



GridPP
UK Computing for Particle Physics

GridPP Project Management Board

Cloud Computing in GridPP4

Document identifier :	GridPP-PMB-146- CloudComputing.doc
Date:	28/01/2010
Version:	1.0
Document status:	Final
Author	PMB

Introduction

At the last Oversight Committee meeting we were asked to address the issue of cloud computing (in the context of GridPP and its future development). We continue to track developments in this area with “The Future of Cloud Computing” conference held in Brussels on 26th January 2010 being the most recent event¹. The status is that industry and government are moving to embrace clouds at a high-level. Clouds build on the developments of Grids but there are notable differences between grids and clouds in the eyes of the EC and the expert group². With grid computing, computing resources can be provisioned as a utility that can be turned on or off. Cloud computing goes one step further with on-demand resource provisioning. US vendors have taken the lead in cloud computing but the market is expected to be large and is still young. Many technology issues need to be addressed and hence there are opportunities for new vendors that can overcome them. Clouds are also challenging the legal frameworks, as Grids have done in the past. Governments are investigating their own cloud computing strategy³, but they should not implement regulation too early since there are still issues to be addressed but a clear international legal framework will be necessary for cloud services. An expert group has been formed and Cloud Computing R&D development will be supported by the future EC work programmes (2011-2013). EC sees a global approach to cloud computing as important and wants to do it together with other regions (e.g. USA).

GRIDPP Context

This is an area of great interest to wLCG in an international context and many studies already exist by our collaborators which show, at present, that the cost of data storage on the cloud and the cost of getting data in/out of the cloud compromises their usefulness. Current surveys from experiments presenting at the CHEP09 Conference indicate that such a solution is technically viable for some limited applications (e.g. “Belle on the Amazon EC2 Cloud”⁴) however the full costs and the overall ability to anticipate and meet peak demand have not been established in the (UK) market. In particular, the economics dictates that great care must be taken to minimise the amount of time data resides on the cloud, yet the vendor control of bandwidth into and out of the cloud creates a form of vendor lock-in that makes the cost unpredictable. Thus, the lack of transparency in the overall vendor pricing/business model is a key concern⁵.

In parallel with commercial offerings there are various developments within our collaboration of universities and associated spin-out companies. Examples include Constellation Technologies (RAL) who collaborate with us on various projects and were a pioneer in global grid infrastructure software leading to the development of cloud computing tools based on open source software used by the particle physics community⁶; In association with this work they have completed development of SuperCloud MT, the management tool for their SuperCloud platform which is currently being tested on the Cambridge University Grid; Oxford running a cloud interface to their internal resources and looking at Grid interfaces for external access; Edinburgh instantiating cloud instances of glite components for their grid training; the HEPiX Group on virtualisation planning for users or experiments sending images to run at sites. Once a market cost has been established these may be in a position to provide tools to run an on-demand cloud computing service at commercial cost.

Following internal discussion, it was decided that the GridPP4 proposal should include a request for some seed money to investigate commercial cloud computing offerings in the UK context. Here, costs and constraints may differ from the existing European and US studies and, in any case, the financial landscape will have changed in two years’ time. The plan, contained in the GridPP4 proposal, is to request £10k as seed money to perform a trial producing Monte-Carlo simulations (by far the simplest of our computing tasks) on a commercial cloud and to reserve £100k of the Tier-2 hardware budget for

¹ See <http://cordis.europa.eu/fp7/ict/ssai/docs/20100126-agendaforweb5.pdf>.

² A useful summary is provided at <http://www.ibm.com/developerworks/web/library/wa-cloudgrid/>.

³ See <http://www.guardian.co.uk/technology/2010/jan/27/cloud-computing-government-uk>.

⁴ See http://www.hpcadvisorycouncil.com/pdf/vendor_content/VPAC_cloud2009.pdf.

⁵ See http://www.theregister.co.uk/2009/04/15/mckinsey_cloud_report/.

⁶ See <http://www.constellationtechnologies.com/>.

competitive tender later in the project in order to smooth out peaks in demands for simulated data. It was anticipated that, once the real cost/convenience of commercial offerings had been understood from the trial, then Tier-2s could also compete for the balance of the funding on a commercial basis. The proposal additionally requires some Tier-2 effort (contained with the overall proposal) to be allocated in order to system manage the queue.

The proposal below enables a more detailed evaluation to be undertaken and is to be considered in the context of the overall GridPP4 proposal. In addition to this undertaking we also plan to collaborate with Dr Will Venters (LSE), who is reviewing cloud computing in the wider context (from a social science perspective) so that we may benefit from an external view. We plan to collaborate with him on these wider aspects prior to implementation of such work.

Proposal

The aim of this Cloud Computing Tier-2 Work Package is to allow peak compute demand to be met prior to e.g. summer conference submission for a period of typically a few months in ~April-June by offloading the least demanding compute use case (Monte Carlo Production) to a cloud. It is difficult to estimate this peak load beyond planned capacity, but we assume this could be as large as an additional 50% load for periods up to a month: the following step-wise assessment will be used to determine this load and our response to it.

1. In April 2011 we will establish a £10k pilot survey initiated in association with one or more of the large LHC experiments. This will enable their Monte Carlo production system to be submitted to commercial vendors such as the Amazon Cloud, or to universities with pre-existing cloud computing infrastructures, using test samples. The cloud job turnaround time including the return of data via commercial networks will be measured. The output will be validated against a similar submission to an existing Tier-2 site. The full cost assessment will include those of the commercial vendor(s) as well as any additional UK-based manpower to submit, manage, maintain and support such a submission queue for use during busy periods.⁷
2. Assuming a successful outcome of the pilot survey, a further £100k will be set aside over the next 3 years to respond to peak demand of the experiments: requests for spend will be brought directly to the weekly PMB (where overall demand can be assessed).

Resources

1. £10k additional funding for a pilot survey in FY 11/12.
2. £100k to be reserved from the Monte Carlo Production budget from FY12-14.
3. Existing manpower at a chosen Tier-2 site to establish and maintain such a system (estimated at ½ FTE for 3 months of each year i.e. 1/8 FTE/year).

⁷ This pilot survey is predicated on the Monte Carlos being reasonably stable and a virtual environment being created that meets the requirements of the experiment(s). This can then be packaged, incorporating a set of web services, for submission to the cloud.