and coding schema.

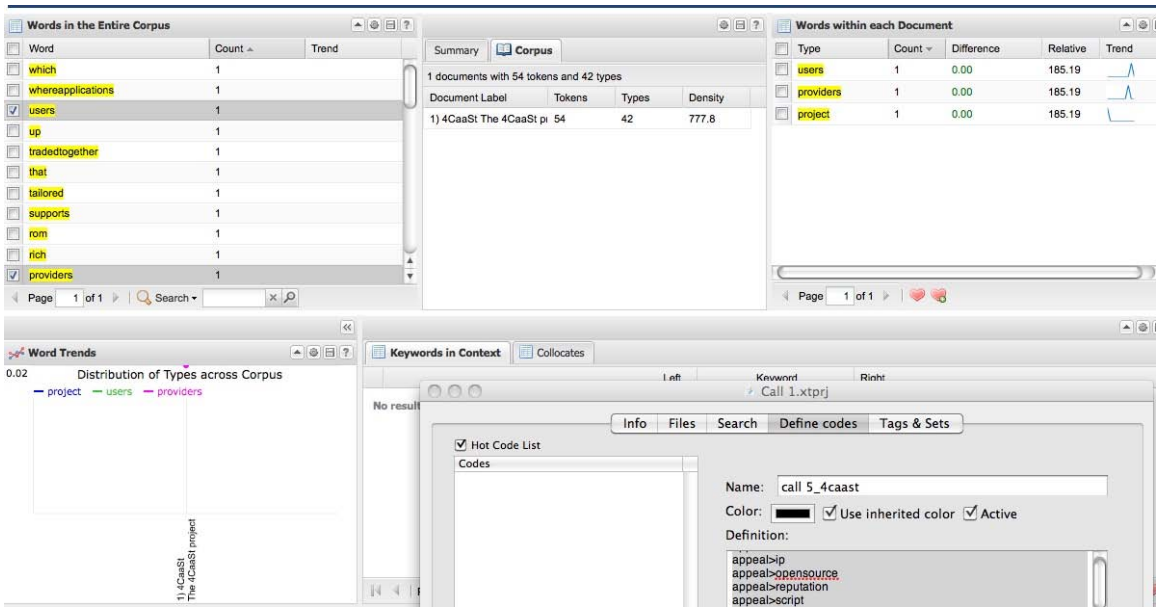

Figure 3 – Screenshot thematic analysis; example 4CAAST project

This resulted in the organisation of the qualitative data into themes or categories associated with users or beneficiaries of the projects fitting the four thematic categories that were developed earlier 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. Note, however, that the thematic categories were developed for Call 1 projects that focused, among other things, more on software development methodologies underpinning the differentiation between 'industry' (market, commercial) and 'developers' (individuals, domain of engineering). As a result, for example, software middleware providers fit the industry cell rather than developers' one unless they impact the domain of software development. These Call 1 categories may therefore need further tweaking which can be done by further research, in particular, focus group that can be employed to validate these categories more extensively.

Table 1 yields insights 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, SOCIETIES aims to develop a complete, integrated Community Smart Space offering new and powerful ways of working, communicating and socialising impacting all communities and industries, while developers such as software engineers and researchers are FITTEST' main beneficiaries.

	Academia	Developers	(End) Users	Industry
4CAAST			X	X
CLOUD-TM		X		
ACSI				X
CONTRAIL			X	X
I2WEB			X	
OMELETTE	X	X		X
SERENOA		X		
SRT-15				X
ALERT		X		
CUMULONIMBO				X
INDENICA				X
OPTIMIS			X	X
SOCIETIES	X	X	X	X
VISION CLOUD				X
CHOREOS				X
FASTFIX		X	X	X
PLAY				X
WEBINOS			X	
CLOUD4SOA		X		X
FITTEST	X	X		
MOSAIC		X	X	
REMICS	X	X		
SOCIOS			X	

Table 1 - Call 5 projects' users/beneficiary dimension

In order to visualise the findings presented in this table a Venn diagram was generated so as to facilitate the understanding about how the projects can be grouped in relation to their users/beneficiaries.

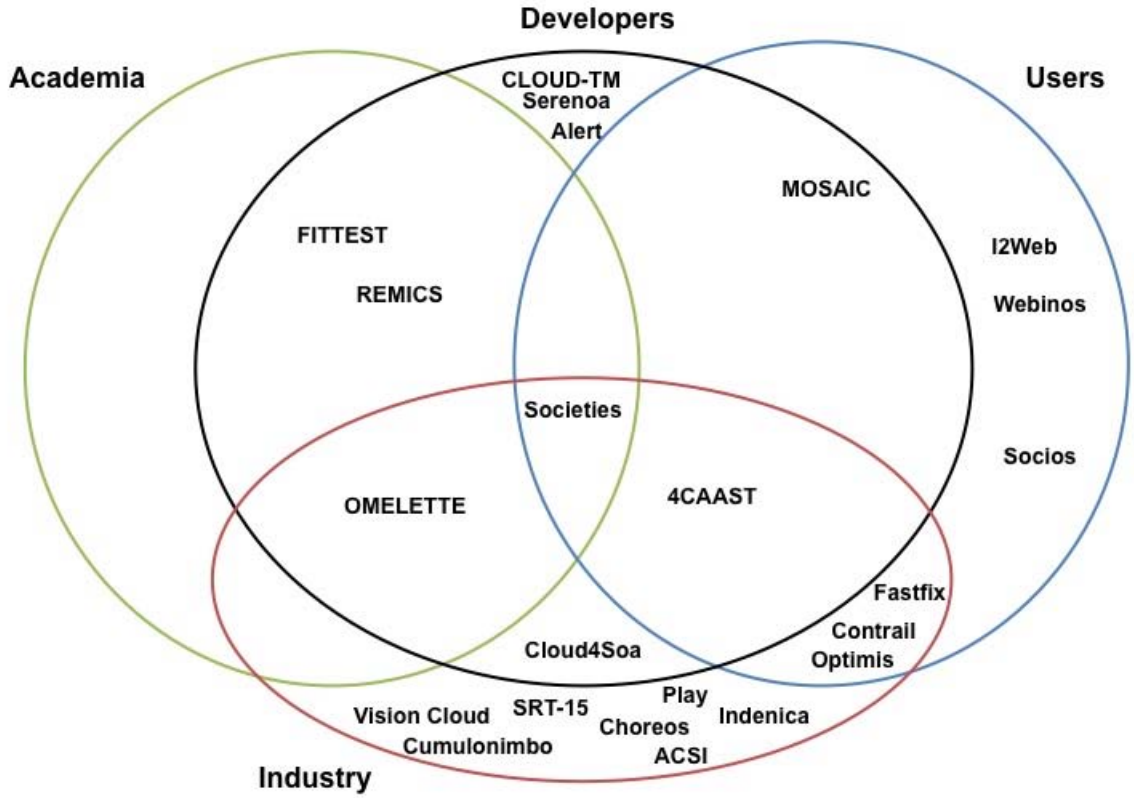

Figure 4 – Venn diagram of Call 5 projects' users/beneficiary dimension

From this figure it can be gathered that the research community is the least targeted beneficiary, while the industry is the most often targeted beneficiary for Call 5 projects. Societies is the only project with the objective targeting all four beneficiaries. Most projects have two or three beneficiary 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 importance.

In the analysis of the Call 5 projects, we also identified some technical topics that can be considered generic: in this way it is possible to recognise a relation among the topics and the different projects. The matrix below is a graphical representation of these connections: each dot indicates that the specific topic is one important technical aspect tackled by the corresponding project. It is important to underline that this representation aims to show only the main technical themes of the projects, considering that every project has a great technical complexity. In the last row of the matrix there is also a graphical view of aggregated results: the circle dimension points out the number of occurrences for each topic. Cloud, virtualization, SOA and semantics are technologies that are studied and used in most of the projects.

