



Resources and Services Virtualization without Barriers (RESERVOIR) Whitepaper

Abstract: RESERVOIR FP7 project is developing breakthrough system and service technologies that will serve as the infrastructure for Cloud Computing and Future Internet of Services by creative coupling of service virtualisation, grid computing, networking, and service management techniques.

The research leading to these results is partially supported by the European Community's Seventh Framework Programme ([FP7/2007-2013]) under grant agreement n° 215605.

1. Overview

Consider a small start-up company that just came up with a great idea for a new Web 2.0 application. They have tested in-house a small-scale version of their system, and now they feel ready to go public into a full Internet scale service. What are the possibilities for expansion available today to this company? Clearly they can raise the funds needed to build their own infrastructure, but how big should this infrastructure be? If they start small and plan to grow with demand they may miss the opportunity – in today’s competitive online market, responsiveness is one of the most important factors in determining success or failure of an online based service. On the other hand if they build a massive infrastructure right from the beginning this may turn out too big of a bet – if the service is only partially successful, the investment will not be returned.

Alternatively, this start-up can partner with a hosting company. The emerging *cloud computing* [1][2][3][4] paradigm for hosting web-based services, as best exemplified by the Amazon Elastic Compute Cloud (EC2) [5] or Google’s App Engine [6], aims to facilitate the creation of innovative internet scale services without worrying about the computational infrastructure needed to support these services. However, these new “cloud computing infrastructure providers” have a scalability problem of their own - what warranties can a single hosting company give to ensure that resources will always be available? In fact, no single hosting company can create a seemingly infinite infrastructure capable of serving increasing number of online-services, each with massive amounts of users and access at all times, from all locations.

To truly cater to the needs of service providers, it is inevitable that the compute cloud is going to be made of a federation of infrastructure providers. Only by partnering with each other, infrastructure providers can take advantage of the diversity factor and achieve the economies of scale needed to provide a seemingly infinite compute utility.

The high-level objective of the RESERVOIR project is to significantly increase the competitiveness of the European ICT industry through the introduction of a next-generation infrastructure for the deployment of complex services on a compute cloud that spans infrastructure providers and even geographies, while ensuring QoS and security guarantees. In doing so, RESERVOIR will provide a foundation for a service-based online economy where resources and services are transparently and flexibly provisioned and managed like utilities.

Cloud computing is just the latest incarnation of a concept that has been around since the 1960’s [7][8] - the emergence of a general-purpose public computing utility. Throughout the story of computing we have seen such utilities appear in one form or another and even though some success stories exist, in particular in the area of high performance scientific computing, where grid computing [9][10][11] made significant progress over the past decade, none of these attempts materialized into a general purpose compute utility that is accessible by anyone, at any time, from anywhere.

What makes cloud computing different is that recent developments in IT such as the appearance and fast adoption of virtualization technology for x86-based servers [12][13][14], as well as industry trends such as the adoption of Software as a Service [15] [16] as an alternative method for delivering functionality both to individuals as well as to companies, have finally created an opportunity/need for a global computing utility: On one hand the reluctance to use online services as a replacement for traditional software is


lessening – the success of companies such as salesforce.com [17] proves that with the right set of security warranties and a competitive price, companies are willing to trust even their most valuable data – customer relations – to an online service provider. On the other hand, virtualization has made it possible to decouple the functionality of a system as it is captured by the software stack (OS, middleware, application, configuration and data) from the physical computational resources on which it executes, this in turn enables a new model for online computing – instead of specially crafted online software we can now think in terms of general purpose online virtual machines that can do anything, i.e., thru virtualization the era of a general purpose compute utility is now within reach.

Essentially, the RESERVOIR project aims to support the emergence of Service-Oriented Computing [18] as a new computing paradigm. In this paradigm, services are software components exposed through network-accessible, platform and language independent interfaces, which enable the composition of complex distributed applications out of loosely coupled components. Services will be a fundamental element of the online economy and will underlie virtually all business and government activities such as electronic commerce within and across organizational boundaries, enterprise applications, telecommunication systems, energy and utilities, healthcare, travel and entertainment and more.

