

GaaS: toward a more “elastic” and sustainable grid environment

Description of the work

The GaaS model combines the advantage of providing users with an usage model that is familiar to the traditional Grid, with the possibility of flexible management of computational resources in a IaaS-like manner. Hence, our model can be classified as a Platform-as-a-Service for extending Grid environments with elastic (i.e., virtual) resources. By using GaaS, users can define new Grid Sites, add computational resources to existing Grid Sites and modify the resources aggregation scheme, i.e., site queues. All GaaS services may be requested by users, with certain roles, from the UI. GaaS actually consists of four services to (a) add new computing resources (WN service: GaaS_WNS); (b) aggregate existing computing resources in a new queue (queue service: GaaS_QS); (c) add a new grid site for an existing VO (grid site service: GaaS_GSS); (d) create a suited runtime environment for a set of applications on existing or new computing resources (application environment service: GaaS_AES). The GaaS prototype, deployed on the S.Co.P.E. infrastructure at University of Naples, is based on the gLite-EMI middleware and on the OpenNebula cloud management system with Xen hypervisor to create virtual machines (VM) that host WNs, CE and IS. The main efforts in the prototype development were (i) the definition of templates for gLite-EMI services configuration, and (ii) the enabling of their fast provisioning. We designed our VM disk provisioning system in order to provide fast VM creation and avoid as much data copy as possible. Our solution is based on the GNU/Linux's Logical Volume Manager (LVM) system and a set of grid roles configuration templates. This approach allows a really fast creation of snapshots starting from a reference LV and the later configuration of the VM with the requested grid role.

Wider impact of this work

GaaS flexibility provides several advantages to traditional Grid infrastructures in different scenarios: WNs can be customized with software tailored to a given set of users, as well as queues can be configured to fulfill a specific computation needs. Moreover, GaaS support the creation of complete Grid sites in order to, i.e., enable a community that has to share resources for the life time of a project, to avoid the burden of configuring from scratch all the required services and resources. Finally it must be taken in account that GaaS, being an on-demand solution, can be also useful to optimize the energy efficiency of large scale systems reducing their overall operational cost and improving their sustainability.

Printable Summary

The work on GaaS (Grid as a Service) started two years ago in order to find a solution to make distributed computing environments, based on Grid model, more “elastic” and more sustainable by exploiting the features of the Cloud model. The Grid model is basically static and users cannot define new Grid sites, add computational resources to existing ones and modify the resources aggregation scheme in accordance to their needs; it is also not possible to dynamically modify the resources number on the basis of the real system workload, in the name of saving energy and to achieve a more efficient and sustainable environment. During the poster session we will describe the steps taken to design and to deploy GaaS prototype on the S.Co.P.E. production infrastructure at University of Naples.

The presentation is aimed at both scientists who use and technicians who manage grid environments because they will gain an understanding of how this work can be re-used for their own.

Primary author: Mr BARONE, G.B. (Università degli Studi di Napoli Federico II)

Co-authors: Mr BOTTALICO, D (Università degli Studi di Napoli Federico II); Mr CARRACCIUOLO, L (Università degli Studi di Napoli Federico II); Mr BIFULCO, Roberto (University of Naples Federico II); Dr BOCCIA, Vania (Università degli Studi di Napoli Federico II)

Track Classification: Virtualised Resources: challenges and opportunities (Michel Drescher: track leader)