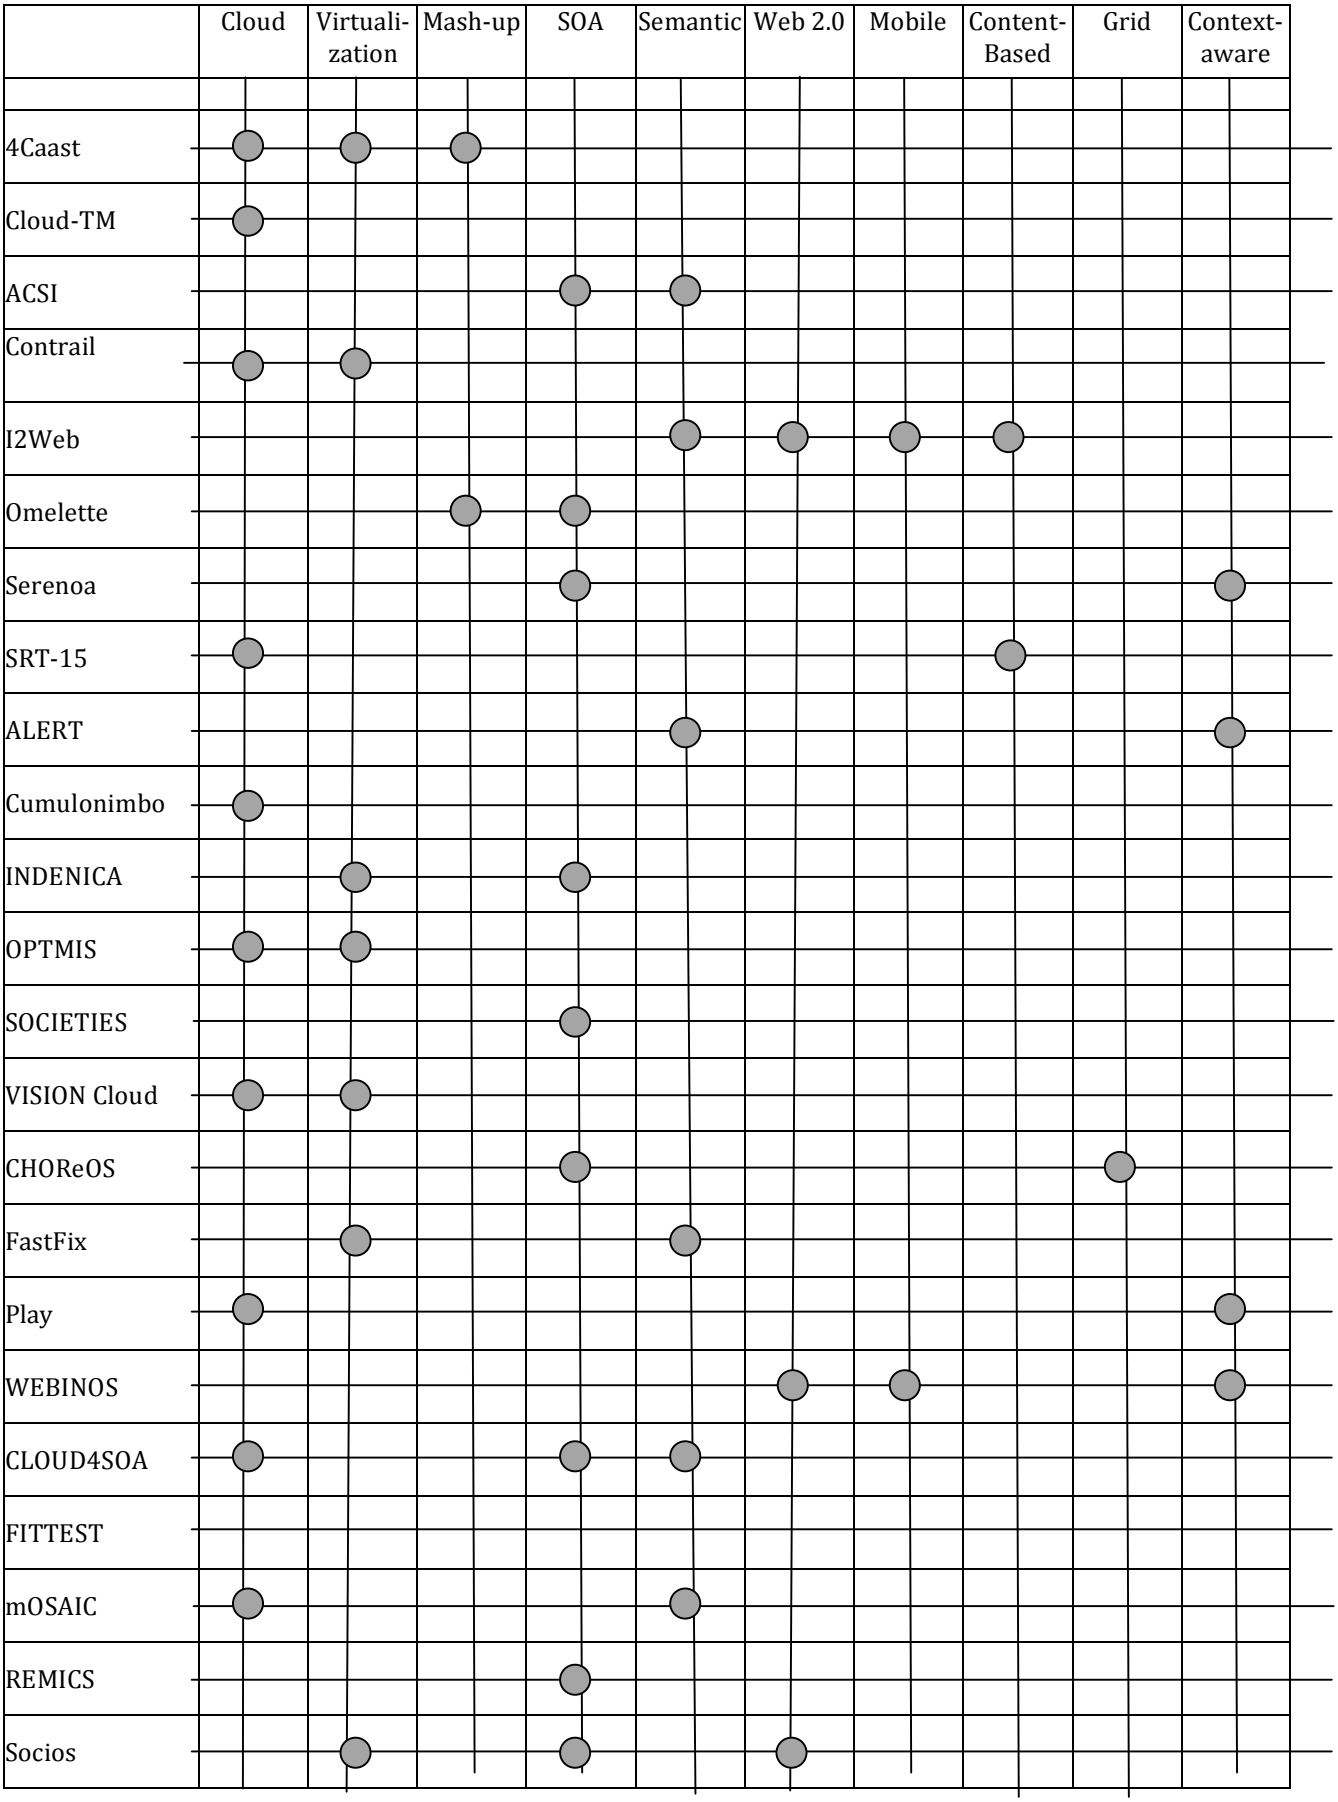

Figure 5 – Call 5 projects' technological dimension

The following tag cloud was created automatically using the text of the technical descriptions of the Call 5 projects. It is simple to see that words like “services” “data”, “cloud” and “information” have many occurrences and represent important topics for the projects analysed: this is also consistent with the previous technical matrix. The tag cloud also shows the prominence of the word “monitoring”, which is an important functionality of several platforms and applications (that are two other “big” words in the cloud), for example in the QoS management that has been taken into account by different projects.

Figure 6 – Call 5 project tag cloud of the technological dimension

This chapter has sought to yield insights into the projects under analysis and the beneficiaries targeted by the 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

In this section we present a short description of the deliverables and papers that were available for Call 5 projects at the time of writing (the last update occurred on March 21st). Thanks to the focus group sessions and communication activities we are in contact with the majority of the projects and we will be monitoring their deliverable production in the future, looking in particular to those documents that present use case definition and deployment.