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

The “Service Oriented Infrastructure (SOI) Equation”

To accomplish the vision, our work will extend, combine and integrate three technologies: virtualisation, grid computing and business service management (BSM). We believe this approach can realize the vision of ubiquitous utility computing, by harnessing the synergy between the complementary strengths of these technologies. Graphically, this research strategy is represented by the “Service Oriented Infrastructure (SOI) Equation”:



Virtualisation and Grid. Virtualisation technology has shown to be useful in overcoming some barriers to commercial adoption of grid technology. However, to fully benefit from virtualisation, RESERVOIR will extend contemporary virtualisation infrastructures to be grid-aware, e.g., by supporting live migration across administrative domains. On the other hand, RESERVOIR will add virtualisation-awareness to the grids, for example, by using low-level monitoring information (available in virtualized systems) for metering and billing. We believe that these extensions will lay a foundation for a dynamic and eas-

easier-to-use service-oriented infrastructure, and that the resultant whole (SOI) is larger than the sum of its parts (virtualisation and grid technology).

Business Service Management (BSM). To fully benefit from the dynamic nature of the RESERVOIR infrastructure, we will develop a uniform policy-driven management layer that will automatically allocate resources to services and monitor execution and utilization to ensure compliance to Service Level Agreements (SLA) by adjusting resource allocation level and location. The new capabilities of the infrastructure will enable us to explore new allocation policies, optimizing over a range of parameters that is wider than what is commonly done today, e.g. the reduction of power consumption.

Having described the base technologies which we will build upon, we can now further elaborate on the novel principles comprising the RESERVOIR infrastructure. The following ideas, based on the spirit of the “SOI Equation”, are the innovative principles of our proposed next-generation SOI:

The Idea of Provisioning Services as Utilities

The vision of RESERVOIR is to enable the delivery of services on an on-demand basis, at competitive costs, and without requiring a large capital investment in infrastructure. Our research is inspired by a strong desire to liken the delivery of services to the delivery of utilities in the physical world. For example, a typical scenario in the physical world would be the ability of an electrical grid in one country to dynamically provide more electric power to a grid in a neighbouring country to meet a spike in demand.

We recognize that provisioning services from disparate service domains on-demand is arguably a more complex problem than the analogous problem of provisioning a utility in the physical world. For example, a service may be a composition of numerous distributed resources, including computing, storage and network elements. Provisioning a new service will consume physical resources, but cannot be allowed to cause an SLA violation of any other running application. Additionally, once initially provisioned, resources composing a service may be moved to other domains based on lower costs, performance or availability considerations. Our research will address those issues and minimize the barriers to delivering services as utilities.

Advanced End-to-End Support for Service-Oriented Computing

Service Oriented Computing (SOC) is a paradigm shift in the way software applications are designed and implemented to support business processes and users. SOC is built on the Service Oriented Architecture (SOA), which, while still emerging, is generally considered to implement (at least) the following principles:

- Services are reusable.
- Services are autonomous, loosely coupled, and platform independent.
- Services need to conform to the requirements of a service contract.
- Services can be combined to form larger, multi-tier solutions.
- Services can be described, published and discovered.

While SOA gives the general architecture to enable Service Oriented Computing, it does not address the fundamental issues required for an actual deployable solution, such as end-to-end security, service deployment, management and orchestration, service billing, and interpretation and monitoring of Service Level Agreement (SLA) conditions. RE-

SERVOIR, as an infrastructure project, will develop the technologies required to address these gaps, making the SOC paradigm a reality in the European economy.

Service and Resource Migration without Boundaries

The logical separation of a computing process and the physical environment on which it is hosted enables the *migration* of the process from one physical environment to another, without affecting the process. Today, this separation is achieved either (1) by limiting the knowledge the computing process has about the environment on which it is running, or (2) by imposing hard constraints on the configuration of the infrastructure.

The first approach, commonly used in today's scientific applications, is restricted to self-contained independent processes that can be easily checkpointed and restarted elsewhere, for example processes that analyze large amounts of data but rarely interact with other processes or users. This restriction renders this approach inappropriate for the type of commercial services which RESERVOIR aims to support.

In the second approach, widely used in today's commercial server virtualisation offerings, the configuration of all the physical resources in the infrastructure has to be identical. For example, the commercial VMotion product from VMware [14] only supports migration when the source and destination hypervisors are on the same subnet and have shared storage. Clearly, these configuration limitations make this approach inapplicable to a large geographically distributed infrastructure that spans administrative domains.

RESERVOIR plans to remove these limitations or *boundaries* and take virtualisation forward to the next level in order to allow the migration of resources across geographies and administrative domains, maximizing resource exploitation, and minimizing costs to the European citizen.

