

# EMOTIVE: THE BSC'S ENGINE FOR CLOUD SOLUTIONS

By Iñigo Goiri, Jordi Guitart, Mario Macias, Jordi Torres, Eduard Ayguadé, Jorge Ejarque, Raul Sirvent, Daniele Lezzi, Rosa M. Badia - Barcelona Supercomputing Center, Spain

Cloud computing, in which IT resources and services are accessed as utilities, has recently emerged as a promising paradigm for exploiting e-Infrastructures.

Cloud computing is strongly based on virtualization, allowing applications to be multiplexed onto a physical resource while isolated from other applications sharing that physical resource. This technology simplifies the management of e-Infrastructures, but also requires additional effort if users are to benefit from it. Cloud computing must hide its underlying complexity from users: the key is to provide users with a simple but functional interface for accessing IT resources “as a service”, while allowing providers to build cost-effective self-managed systems for transparently managing these resources. System developers should be also supported with simple tools that allow them to exploit the facilities of cloud infrastructures.

## ELASTIC MANAGEMENT OF TASKS IN VIRTUALIZED ENVIRONMENTS

The Barcelona Supercomputing Center (BSC) is contributing to the e-Infrastructures research community with the EMOTIVE (Elastic Management of Tasks in Virtualized Environments) middleware, mainly the result of work by the BREIN and SORMA projects.

EMOTIVE provides users with elastic fully customized virtual environments in which to execute services. Further, it simplifies the development of new middleware services for Cloud systems by supporting resource allocation and monitoring, data management, live migration, and checkpoints.

## LAYERED ARCHITECTURE FOR IMPROVED SUPPORT

EMOTIVE comprises several layers and provides users with basic tools for supporting tasks:

- the Virtualization Resource Management and Monitoring (VRMM) layer wraps each virtualized node and monitors its state
- within the VRMM layer, the Virtualization Manager (VtM) provides an application-specific virtual machine (VM) for each application; the application is given full control of its execution environment without introducing risk to the underlying system or other applications
- the VMs are created on demand, according to application requirements, and are consolidated in the provider's physical resources to optimise their use
- a Resource Monitor (RM) monitors task and resource status, including historical information.

EMOTIVE is currently based on Xen virtualization technology, with plans to support several virtualization solutions using the libvirt toolkit. It can manage local and remote virtual machines (i.e. those running in a provider's nodes or in a third-party providers' such as Amazon EC2). In this way, users can benefit from several Cloud implementations when running applications.

EMOTIVE also provides facilities for supporting resource management in Cloud environments. These include:

- VM placement and fine-grain dynamic resource distribution based on Service Level Agreements (SLAs),
- efficient live migration of VMs across provider nodes,
- a checkpoint/recovery system to resume task execution upon hardware failure (thus achieving fault-tolerance),
- ability to create additional VMs on external clouds when the local provider is overloaded, and
- data management services for supporting VM creation and the migration and checkpoint mechanisms, and also to allow users to provide input (i.e. data stage-in) and retrieve output (i.e. data stage-out).

Middleware services for Cloud systems can also be developed on top of EMOTIVE. For example, within the BREIN IST project, BSC has developed a Semantically-Enhanced Resource Allocator (SERA) that distributes resources using semantic information. Similarly, within the SORMA IST Project, BSC has developed an Economically-Enhanced Resource Manager (EERM) that manages resources using economic algorithms.

## PARTICIPATION IN NUBA

Future development of EMOTIVE will also be fostered through participation in NUBA, the “Normalized Usage of Business-oriented Architectures” strategic research program, funded by the Spanish Ministry of Industry, Tourism and Trade. NUBA aims to advance state-of-the-art in business models and technology to deploy a federated Cloud platform that will simplify the deployment of business services.

## OUTCOMES OF THE BELIEF WORKSHOP ON DATA PROVENANCE

*By Eleni Toli, Michael Pantazoglou and Dimitra Keramida, National Kapodistrian University of Athens, Greece*

A hot topic for scientific communities is the challenge of dealing with massive volumes of digital research objects.

For such data to be of long term worth, they must be traceable in terms of provenance (data origin) and authenticity (data validity). Achieving this would also extend return on e-Infrastructure investment and increase competitiveness. However, managing provenance and authenticity data also poses significant challenges in terms of formal modelling, storage, and maintenance.

### THE ATHENS CORE

The 5th BELIEF brainstorming workshop, held in Athens, April 2009, aimed to tackle these issues, furthering discussions started at the 6th eConcertation Meeting in Lyon, November 2008. The Athens workshop led to several tangible results and recommendations, including the need to formally define a minimal set of provenance information and to develop a roadmap towards data provenance. Together, these actions form what is called the “Athens Core”.

The need for high-level definition of a roadmap towards data provenance is part of this “core”; the roadmap would involve Communicable Information Packages (CIPs) and their definition, implementation, testing, real-world application, and so on.

The “Athens Core”, if further elaborated and accepted, could act as a cornerstone for future research and standardisation activities in data provenance and authenticity.

Other important technical outcomes included the need for:

- Abstract definition of data provenance information as a “relationship”, with a graph model
- Layered viewing of provenance information to reflect their granularity and differences in stakeholder perspectives