Tab. 2 – Short description of projects' deliverables and publications

4CaaSt	**D1.1.1 Analysis of the State of the Art and Definition of Requirements** The deliverable is divided into two parts: part D1.1.1a describes the State of the Art, while user and technical requirements are described in D1.1.1b. The deliverable *D1.1.1a* defines evaluation criteria for the SaaS, PaaS and IaaS layer and applies them to a considerable number of players in the respective fields. An overview on related scientific work and related standards accompanies it. In the second part of the document, an outlook is provided on the work package specific state of the art and of the activities that are going to be driven by the technical work packages. This document covers about 30 SaaS, PaaS and IaaS solutions with specific evaluations based on the evaluation criteria. It thereby provides a comprehensive overview of the current market and combines that with an outlook on planned 4CaaSt activities. The deliverable D1.1.1b introduces the concept of User Stories, their elicitation from the work packages and End-To-End Use Cases, their initial prioritization and planned evolution, as well as a set of harmonized roles used to structure all User Stories. This document contains about 120 User Stories with priorities assigned.
CLOUD-TM	**D1.1 User Requirement Report** In this deliverable, the key functional requirements that will be supported by the Cloud-TM middleware platform, and in particular by its two main logical modules, the Data Platform and the Autonomic Manager, are specified. The document provides a series of case studies for the Cloud-TM platform that are based on two pilot applications, Al-Colony and GeoGraph. Al-Colony is a particular class of Massively Multiplier Online Game called office games, while GeoGraph is a framework for the development of Location-based Mobile Social Networking. These case studies provide concrete examples that motivate the main requirements described in the deliverable. In particular, the deliverable analyses in each case the requirements of the applications in terms of consistency, importance of in-memory or persistent data, possibility to define the desired level of quality, self-optimization and disaster recovery. **D1.2 Enabling technologies** This deliverable describes some of the key tools and technologies that the project partners have identified as potential candidates to be integrated into, or influence the development of, the Cloud-TM platform. In particular it surveys some of the most important technologies in the following areas: software transactional memory layers and related mechanisms, frameworks for distributed computing and tools for Cloud Computing.
ACSI	Marlon Dumas, Luciano Garcia-Banuelos1, and Artem Polyvyanyy, **Unraveling Unstructured Process Models**, Proceedings of the 2nd International Workshop on the Business Process Modeling Notation (BPMN), Potsdam, Germany, October 2010. Springer. This paper tackles the problem that many process models found in practice are not well-structured, and it is not always feasible or even

	desirable to restrict process modellers to produce only well-structured models (BPMN models when splits and joins are always paired into single-entry-single-exit blocks). Also, not all processes can be captured as well-structured process models. An alternative to forcing modelers to produce well-structured models, is to automatically transform unstructured models into well-structured ones when needed and possible. This paper summarizes some of the results on automatic transformation of unstructured process models into structured ones, including the initial results on an on-going research effort aiming at developing a complete method for structuring BPMN process models that is implemented in a tool called BPStruct. Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Riccardo Rosati, **View-based Query Answering in Description Logics: Semantics and Complexity**, Journal of Computer and System Sciences, 2011. This paper deals with the problem of answering a query based only on the precomputed answers to a set of views (View-based query answering). This problem that has been investigated in databases, it is unexplored in the context of Description Logic ontologies which may express several forms of incomplete information, and this poses challenging problems in characterizing the semantics of views. The authors present a general formal framework for view-based query answering in DL ontologies: users pose queries to a system, whose knowledge is represented by an ontology expressed in a given DL, the system associates to each user (or, class of users) a set of views, whose extensions are computed as certain answers to the ontology. The system answers user queries coherently with the ontology, though hiding information not implied by the views. It's showed also that different definitions can be provided for the semantics of view-based query answering, each one capturing specificc properties for the notion of solution. The framework defined is related to the problem of privacy-aware access to ontologies, by illustrating how view-based query answering is able to conceal from the user the information that are not logical consequences of the associated authorization views. At the end the paper provides decidability results, algorithms, and data complexity characterizations for view-based query answering in several Description Logics, ranging from those with limited modelling capability to highly expressive ones.
OMELETTE	Florian Daniel, Agnes Koschmider, Tobias Nestler, Marcus Roy, Abdallah Namoun, **Toward Process Mashups: Key Ingredients and Open Research Challenges,** Proceedings of Mashups 2010, December 2010, ACM. Mash-ups as applications have specialized into data mashups, service mashups, or user interface mashups, although these terms lack a common agreement on definitions, while other types of mashups can still be identified. In fact, recently the term process mashup emerged, yet, again, its meaning is everything but clear. This paper try to understand what process mashups are identifying three dimensions that distinguish process mashups from most of the current types of mashups: support for multiple users, support for navigation among multiple connected, pages, and support for workflows. In the virtual space, defined by these three dimensions the authors identify new types of mashups (Simple Mashup, Multi-page Mashup, Guided Mashup, Page Flow Mashup, Shared Page Mashup, Shared Spaces Mashup, Cooperative Mashup, Process Mashup). Every type of these new mashups is discussed showing suitable application scenarios and specific support tools. José Ignacio Fernández-Villamor, Jacobo Blasco-García, Carlos A. Iglesias, Mercedes Garijo, **A Semantic Scraping Model for Web Resources - Applying Linked Data to Web Page Screen Scraping,** Proceedings of the 3rd International Conference on Agents and Artificial Intelligence, 2011. This paper tries to propose a solution for semantic web relating with the scraping technology. In particular the authors propose a

	generic framework for web scraping for the extraction of RDF graphs that represent content in HTML documents. The framework allows defining services based on screen scraping by linking data from RDF graphs with contents defined in HTML documents. This framework is structured in three levels: scraping services, semantic scraping model and syntactic scraping. The first level provides an interface to generic applications or intelligent agents for gathering information from the web at a high level. The second level defines a semantic RDF model of the scraping process, in order to provide a declarative approach to the scraping task. Finally, the third level provides an implementation of the RDF scraping model for specific technologies. The work has been validated in a scenario that illustrates its application to mashup technologies: The scenario had the goal of showing the most commented sports news on a map, according to the place they were taken.
SERENOA	**D2.1.1. Context Aware Design Space and Context Aware Reference Framework** This deliverable describes the CADS and the CARF: two important components to the development of the framework for adaptation of services front-end according to context awareness. CADS stands for Context-aware Design Space, a graphical representation that provides an overview of the possible dimensions for adaptation and their respective abstraction levels. CARF consists of a Context-aware Reference Framework that refines and details the dimensions of the CADS. These methods are complementary and represent a unified approach that joins context-information and adaptation concepts, such as dimensions, levels and techniques. This deliverable defines the methodologies and also presents the background information required to understand the context and motivation for the development and application of them, as well as examples of use and applications of these methods. The investigation identifies and describes in details the adaptation elements and the related techniques. **D1.3.1. Serenoa Roadmap** This deliverable provides an overview of the SERENOA Roadmap. It presents the current main challenges/trends, both from a technological and a scientific perspective. The technological trends relevant for the SERENOA project are those related to Interactive devices in the mass market: Software for service front ends, Software for adaptation in service front ends (including authoring tools). In terms of interactive devices there are two main trends: considerable increase of the number of users, and increasing variety of interaction capabilities. The deliverable provides also the vision that each SERENOA partner owns of the Project, as well as the contribution that each partner plans to provide in order to support its concrete realization.
OPTIMIS	**D1.1.1.1 Requirement Analysis** The purpose of this document is to capture and structure the project requirements in order to ensure a number of qualities such as precision, completeness, consistency, traceability, and testability. The requirements are elaborated based upon: A full project glossary capturing the domain vocabulary. A scenario approach to discover functional requirements of various stakeholders, using precise UML 2.0 sequence diagrams, both for normal and exceptional cases. A goal oriented methodology, to structure those scenarios into goal trees with a clear distinction between requirements under the responsibility of some component to design, assumptions on the existing environment, and the ability to check for consistency and completeness. The ISO-9126 standard to document all the project non-functional requirements. The definition of Traceability mechanisms to reference requirements and also show the coverage between goals and scenarios, directly useful for verification and validation test cases.

Comment: particularly important is the use of the ISO 9126 specifications in order to build up a comparable metric in order to assess the achievement of the (requirements) goals.