Federated Heterogeneous Infrastructure and Management

Commercial virtualisation systems typically offer non-standard management interfaces which are limited to their proprietary technologies and environments. In contrast, RESERVOIR will introduce an abstraction layer that will allow us to develop a set of high level management components that are not tied to any specific environment. To demonstrate the applicability of this generic management layer, we will use at least two different virtualisation technologies: Virtual Machines (VMs) and Virtual Java Service Containers (VJSCs). Furthermore, RESERVOIR will collaborate with standardization bodies to create standard interfaces that will enable interaction between distributed sites or grid environments, allowing the federation of infrastructures.

The Big Picture – RESERVOIR in the greater context

RESERVOIR research on virtualisation and management of services enables and unifies some of the emerging trends in Future Internet development as enabled by FIA (Future Internet Assembly) program (<http://www.fi-bled.eu/>)

- *Internet of Things*, where every electronic object will be an lively partaker in the network;
- *Internet of Services* (Service-oriented Internet), where applications and services actively reside and collaborate in the network, and data / context data becomes an dynamic object in the network; Services can be searched, identified and composed as utilities into business process components

- *Internet of Networks*, where connectivity, mobility, cooperation, mobility, service-awareness and ubiquity are depended upon.
- *Internet of Content*, where most of the contents are generated by end-users/consumers
- *Internet of Collaboration, Privacy and Anonymity*, where people and services ought to realize how to extend trust to others;
- *Self-aware Internet*, where different Internet systems and services shall manage their own life-cycle, operation and interactions without human intervention and without centralised governance.

RESERVOIR pays its part in the development of the NEXOF - Open Service Framework (<http://www.nessi-europe.com>). The Networked European Software and Services Initiative (NESSI) aims to create and evolve a unified research agenda, based on a multidisciplinary approach, for European research in Services and their foundations. This research agenda defines and promotes the wide adoption of technologies, strategies and deployment policies fostering new, open, industrial solutions and societal applications that enhance the safety, security and well-being of citizens.

At the time of writing this document the FP7 Call 1 projects are just at their starting point, nevertheless there is clear direction towards research on services and research on putting and operating services in relationship with their context of use. Some of the related FP7 projects are:

- EzWEB project is addressing the defining a reference architecture and implementation of an open platform supporting the retrieval, combination and utilization of front-end layer components in a next- generation, global Service-Oriented Architecture;
- MASTER project is addressing the provision of methodologies and infrastructures that facilitate the monitoring, enforcement, and audit of quantifiable indicators on the security of a business process, and that provide manageable assurance of the security levels, trust levels and regulatory compliance of highly dynamic service-oriented architectures in centralised, distributed (multi-domain), and outsourcing contexts;
- SLA@SOI project is addressing the definition of a holistic approach for the management of service level agreements (SLAs) and to implement an SLA management framework that can be easily integrated into a service-oriented infrastructure (SOI);
- SOA4ALL project is addressing the definition and implementation of a comprehensive framework and infrastructure that integrates four complimentary and revolutionary technical advances (Web principles, Web 2.0, Semantic Web, and Context management) into a coherent and domain independent service delivery platform.

In addition RESERVOIR is monitoring the developments in a number of research initiatives, which are exploring ways to overcome the ossification of the current Internet, including:

- OGF (Open Grid Forum - <http://www.ogf.org>) is an international research and standardisation community dedicated to accelerating grid adoption to enable business value and scientific discovery by providing an open forum for grid innovation and developing open standards for grid software interoperability.
- FIND (Future Internet Design: <http://www.nets-find.net>) is a new long-term initiative of the NSF (National Science Foundation) research program. FIND considers what the requirements should be for a global network of 15 years from now, and how it could build such a network based on clean-state approach. The philosophy of the program is to help conceive the future by momentarily letting go of the present - freeing our collective minds from the constraints of the current state.
- GENI (Global Environment for Network Innovation Program: <http://www.geni.net/>) is a research program addressing serious problems facing today Internet and the systems it supports: inadequate security, reliability, manageability and evolvability. GENI aims to provide a facility where large-scale experiments could be carried out, enabling proposals to be realistically evaluated. In so doing, GENI would transform the way science is done in this field.

The RESERVOIR Consortium

To ensure maximum impact, the RESERVOIR Consortium was established with a good representation of industry leaders, small companies, academia and a standards organization, as follows:

International Business Machines (IBM)

The IBM Haifa Research Lab (HRL) has been first set up as a Scientific Center in 1972. The lab has grown steadily, to today's size of over 500 researchers and engineers, working in areas such as systems management, virtualisation technologies, storage systems, verification technologies, problem determination, information retrieval, optimization technologies, and ICT services.

