



Contribution ID: 3

Type: **not specified**

Building a Grid-of-Clouds, Or: How One HEP Experiment Is Evaluating Strategies to Incorporate "The Cloud" into the Existing Grid Infrastructures

Wednesday, 28 March 2012 11:00 (20 minutes)

Description of the Work

In mid-2011 the ATLAS experiment formed a Virtualization and Cloud Computing R&D project to evaluate the new capabilities offered by these software and standards (e.g. Xen, KVM, EC2, S3, OCCl) and to evaluate which of the existing grid workflows can best make use of them; this effort is being coordinated by CERN IT Experiment Support. In parallel, many existing grid sites have begun internal evaluations of cloud technologies (such as Open Nebula or OpenStack) to reorganize the internal management of their computing resources. In both cases, the usage of standards in common with commercial resource providers (e.g. Amazon, RackSpace) would enable an elastic adaptation of the amount of computing resources provided in relation to the varying user demand.

In the topic of workload management, we have evaluated a few strategies to add cloud-based sites to the ATLAS PanDA workload management system (WMS). In particular, we have developed a lightweight "cloud factory" service which manages deployed VM instances and can be used by central grid operators for central production or by individual users to perform urgent data analyses on (chargeable) cloud resources. We present results of running sample analyses on virtualized/cloud resources at CERN and in StratusLab. Next, in the topic of cloud storage access and management we will present tests of remote XROOTD access to input data over the WAN, movement and management of EC2-resident data, and strategies to instantly deploy cloud-resident storage elements.

Conclusions

The Virtualization and Cloud Computing R&D project in the ATLAS experiment at CERN is evaluating techniques to incorporate these technologies to the existing grid infrastructure. This work will present the current status of this project in relation to the workload and data management services used by this experiment.

Impact

Virtualisation and Cloud Computing bring new features to homogenize the infrastructure layer and improve resource scalability. For the HEP VOs, elastic scaling of resources could be employed to better match provisioned resources to the dynamic demand, thereby decreasing costs and improving the user experience.

We also recall that one of the great strengths of grid computing is that it enables computing resource funding to be spent locally (to the benefit of the local economy) while providing the technology to pool global computing facilities to solve the Grand Challenges in computing. Cloud Computing does not sacrifice that strength. Indeed, sites are beginning to make their local resources available via a cloud API such as EC2, enabling both local and remote users to use the facilities using an API that is shared in common with Industry. By making it

easy to target applications at both traditional “academic” resources and new commercial computing centres, the users can flexibly adapt according to budgetary and urgency constraints.

Overview (For the conference guide)

Emerging standards and software often marketed as “Cloud Computing” bring attractive features to improve the operations and elasticity of scientific distributed computing. At the same time, the existing European Grid Infrastructure and Worldwide LHC Computing Grid (WLCG) have been highly customized over the past decade or more to the needs of the VOs and are operating with remarkable success. It is therefore interesting not to replace The Grid with The Cloud, but rather to consider strategies to integrate cloud resources, both commercial and academic, into the existing grid infrastructures, thereby forming a Grid-of-Clouds. This work will present the efforts underway in the CERN IT Experiment Support Group along with the ATLAS Experiment to adapt existing grid workload and storage management services to cloud computing technologies.

Primary authors: Dr VAN DER STER, Daniel (CERN IT-ES); Mr BARREIRO MEGINO, Fernando (CERN IT-ES)

Presenter: Dr VAN DER STER, Daniel (CERN IT-ES)

Session Classification: Clouds: Users

Track Classification: Operational services and infrastructure