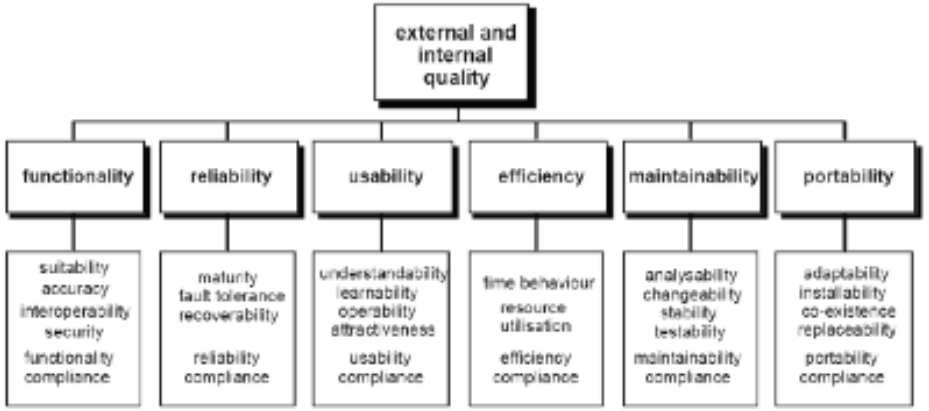

D1.2.1.1 Architecture Design document

This document contains a high-level view of the OPTIMIS Toolkit, including definition of the main architectural components, identification of interfaces among these, and analysis of communication patterns. The component descriptions reference the respective requirements (based on Deliverable 1.1.1.1 Requirements report) that are realized by the components. The focus is on software tools for SPs and IPs, including the OPTIMIS Base Toolkit for monitoring and data mining, as well as handling of aspects related to trust, risk, eco-efficiency, and cost. On top of this base toolkit, higher-level functionalities for, e.g., service construction (including development, integration, and composition), service deployment optimization, admission control, Virtual Machine (VM) placement, data management, elasticity, etc. Selection, composition, and policy customization of these components enable SPs and IPs to provision sustainable IT services capable of satisfying key societal and economical needs.

To validate that the components of the toolkit are appropriately designed, this architecture document also contains illustrations of how the toolkit components can be used in the proposed many-cloud scenarios (federated clouds, multi-clouds, bursting, and brokerage). Another purpose of this architecture report is to guide the design of specific components of the toolkit, and thus act as a starting point for the detailed design work to be carried out in the subsequent months, resulting in the detailed design meta-deliverable (D1.2.2.1).

D1.1.2.1 Verification and Validation Cases

This document on the Verification and Validation (V&V) Cases provides in high-level detail the testing roles, tasks and test cases for the OPTIMIS project. Data information from the deliverable D1.1.1.1 “Requirements Analysis” was used as input. This document will be used at various stages of the project to demonstrate that the project output fits the initial goals. (V&V) is a systems engineering discipline which evaluates software in a systems engineering context, throughout the development and maintenance life cycle.

Comment: based on the requirements deliverable, describes how will be verified and validate the three use cases. Starting from the

	ISO 9126, will be defined a set of indicators and, for each test case, will be recorded the measurement of the indicators identified (in a already developed template).
CHOReOS	**D1.1 CHOReOS State of the Art, Baseline, and Beyond** The purpose of this document is to clarify the baseline, progress, and state-of-the-art that CHOReOS will address. As such, the deliverable provides an updated version of the content provided in the project's DoW. The document is structured in 4 sections, which respectively concentrate on WP1 to WP4. Each section is further decomposed into a number of subsections, one for each of the main scientific and technological areas investigated (as identified by the WP tasks) for which they set: (i) State of The Art (SoTA), (ii) baseline built upon and (iii) progress to be made by CHOReOS. The goal of WP1 is to unambiguously characterize the paradigms underlying the Future Internet according to the CHOReOS vision, with a special emphasis on enabling choreographies in such Ultra-Large-Scale (ULS) networking environments. The Future Internet demands a dynamic and user-centric development process from design to run-time, spanning all the software life cycle. WP2 is about defining such a process and devising its related methods and tools to enable the dynamic development of adaptable, QoS-aware, highly scalable choreographies. The execution of choreographies within the Future Internet heavily relies on adequate middleware support that must: (i) leverage the diversity of networks and networked resources that get federated, and (ii) implement the service-oriented abstractions subsumed by the CHOReOS dynamic development process and associated methods and tools. Providing such a support forms the core objective of WP3. The CHOReOS work package WP4 investigates strategies and mechanisms necessary to establish and to exercise governance in a wide heterogeneous inter-organization setting where integration is defined via large scale service choreographies.
CLOUD4SOA	Nikolaos Loutas, Vassilios Peristeras, Thanassis Bouras, Eleni Kamateri, Dimitrios Zeginis, Konstantinos Tarabanis, **Towards a Reference Architecture for Semantically Interoperable Clouds** This paper, presented during the IEEE Second International Conference on Cloud Computing Technology and Science, 2010, focuses on the emerging problem of semantic interoperability between heterogeneous cooperating Cloud platforms. It tries to pave the way towards a Reference Architecture for Semantically Interoperable Clouds (RASIC). To this end, three fundamental and complementary computing paradigms: Cloud computing, Service Oriented Architectures (SOA) and lightweight semantics are used as the main building blocks. The open, generic Reference Architecture for Semantically Interoperable Clouds introduces a scalable, reusable and transferable approach for facilitating the design, deployment and execution of resource intensive SOA services on top of semantically interlinked Clouds. In order to support the development of semantically interoperable Cloud systems based on RASIC, the model of a common Cloud API is also specified.
MOSAIC	**D1.1 Architectural design of mOSAIC's API and platform** The mOSAIC project considers the need of Sky computing and open-source solutions, the need of semantics in order to express diversity and to focus on with data intensive applications, in Cloud APIs as solution to programme e-infrastructures and in patterns as way to attract new user communities. This deliverables presents the main existing weaknesses of the current research and practice in Cloud computing and the main beliefs and vision of the mOSAIC team that leads their R&D activities. It describes also the main architectural layers of the project technical solution, its components and details about the technologies that are used in its implementation. The main concepts that are exposed are related to the set of unified APIs, Cloud components, platform services, Cloud agency, Cloud

	ontology, Semantic engine, Cloud usage patterns, or SLA management. **D3.1 Cloud usage patterns** This deliverable presents cloud usage patterns that are identified from an application specific view. It is focused mainly on a resource-view of Cloud usage patterns, from application point of view, resource usage patterns, virtualisation techniques, etc. Cloud usage patterns are designed as workflows of services that are using mOSAIC API. The mOSAIC platform must evolve as an open environment, minimizing vendor lock-in and increasing customer choice. Therefore, this text highlights requirements that need to be standardized in a cloud environment to ensure interoperability, ease of integration and portability of existing and future applications. It must be possible to implement all of the scenarios described in this deliverable without using closed, proprietary cloud technologies. It also presents in Annex the most typical cloud use cases and their respective characteristics and needs by the stakeholders which have been defined in the "Cloud Computing Use Cases. White paper." Beniamino di Martino, Dana Petcu, Roberto Cossu, Pedro Goncalves, Tam as Mahr, and Miguel Loichate, **Building a Mosaic of Clouds**, 2010 This article presents the concepts and expectations of the mOSAIC Project which intends to create and promote an open-source Could application programming interface and a platform targeted for developing multi-Could oriented applications in order to overcome the limits of actual cloud computing offers. Aversa, R.; Di Martino, B.; Rak, M.; Venticinque, S., **Cloud Agency: A Mobile Agent Based Cloud System** This article, presented during the Complex, Intelligent and Software Intensive Systems (CISIS) International Conference, 2010, describes the solution developed by mOSAIC project regarding the integration of Cloud on GRID architecture with a mobile agent platform. The architecture proposed offers Virtual clusters with full administrative control to final users, adopting an existent GRID architecture and especially its security infrastructure. The mobile agent platform is able to dynamically add and configure services on the virtual clusters. The experience presented shows that the mobile agent paradigm well fulfils the dynamic properties of the Cloud paradigm and could be a good choice to simply develop application and services able to dynamically adapt themselves to the virtualized environment. Petcu, D., **Identifying Cloud computing usage patterns** This article, presented during the IEEE International Conference 2010, describes the aims of a generic API for Cloud-based applications, discuss the existing Cloud use cases, and exposes the design of a collection of Cloud usage patterns.
REMICS	**D2.1. State of the art on modernization methodologies, methods and tools** This deliverable presents the analysis of the current state-of-the-art of methodologies, methods and tools for application modernization and Service Clouds paradigms. Each section provides an analysis of related projects, products, standards and approaches in addition to conclusions and pointing out shortcomings and ideas for REMICS research. The technologies areas covered are: Software development methodologies with focus on model-driven, service-oriented, migration and agile methodologies; Model-driven modernization of legacy systems; Cloud platforms and SOA; technologies as well as design patterns; Model execution and Models@runtime; Model-driven interoperability; Model-driven verification and testing.