Telefónica Investigación y Desarrollo (TID)

Telefónica I+D is the innovation company of the Telefónica Group, which is one of the world's largest telecom companies. It has over 220.000 employees, and manages over 190 million customers. Telefónica I+D was formed in 1988, with the aim of strengthening the Group's competitiveness through technological innovation. It has facilities in Madrid, Barcelona, Valladolid, Granada and Huesca and new subsidiaries has been created recently in Brazil and Mexico. Telefónica I+D is 100% owned by Telefónica. It employs over 1225 persons, 93% holding a University degree. It is the most important private R&D company in Spain, in terms of both activities and resources. R&D activities focus on strengthening landline, mobile, multimedia and interactive telephone services, service and network management systems, operation and business support systems, and emerging ICT technologies.

Centre d'excellence en technologies de l'Information et de la communication (CETIC)

CETIC, and SME, is an applied research centre, focused on Information and Communication Technologies, created in 2001 as an initiative of the Polytechnic Faculty of Mons, the University of Namur, and the Catholic University of Louvain. CETIC positions itself as a R&D centre "serving the industry", being a connecting agent between academic research and enterprises, acting for technology transfer in software engineering and electronic systems. The research areas of CETIC include: (a) software engineering, with a focus on processes and quality, requirements engineering, formal methods, and IT systems reverse engineering); (b) distributed systems and grid computing, with a focus on high availability, mobility and collaboration; and (c) embedded communication systems with a focus on hardware-software co-design.

University College of London (UCL)

UCL is one of the premier universities in the United Kingdom, with a strong emphasis on research. It has some 5000 staff, 9000 undergraduate and 3000 postgraduate students covering all areas including Science, Engineering and Medicine. The RESERVOIR project activity will be carried out in the Departments of Electronic Engineering (UCL-EE) and Computer Science (UCL-CS).

Universidad Complutense de Madrid (UCM)

The Universidad Complutense de Madrid (<http://www.ucm.es>) is the largest and most renowned university in Spain, with over 700 hundred years of history. The research group involved in RESERVOIR is the Distributed Systems Architecture Research Group (<http://dsa-research.org>), which is part of the Computer Science Faculty. The main activity of the group is research and development of Grid and virtualization technology, its deployment to build computational infrastructures and its application to solve large-scale scientific, engineering, and business problems.

Elsag Datamat (ED)

Elsag Datamat (www.elsagdatamat.com) is the new business concern established to act as a centre of excellence for the design and production of systems, services and solutions in automation, security, transport, defence & space, information technology. Resulting from the integration of Elsag and Datamat in 2007, Elsag Datamat is a company in Finmeccanica group, one of the world's leading industrial conglomerates.

Sun Microsystems (Sun)

A singular vision -- "The Network Is the Computer" -- guides Sun in the development of technologies that power the world's most important markets. Sun's philosophy of sharing innovation and building communities is at the forefront of the next wave of computing: the Participation Age. Sun can be found in more than 100 countries and on the Web at <http://sun.com/>. Sun's contribution to RESERVOIR will be hosted by the Grid Engineering Centre in Regensburg, Germany. The engineering centre in Regensburg is part of Sun Microsystems GmbH, a Germany registered subsidiary of Sun Microsystems Inc.

Thales

Thales is a global electronics group serving professional markets in three areas: defense, aerospace and security. Its activities include prime contracting for large-scale programmes, complex system architecture, and the supply of equipment and systems and related support services. Thales engineers draw on a solid foundation of generic and dual civil/military technologies focused on real-time management and distribution of information.

Università della Svizzera Italiana (University of Lugano) (USI)

The University of Lugano has been founded in 1996 and is part of the Swiss University System that includes Federal and Cantonal Universities. It is composed of 4 Faculties: Economics, Communication Sciences, Architecture and Informatics. The University of Lugano employs more than 400 faculty members, and offers courses to more than 2,000 students coming from 30 different countries. It offers joint degrees the Università Bocconi (in Economics) and Politecnico di Milano (in Architecture), two prestigious Italian Universities.

Umeå University (Umeå)

Umeå University (Umeå) is the largest and oldest university in northern Sweden (27000 students, 3900 employees). The Grid computing research group at Umeå University includes researchers in two units: the High Performance Computing Center North (HPC2N) and the Department of Computing Science. HPC2N is a national center for scientific and parallel computing, providing computing and competence resources to the research community. This includes a set of major computing and storage resources, in-depth user support, a range of training activities, and an extensive network. The department of Computing Science is conducting research in a broad spectrum of topics, including Grid infrastructure.