CONCLUSIONS

The analysis of the documentation available for Call 5 projects, even if less rich than that which was available for Call 1 projects, confirms our approach developed so far in terms of methodology definition and data gathering instrument (namely the second questionnaire). The second questionnaire was at the centre of the focus group sessions and, thanks to the feedback collected on those occasions, the SEQUOIA consortium is now almost ready to distribute it. As planned, this will be done using the online application already used for the preliminary questionnaire.

Moreover, Call 5 projects showed a vivid interest in SEQUOIA and their participation in the five sessions of the focus group organised while writing this deliverable informed our work in a positive way. We opened up important channels of communication with Call 5 projects, and this will result in a more efficient exchange of information. In fact, the SEQUOIA researchers will constantly monitor the production of documentation of Call 5 projects, but the direct contacts developed will make also possible to receive, or ask for, updates and more precise information directly from the projects.

With this deliverable the documentation mapping can be considered concluded; analysing the materials produced by Call 5 projects (even if only at the beginning of their life-cycle) was very interesting and added important information to that which we had previously gathered about Call 1 projects. Thanks to the focus groups and to the second questionnaire that will be soon be launched, we will know more and more about the individual projects and we will be able to finalise the SEQUOIA self-assessment methodology that - presented to the projects during the focus group - has now reached an interesting level of development.

ANNEX A: PROJECTS DESCRIPTION AND COMPLETE LIST OF DELIVERABLES[1]

4CAAST
Building the PaaS Cloud of the Future

The 4CaaSt project aims to create a PaaS Cloud platform which supports the optimized and elastic hosting of Internet-scale multi-tier applications. 4CaaSt embeds features that ease programming of rich applications and enable the creation of a business ecosystem where applications from different providers can be tailored to different users, mashed up and traded together.

Cloud computing is essentially changing the way services are built, provided and consumed. Despite simple access to Clouds, building elastic services is still an elitist domain and proprietary technologies are an entry barrier especially to SMEs and consequently, it remains largely within the domain of established players. The project will bring significant benefits to the European economy. It will provide an easy to use Infrastructure for a More Competitive Environment, greatly simplifying design and delivery of tailored services and compositions.

Shifting focus from technology to business aspects will accelerate the Massive Uptake of Value Added Services taking full advantage of Telco communication paradigms. Thus, 4CaaSt ultimately contributes to the Future Internet by its ability to serve the service demands of tomorrow in a highly efficient and scalable manner. It brings Europe's Telco industries the unique opportunity to expand their traditional business towards integrated IT/Telco service platforms as well as allowing Cloud Platform providers to expand the scope of the way the platform can be utilised.

Deliverables:

- D1.1.1 Analysis of the State of the Art and Definition of Requirements, 31/12/2010
- D10.3 Collaboration Plan, 11/02/11
- D10.1 Website and promotional Material, 12/02/2011

[1] Where deliverables are not indicated either there are not deliverables available at the time of writing (but the projects' websites are planned in a way in which seams like deliverables will be available in the future) or deliverables are not going to be pubblished on the projects' website.

Cloud-TM:
A novel programming paradigm for cloud computing

Cloud-TM aims at defining a novel programming paradigm to facilitate the development and administration of cloud applications. It will develop a Self-Optimizing Distributed Transactional Memory middleware that will spare programmers from the burden of coding for distribution, persistence and fault-tolerance, letting them focus on delivering differentiating business value. Further, the Cloud-TM platform aims at minimizing the operational costs of cloud applications, pursuing optimal efficiency via autonomic resource provisioning & pervasive self-tuning schemes.

Deliverables:

- D5.1. Website, 01/07/2010
- D5.2. Project Presentation, 01/07/2010
- D1.1. User Requirements Report, 30/11/2010
- D1.2. Enabling Technologies, 30/11/2010
- D5.3. Initial Collaboration Plan. 30/11/2010

ACSI - Artifact-Centric Service Interoperation

ACSI will serve to significantly reduce the effort and lead-time of designing, deploying, maintaining, and joining environments that support service collaborations.

The project will develop scientific advances, techniques, and tools to simplify the design and deployment of infrastructure to support service collaborations, the ability of services to join such collaborations, and the evolution of such collaborations as the marketplace and competitive landscape change. This will be achieved by developing a rich framework around the novel notions of dynamic artifacts and interoperation hubs.

Deliverables:

- D6.1. Public Website, 29/08/2010
- D6.5. Collaboration Activities Plan, 28/10/2010

Publications:

- Simplifying Schema Mappings, Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Moshe Y. Vardi. Proc. of the 14th Int. Conf. on Database Theory (ICDT 2011), 2011.
- View-based Query Answering in Description Logics: Semantics and Complexity, Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Riccardo Rosati, Journal of Computer and System Sciences, 2011. To appear.
- Evolution of DL-Lite Knowledge Bases, Diego Calvanese , Evgeny Kharlamov, Werner Nutt, and Dmitriy Zheleznyakov, Proceedings of the 9th Int. Semantic Web Conf. (ISWC 2010), 2010. To appear.
- Unraveling Unstructured Process Models, Marlon Dumas, Luciano Garcia-Banuelos1, and Artem Polyvyanyy, Proceedings of the 2nd International Workshop on the Business Process Modeling Notation (BPMN), Potsdam, Germany, October 2010. Springer.
- Conjunctive Artifact-Centric Services, Piero Cangialosi, Riccardo De Masellis, Giuseppe De Giacomo, and Riccardo Rosati, Proceedings of the 8th International Conference of Service Oriented Computing (ICSOC'10), 2010. To appear.
- Artifact-centric Service Interoperability in the Flanders Research Information Portal, De Leenheer, P, Proceedings of the 1st International Workshop on Future Internet Systems, collocated with FIS2010, CEUR.
- Two-Player Game Structures for Generalized Planning and Agent Composition, Giuseppe De Giacomo, Paolo Felli, Fabio Patrizi, and Sebastian Sardiña, Proceedings of the 24th Conference on Artificial Intelligence (AAAI'10), 2010.
- Inconsistency-tolerant Semantics for Description Logics, Domenico Lembo, Maurizio Lenzerini, Riccardo Rosati, Marco Ruzzi, and Domenico Fabio Savo, Proceedings of the fourth International Conference on Web Reasoning and Rule Systems (RR 2010), 2010.

CONTRAIL
Open Computing Infrastructures for Elastic Services

After decades in which companies used to host their entire IT infrastructures in-house, a major shift is occurring where these infrastructures are outsourced to external operators such as Data Centers and Computing Clouds. To allow open access to shared computing resources, the vision of the CONTRAIL project is that any organization should be able to be both a Cloud provider when its IT infrastructure is not used at its maximum capacity, and a Cloud customer in periods of peak activity.

The goal of the CONTRAIL project is to design, implement, evaluate and promote an open source system in which resources that belong to different operators are integrated into a single homogeneous Federated Cloud that users can access seamlessly.

CONTRAIL will vertically integrate an opensource distributed operating system autonomous resource management in Infrastructure-as-a-Service (IaaS) environments, and level services and runtime environments as foundations for Platform-as-a-Service (PaaS).

I2Web
Inclusive Future-Internet Web Services

The Future Internet Community that will be more mainstream in people's lives, may further isolate excluded groups. I2Web will provide tools to develop inclusive Future Internet services that will overcome this widening divide. To enable the Future Internet to be very extensively used by people with disabilities and the elderly, the inclusiveness of its Services Front Ends will be of paramount importance. I2Web particularly responds to immediate challenges of the Future Internet: ubiquitous and mobile Web, media convergence and user-generated content, in combination with cloud computing, Web 2.0 developments, Social Networking, User-Centred Design and Inclusive Design principles.

I2Web will deliver an open, scalable, dependable service platform architecture and online space, enabling automatic service description, discovery, composition, and negotiation with a multiplicity of reusable inclusive services, which may be mobile or nomadic, multi-device, multimodal and multi-context.

Omelette
Open Mashup Enterprise service platform for LinkEd data in The TElco domain

OMELETTE aims to provide new platforms for service composition addressing mass users served by the telco industry. It will provide support for the development of telco services, their flexible usage in mashup environments and their deployment on an open interoperable mashup delivery platform. It will aim at breaking the barriers between the web, telecommunication and hybrid services, taking into account the convergence of IT/telecom/content systems.

In particular, OMELETTE aims at researching on the development, management, governance, execution and conception of converged services with a specific focus on the telco domain. It will create a sound model of mashups that follows the REST architectural style (also supported by standard widget technology), as well as a standard specification of a mashup-containing platform that may guarantee portability and interoperability among different vendors and versions. These concepts will be based on a solid theoretical model of mashup foundations and the specific requirements gained from the telco domain.

Publications:

- A Semantic Scraping Model for Web Resources - Applying Linked Data to Web Page Screen Scraping, José Ignacio Fernández-Villamor, Jacobo Blasco-García, Carlos A. Iglesias, Mercedes Garijo, Proceedings of the 3rd International Conference on Agents and Artificial Intelligence, 2011.
- Toward Process Mashups: Key Ingredients and Open Research Challenges, Florian Daniel, Agnes Koschmider, Tobias Nestler, Marcus Roy, Abdallah Namoun, Proceedings of Mashups 2010, December 2010, ACM.
- Wisdom-Aware Computing: On the Interactive Recommendation of Composition Knowledge, Soudip Roy Chowdhury, Carlos Rodríguez, Florian Daniel and Fabio Casati, Proceedings of WESOA 2010, December 2010, Springer.

SERENOA
Multidimensional, context-aware adaptation of Service Front-Ends

Serenoa is aimed at developing a novel, open platform for enabling the creation of context sensitive service front-ends (SFEs). A context-sensitive SFE provides a user interface (UI) that exhibits some capability to be aware of the context and to react to changes of this context in a continuous way. As a result such a UI will be adapted to a person's devices, tasks, preferences, and abilities, thus improving people's satisfaction and performance compared to traditional SFEs based on manually designed UIs.

Deliverables:

- D7.3. Project Factsheet, 06/2010
- D7.1.1. Up and running Web site, constantly updated, 29/09/2010
- D1.1.1. Requirements Analysis, 25/02/11 (Non-Public)
- D1.3.1. Serenoa Roadmap, 25/02/11
- D2.1.1. Context Aware Design Space and Context Aware Reference Framework, 25/02/11
- D6.1.1. Exploitation Plan, 25/02/11 (Non-Public)
- D7.1.2. Dissemination Actions, 25/02/11
- D7.2.1. Collaboration Plan, 25/02/11

SRT-15
Subscription Routing Technology for 2015

The goal of the SRT-15 research project is to build a scalable platform for connecting business applications and services by enabling the discovery and integration of dynamic enterprise services on the Internet. SRT-15 will allow for dependable and scalable cloud-based processing of data coming to and from a variety of heterogeneous enterprise services spread across multiple distributed locations. In order to be able to embrace the change in the enterprise information processing landscape, SRT-15 relies on technologies that support rapid change: cloud computing, content-based routing, and complex event processing.

Publications:

- SRT-15 Press Release: Un progetto innovativo per una piattaforma di servizi. Romano, L, DeMari D. 2011.
- SRT-15 Press Release: Yahoo! participa en un proyecto paneuropeo de cloud computing, SRT-15. Junqueira, F, Serafini M. 2011.
- SRT-15 Press Release: L'Institut d'informatique reçoit plus de 700'000 francs de l'Union européenne, Rivière E, Felber P., 2010.

Alert
Active support and reaL-time coordination based on Event pRocessing in open source software development.

The aim of the ALERT project is to develop methods and tools that improve FLOSS coordination by maintaining awareness of community activities through real-time, personalized, context-aware notification. ALERT will create an active collaboration platform, i.e. a virtual actor would interact with other developers, process and recognize various kinds of interactions, suggest actions on the basis of these and remember and bring past interactions into the developers' attention, thus enabling developers to work better together. The main objectives of the project are:

- Efficient modelling of the more reactive coordination in FLOSS development that will improve the awareness of the group work and the responsiveness of individuals.
- Efficient management of the awareness of team members that will enable interesting parties to be notified based on their interest/expertise as soon as something relevant happens without overloading them, interfering with and slowing down their work.
- Efficient management of information relevant for FLOSS teams, including the semantic integration of information and its flow between all stakeholders that will support better understanding of the situations which a developer should react on.
- Personalized and task-based access to information, by allowing developers to focus on activities to be performed to achieve a specific shared task and/or by including information about the presence and activities of other developers in the FLOSS.
- Pilot, evaluate and impact the ALERT system in three important FLOSS communities (Mandriva, KDE and Morfeo) and disseminate the results in other FLOSS and relevant research communities.

In order to achieved the above mentioned objectives, following research outcomes will be produced:

- Set of ontologies that model the conceptual dependencies between community, content and interactions.
- Conceptual model of an (re)active collaboration platform.
- Model for declarative specification of alerts.
- Intelligent, pattern-based engine that can correlate, aggregate, and filter events in order to provide triggers for relevant notifications.
- Module for integration of information from structured and unstructured sources.
- Interaction highway set of services (e.g. search, recommendation, visualisation, etc.).
- Module for capturing explicitly or implicitly the interest/expertise of the developer.
- Module for capturing relevant information for the task at hand.

CumuloNimbo
A Highly Scalable Transactional Multi-Tier Platform as a Service

CumuloNimbo will deliver a new generation of cloud Platforms as a Service (PaaS) that will provide consistency, availability, and simpler programming abstractions, such as transactions. Current practices in cloud computing sacrifi ce data consistency for scalability, thus increasing the complexity of building applications on top of such platforms when strong consistency guarantees are necessary.

CumuloNimbo will go beyond the state of the art by providing a scalable PaaS with full transactional guarantees, that is, without requiring any data partitioning. Th e support of full transactional guarantees will provide full coherence to applications developed on top of CumuloNimbo resulting in a platform amenable to be used by regular application developers with modest IT skills.

CumuloNimbo will deliver a cloud computing PaaS that that will exhibit high scalability while guaranteeing full consistency. CumuloNimbo will follow a radically diff erent approach from current state of the art approaches in cluster and cloud environments. Cluster environments have been based on replication approaches that provide strong consistency but at the cost of limited scalability. For this reason cloud environments rely exclusively on data partitioning and therefore do not provide data consistency across data partitions.

The PaaS sector is foreseen to reach a market of $16 Billion by 2016 according to Forrester Research. Soft ware vendors in Europe need to pursue the cloud market to keep their market shares. For this, they need access to a scalable PaaS. CumuloNimbo will off er such a PaaS system with the clear distinguished added value that lies in providing high scalability whilst guaranteeing transactional consistency as opposed to the current state of the art.

As a result of the fragmentation and the increased dependency on external service and/or platform vendors that is mentioned above, companies will, on the one hand, be forced to create Virtual Service Platforms to shield their application infrastructures from those influences as much as possible. On the other hand, for domains where little or no appropriate platforms exist today, advanced interoperability and portability constraints will have to be taken into account when designing new service platforms. Depending on the actual application domain, the shape of such a Virtual Service Platform may vary significantly, ranging from a thin layer providing abstraction over existing services to the other extreme where it may be a large building block integrating in-house and external services, existing legacy components, and assets. INDENICA will research on the problems outlined above to support the development of virtual domain-specific platforms while emphasizing openness and interoperability.
In order to support the development of these interoperable infrastructures, INDENICA aims to provide methods, architectures, components, tools and assets for reuse-based creation of the adapted platforms. Effectively, INDENICA will research on a product line for the development of domain-specific platforms that are interoperable by design. Members of the product line (i.e., platforms) can be modified after deployment, thus providing capabilities to make decisions that are currently made at design time at runtime.

In particular, INDENICA strives for the following research results:

- It aims to simplify the overall complexity of service platform development caused by the mentioned fragmentation. This will be achieved through methods, architectures and tools that support the development of context and domain-specific platforms along with the services.
- It targets to support platform convergence and interoperability to avoid the increased dependency on external service and platform vendors. This will be achieved by providing a common basis for families of platforms which integrate system management and interoperability capabilities right from the start.
- It aims to provide a reusable infrastructure for platform development and support the automatic deployment and the monitoring, governance, and adaptation of services in a Virtual Service Platform.

Optimis
Optimized Infrastructure Services

The motivation for OPTIMIS is the vision that hybrid clouds will become commonplace, realized by private clouds interacting with a rich ecosystem of public and other cloud providers. OPTIMIS is aimed at enabling organizations to automatically externalize services and applications to trustworthy and auditable cloud providers in the hybrid model. Consequently, OPTIMIS believes its activities will support and facilitate an ecosystem of providers and consumers that will benefit from the optimal operation of services and infrastructures. The optimization covers the full lifecycle of services and their interactions.