SAP

SAP has grown to become the world's leading provider of e-business software solutions. With 12 million users, 96,400 installations, and more than 1,500 partners, SAP is the world's largest inter-enterprise software company and the world's third-largest independent software supplier, overall. SAP solutions help enterprises of all sizes around the world to improve customer relationships, enhance partner collaboration and create efficiencies across their supply chains and business operations. SAP employs over 39,300 people in more than 50 countries.

SAP Research is the technology research department of SAP and as an integral part of SAP's R&D activities, SAP Research is responsible for identifying, researching, understanding, developing and evaluating new and emerging technologies, processes and e-business solutions that strategically influence the future of SAP business applications.

University of Messina (UniMe)

The University of Messina was founded in 1548 and counts about 40000 students and 1600 teachers. The Centre on Information Technologies and Their Application (CIA) is a new structure where advanced ICT research is developed and applied. The main scientific activity is focused on studying distributed and parallel systems. The CIA is one of the GRID computing nodes running in south Italy. The centre has also recently established

two new research Labs: RFIDLab, jointly with ORACLE and INTEL on study and application of RFID technology, and The MobiLab, jointly with HP on mobility and open source.

OGF.eeig

OGF.eeig, which represents OGF in the RESERVOIR project, is the European Chapter of the Open Grid Forum (OGF) Economic Interest Group. It has been established as a European-based entity designed to extend the OGF model to Europe in alignment with OGF's mission, strategy and leadership to stimulate, co-ordinate and harmonize networked actions across Europe and globally.

More information

For more information on RESERVOIR you can visit our web site at www.RESERVOIR-fp7.eu

Alternatively, you may directly contact our project coordinator:

Eliot Salant

salant@il.ibm.com

References

- [1] *The Information Factories*, Wired Magazine, Issue 14.10, October 2006
http://www.wired.com/wired/archive/14.10/cloudware.html?pg=1&topic=cloudware&topic_set=
- [2] *Reflections on Cloud Computing*, Irving Wladawsky-Berger's Blog, March 2008,
<http://blog.irvingwb.com/blog/2008/03/reflections-on.html>
- [3] *Understanding Cloud Computing*, Wallis Paul, Keystones and Rivets, February 2008,
<http://www.keystonesandrivets.com/kar/2008/02/cloud-computing.html>
- [4] *The Big Switch – Rewiring the World from Edison To Google*, Nicholas Carr, published by W. W. Norton, January 2008
- [5] Amazon Elastic Compute Cloud (Amazon EC2) web site,
<http://www.amazon.com/gp/browse.html?node=201590011>
- [6] What is Google App Engine
<http://code.google.com/appengine/docs/whatisgoogleappengine.html>
- [7] *The MAC system : The computer utility approach*, R. M. Fano, in IEEE Spectrum, vol. 2, pp. 5644, January 1965.
- [8] *Time-sharing on computers*, R. M. Fano and F. J. Corbato, in Scientific American, September 1966.
- [9] *What is the Grid? A Three Point Checklist*, Ian Foster, 2002,
<http://www-fp.mcs.anl.gov/~foster/Articles/WhatIsTheGrid.pdf>
- [10] The Globus Toolkit
<http://www.globus.org/toolkit/>
- [11] The Open Grid Services Architecture
<http://www.globus.org/ogsa/>
- [12] *Server Virtualisation: Doing More with Less*, Leon Erlanger, Inforworld Report, September 11 2006,
http://www.inforworld.com/article/06/09/11/37FEvirtcaseserv_1.html
- [13] *Xen and the Art of Virtualisation*, P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt and A. Warfield, in Proceedings of the 19th ACM Symposium on Operating Systems Principles, Bolton Landing, NY, USA
- [14] The VMware Web Site
www.vmware.org
- [15] *Turning Software into a Service*, Mark Turner, David Budgen, Pearl Brereton, Computer, vol. 36, no. 10, pp. 38-44, Oct., 2003
- [16] *“The Different Faces of IT as Service”* I. Foster, S. Tuecke -
http://www.ggf.org/documents/Diff_Faces_foster.pdf
- [17] The salesforce.com web site
<http://www.salesforce.com/>
- [18] *Service-Oriented Computing – A Research Roadmap*, M. Papazoglou, P. Traverso, D. Schahram, F. Leymann, B. Kraemer, Dugstuhl Seminar on Service Oriented Computing, 2006, <http://drops.dagstuhl.de/opus/volltexte/2006/524/>.