Main innovations

- OPTIMIS will identify, capture and codify what an optimized cloud ecosystem driven by trust, risk, eco-efficiency and cost will look like.
- The OPTIMIS framework and toolkit will simplify service construction, and support deployment and runtime decisions based on prior evaluation of providers.
- OPTIMIS will facilitate the use of resources based on economic and eco-efficiency goals while achieving a dynamic and proactive management of cloud infrastructures.
- OPTIMIS deliverables will enable clouds to be composed from multiple services and resources. It will support service brokerage via interoperability, and is architecture-independent.
- OPTIMIS will identify value networks for these cloud markets and recommend suitable legal and regulatory guidelines for operation.
- Expected outcomes and use cases

OPTIMIS will deliver a specification and toolkit that can be used to construct these next-generation cloud architectures. There will be yearly code drops with incremental functional enhancement.

The project will produce three use cases:

- The Cloud Programming Model use case
- The Cloud Brokerage use case
- The Cloudbursting use case

Deliverables:

- D1.1.1.1. Requirement Analysis, 30/09/2010
- D1.2.1.1. Architecture Design document, 30/09/2010
- D1.1.2.1. Verification and Validations Cases, 15/11/2010
- D7.1.1. Cloud Ecosystem: Taxonomy and Value Networks, 30/09/2010
- D7.1.2. Cloud Business: a deeper dive into the SlaaS market, 31/01/2011
- D5.1.3.1. Report on Cloud License Management existing technologies and Requirements, 31/01/2011
- D7.2.1.1. Cloud Legal Guidelines, 31/01/2011

Societies
Self Orchestrating CommunIty ambiEnT IntelligEnce Spaces

SOCIETIES will radically improve the utility and scope of future Internet services by merging social computing and pervasive computing through the design, implementation and evaluation of an open scalable service architecture and platform for self-orchestrating Community Smart Spaces.

SOCIETIES Objectives

The goal of SOCIETIES will be achieved through four key objectives:

- To facilitate the creation, organisation, management and communication of communities via Community Smart Spaces, where pervasive computing is integrated with social computing communities;
- To provide an enhanced user experience for both individuals and entire user communities, based on proactive smart space behaviour and dynamic sharing of community resources across geographic boundaries;
- To design and prototype a robust open and scalable system for self-orchestrating Community Smart Spaces;
- To evaluate, through strong involvement of end-users, the usefulness and acceptance of the developed CSS software via three user trials with the following groups:
    - Enterprise Users: Enterprise communities play an important role in bringing together people, goods and services within global markets, local ecosystems or large organisations. The CSS concept will bridge the gap between smart IT systems and established enterprise community activities.
    - Students: Students adapt easily to new technology, and since communication and social networking play an integral role in their lives, they are most likely to adopt CSSs, using them in ways both foreseen and unforeseen.
    - Disaster Relief Experts: The ability to rapidly form a disaster management community from all the closely located relief teams can help save lives, property, and the environment.

Vision Cloud
Virtualized Storage Services Foundation for the Future Internet

The goal of VISION Cloud is to introduce a powerful ICT infrastructure for reliable and effective delivery of data-intensive storage services, facilitating the convergence of ICT, media and telecommunications. This infrastructure will support the setup and deployment of data and storage services on demand, at competitive costs, across disparate administrative domains, while providing QoS and security guarantees.
The VISION Cloud project is based on the following ideas:

- Raising the Abstraction Level of Storage: encapsulating storage into objects with user defined and system-defined attributes (a.k.a. “metadata”). This metadata supports effective access, management and manipulation of the storage and stored data and derivation of new information. With that, we enable high scalability and simplification of all storage and data functions.
- Data Mobility and Federation without Boundaries: substantially extending the limited data migration capabilities available in contemporary infrastructures to migrate and federate data across geographically distributed administrative domains, ensuring comprehensive and transparent data interoperability and overcoming the problem of data lock-in.
- Computational Storage: providing an architected solution for bringing computing to storage, enabling secure execution of computational tasks near their data, applicable both to user-driven computational tasks as well as autonomous data derivation and transformation.
- Content Centric Storage: facilitating access to data by content and its relationships, rather than details of the underlying storage containers, further supported by mechanisms to define domain-specific storage optimisations, and altogether leading to highly simplified and efficient storage access.
- Advanced Capabilities for Cloud-Based Storage: combining the base concepts of cloud computing with additional capabilities of storage virtualisation, advancing the state of the art by creating a platform where data-intensive services can be provided securely, at the desired level of service, and at competitive costs.

Choreos
Large Scale Choreographies for the Future Internet

CHOReOS will implement a framework for scalable choreography development. The goal is to enable domain experts to develop decentralized ultra-large scale (ULS) solutions composed of heterogeneous services that are adaptable and QoS (Quality-of-Service) aware. Prior to this, these solutions were only possible with the support of dedicated IT professionals to provide the skills needed for architectural design and software engineering.

CHOReOS will deliver formally grounded abstractions and models, dynamic choreography-centric development processes, governance and service-oriented middleware manipulated via an Integrated Development Runtime Environment (IDRE) aimed at overcoming the ULS impact on software system development.

Here is what to expect from the CHOReOS project:

- Abstractions and models
- A choreography-based development environment
- Service-oriented middleware
- Governance mechanisms

Deliverables:

- D1.1. CHOReOS State of the Art, Baseline, and Beyond, 21/01/11

Publications:

- Towards Verification and Validation of Choreographies, Felipe M. Besson, Pedro M. B. Leal, Fabio, University of S˜ao Paulo (USP), Technical Report No: RT-MAC-2011-01

FastFix

Fastfix
Monitoring Control for Remote Software Maintenance

FastFix will enable time- and cost-efficient maintenance and support services, by monitoring software applications, replicating semantic execution failures, and automatically generating patches.
Software maintenance and support services are key factors to customers' perception of software quality. Customers are more demanding about these services, while contribution of maintenance to products total cost of ownership should be reduced.
Reducing maintenance costs is even more crucial for SME's. Their software products often cover specific niches. To take full advantage of this, SMEs must get access to global markets, what is usually hampered by the lack of an adequate on-site customer support.
Software vendors need a system to remotely provide a high quality support service to their customers, improve user experience and facilitate corrective, adaptive and preventive maintenance – of both new and existing software products.
FastFix results will include a platform and a set of open source tools to on-line monitoring of execution environments, gathering semantic information on application and user behaviour. This information is sent in real time to a support centre, taking special care on privacy and security issues. Using event correlation techniques, FastFix identifies failure symptoms, performance degradation or changes in user behaviour and allow for failure replication, patch generation and patch deployment, resulting in a self-healing software application.

Main objectives are to develop:
(1) tools to gather context information on user and application,
(2) a run-time with minimum impact on application performance,
(3) a secure method to send this information to a centralized fault analysis platform,
(4) a tool to detect software failures, undesirable execution trends and performance degradation,
(5) a platform to replicate failure conditions within a virtual machine, and
(6) a tool to generate change strategies and necessary patches.

Play
Pushing dynamic and ubiquitous interaction between services Leveraged in the Future Internet by ApplYing complex event Processing

The PLAY project will develop and validate an elastic and reliable architecture for dynamic and complex, event-driven interaction in large highly distributed and heterogeneous service systems. Such architecture will enable ubiquitous exchange of information between heterogeneous services, providing the possibilities to adapt and personalize their execution, resulting in the so-called situational-driven process adaptivity.

The main goal is to provide an open highly distributed Platform for event-driven interaction between services that scales at the Internet level based on a federated architecture able to address the multiplicity and the heterogeneity of service networks and address Quality of Service (QoS) requirements (such as dependability, security and scalability) in very large scale Internet based networks.

In a nutshell, the Platform will enable combining events from any service and pushing to any service on a global scale. In more general terms, PLAY will lay the foundation for the event-driven, push-based Future Internet, which enables sensing changes in the environment/context and respond correspondingly, including affecting running business processes in an ad-hoc manner.

webinos
Secure Web Operating System Application Delivery Environment

Promoting a “single service for every device” vision, webinos will move the existing baseline of web development from installed applications to services, running consistently across a wide range of connected devices, ensuring that the technologies for describing, negotiating, securing, utilizing device functionalities and adapting to context are fit for purpose.

Innovations in contextual description will be broad covering but not limited to device capabilities, network access, user identity and preferences, location, behaviourally induced properties and finally the more complex issue of the users’ social network context and social media engagement.

webinos will boost the industry migration towards web-based services. webinos can back this by providing inter-operable, standardised, open source technology utilizable across domains with direct commercially exploitable value. webinos will also act as an industry catalyst to encourage collaboration and discourage fragmentation in this space. There are strong industry moves towards Internet friendly and Internet integrated offerings, and there exists a window of opportunity to place the webinos technology on a robust open foundation that will remove economic barriers to engagement, embody policy on data privacy in concrete technology and creating a centre of web centric expertise.

Deliverables:

- Industry landscape, governance, licensing and IPR frameworks (coming soon), 03/11
- Requirements, 03/11
- User Expectations on Security and Privacy, 02/11
- Use cases and scenarios (coming soon), 01/11

Cloud4SOA
A Cloud Interoperability Framework and Platform for user-centric, semantically-enhanced service-oriented applications design, deployment and distributed execution

Cloud4SOA focuses on resolving the semantic interoperability issues that exist in current Clouds infrastructures and on introducing a user-centric approach for applications which are built upon and deployed using Cloud resources. To this end, Cloud4SOA aims to combine three fundamental and complementary computing paradigms, namely Cloud computing, Service Oriented Architectures (SOA) and lightweight semantics. The project will propose a reference architecture and deploy fully operational prototypes.

The Cloud4SOA project contributes towards Cloud interoperability through the following research and innovation related activities:

- Identify and analyse the semantic interoperability problems within the Cloud by carrying out a comprehensive analysis on both IaaS and PaaS layers, while including stakeholder experience, requirements and feedback. The Cloud Semantic Interoperability Framework will be the result of this exercise, acting as a roadmap for creating such user-oriented platforms.
- To resolve the most problematic semantic interoperability issues through the Cloud4SOA Reference Architecture which introduces an open, generic architecture for a semantically interoperable Cloud, capitalizing on Service Oriented Architecture, lightweight semantics and user-centric design and development principles.
- Implement, deploy and evaluate the usability, applicability and adaptability of the Cloud4SOA Reference Architecture by developing a proof-of-concept platform which will be applied in different scenarios and setups across 3 European industrial and business showcases.
- In addition to its collaboration with standards bodies, Cloud4SOA will reach out to the SME participants in the European Cloud market, including service (IaaS) and middleware (PaaS) providers, as well as user demands. By highlighting known and potential barriers in interoperability and portability, their consistent feedback is key towards such a framework and platform. This dynamic keeps the project practical and relevant throughout its lifecycle, needed to accurately address this rapidly evolving Cloud market.

Publications:

- "Towards a Reference Architecture for Semantically Interoperable Clouds," Cloud Computing Technology and Science, Nikolaos Loutas, Vassilios Peristeras, Thanassis Bouras, Eleni Kamateri, Dimitrios Zeginis, Konstantinos Tarabanis, IEEE International Conference on, pp. 143-150, 2010 IEEE Second International Conference on Cloud Computing Technology and Science, 2010.

FITTEST
Future Internet Testing

The Future Internet will be a complex interconnection of services, applications, content and media, on which our society will become increasingly dependent for critical activities such as social services, learning, finance, business, as well as entertainment. This presents challenging problems during testing; challenges that simply cannot be avoided since testing is a vital part of the quality assurance process. FITTEST aims to attack the problems of testing the Future Internet with Search Based Testing.

FI applications will be characterized by an extreme level of dynamism. Most decisions made at design or deployment time are deferred to the execution time, when the application takes advantage of monitoring (self-observation, data collection from the environment and logging of the interactions) to adapt itself to a changed usage context.

The realization of this vision involves a number of technologies, including:

- Observational reflection and monitoring, to gather data necessary for run-time decisions.
- Dynamic discovery/composition of services.
- Hot component loading and update.
- Structural reflection, for self-adaptation.
- Asynchronous communication.
- High configurability and context awareness.
- Composability into large systems of systems.

mOSAIC
Open-Source API and Platform for Multiple Clouds

The mOSAIC project aims to develop an open-source platform that enables applications to negotiate Cloud services as requested by their users. Using the Cloud ontology, applications will be able to specify their service requirements and communicate them to the platform via an API. The platform will implement a multi-agent brokering mechanism that will search for services matching the applications' request, and possibly compose the requested service if no direct hit is found. The platform will facilitate competition between Cloud providers, who, in return, will be able to reach customers they could not reach before.

Using mOSAIC solutions, Cloud-based application developers, maintainers and users will be able:

- to specify the service requirements in terms of a Cloud ontology and communicate them to the platform via the innovative API;
- to postpone their decision on the procurement of Cloud services until run-time;
- to find best-fitting Cloud services to their actual needs and efficiently outsource computations.

The platform will:

- implement a multi-agent brokering mechanism that will search for services matching the applications' request, and possibly compose the requested service if no direct hit is found.
- will facilitate competition between Cloud providers, who, in return, will be able to reach customers they could not reach before.

Deliverables:

- D1.1 Architectural design of mOSAIC's API and platform, 28/02/11
- D3.1 Cloud usage patterns, 07/03/11
- D5.2 Press releases, 07/03/11
- D6,1 Project information package and control procedures, 03/03/11

Publications:

- Building a Mosaic of Clouds, Beniamino di Martino, Dana Petcu, Roberto Cossu, Pedro Goncalves, Tamas Mahr, Miguel Loichate, 2010
- Identifying Cloud computing usage patterns, Petcu, D., 2010
- Cloud Agency: A Mobile Agent Based Cloud System, Aversa, Di Martino, Rak, Venticinque, 2010

REMICS
Reuse and Migration of legacy applications to Interoperable Cloud Services

Cloud computing and Service Oriented Architecture (SOA) are recognized game-changing technologies for a cost-efficient and reliable service delivery. Software as a Service paradigm becomes more and more popular enabling flexible license payment schemas and moving the infrastructure management costs from consumers to service providers. However, building a SaaS system from scratch may require a huge investment in time and efforts. Moreover, the organizations legacy systems are difficult to reuse due to platform, documentation and architecture obsolescence.

The goal of REMICS is to develop advanced model driven methodology and tools for REuse and Migration of legacy applications to Interoperable Cloud Services. Service Cloud paradigm stands for combination of cloud computing and SOA for development of Software as a Service (SaaS) systems.

Deliverables:

- D2.1. State of the art on modernization methodologies, methods and tools, 10/03/11
- D7.1. Public Website, 17/03/11
- D7.2. Plans for collaboration with IoS projects and society impact, 10/03/11

Publications:

- Model Driven Reverse Engineering: Increasing Legacy Technology Independence, F.Barbier, G.Deltombe, O.Parisy, K.Youbi, proceedings of Second India Workshop on Reverse Engineering in The 4th India Software Engineering Conference, Thiruvanantpuram, India, CSI ed., February 23, 2011
- Reuse and Migration of Legacy Systems to Interoperable Cloud Services- The REMICS project, Mohagheghi Parastoo, Berre Arne Jørgen, Andrey Sadovykh, Franck Barbier, Gorka Benguria, In Proceedings of Mda4ServiceCloud'10 at ECMFA 2010, June 15 2010.

SocIos
Exploiting Social Networks for Building the Future Internet of Services

SocIoS views the democratization of the service lifecycle management as a major step towards the Future Internet of Services. The focal point of the SocIoS concept is the provision of the necessary tools to common Internet users that allow them to compose, provide and consume application services over abstract resources on the Web.
Social Networks appear to be the perfect "sandbox" for realizing this vision. Internet users inhabit these environments in masses, and the Social Networks themselves, apart from an environment for socialising, also provide tools for application development; thus, already allowing home users to build applications that have a business character and utilize –if not generate- the UCC and social graph.

SocIoS comes to cover the aspect of service application composition, provision and consumption using SN environments as containers of social graphs and UCC while at the same time supporting the business nature of service provision. Practically, SocIoS will achieve these by implementing:

- A SOA infrastructure that will act as a virtualization layer on top of social networking containers
- A toolset for supporting the business extensions of the developed applications; mainly through SLA support
- An API that will grant end users with single point access to the underlying functionality, including a set of core wrapper services
- A user interface that will simplify the use of the SocIoS API for easier application development and deployment by the home user
- A toolset for third-party services support so they can "blend" with the services developed and deployed using SocIoS, for more effectively delivering applications that exploit the UCC and social graph.

SocIoS will enable common users to access SNs and their applications as well as services created by other users through a single API and combine them in order to create a new application with enhanced functionality tailored to their needs. For SN platforms compatible with SocIoS (BounceIt® and platforms supported by OpenSocial®), it will be possible for the user to include the application in their SN profile. Apart from the common users, independent developers are also potential end-users of the API as it will provide them an easy way of accessing the various services available in the platform.
As the picture depicts, different artefacts are consumed by different end-user categories as SocIoS proposes the provision of a framework addressing the value chain from end-to-end.