

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

Daniel Colin van der Ster¹

E-mail: daniel.colin.vanderster@cern.ch

Fernando Harald Barreiro Megino

E-mail: fernando.harald.barreiro.megino@cern.ch

Ian Gable

E-mail: igable@uvic.ca

Michael Paterson

E-mail: mhp@uvic.ca

Ryan Taylor

E-mail: rptaylor@uvic.ca

Rodney Walker

E-mail: Rodney.Walker@physik.uni-muenchen.de

On behalf of the ATLAS Collaboration

Abstract

The ATLAS Computing Model was designed around the concept of grid computing and since the start of data taking, this model has proven very successful. However, new paradigms in computing, namely virtualization and Cloud Computing, present improved strategies for managing and provisioning IT resources that could allow ATLAS to more flexibly adapt and scale its storage and processing workloads on varied underlying resources. This work will present the current status and use-cases of the Cloud Computing R&D project in ATLAS Distributed Computing and the initial results found while running ATLAS simulation and analysis jobs in cloud resources.

*EGI Community Forum 2012 / EMI Second Technical Conference,
Munich, Germany
26-30 March, 2012*

¹ Speaker

1. Introduction

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 give an overview of the existing Cloud panorama and present the efforts underway in the CERN IT Experiment Support group along with the ATLAS experiment [1] to adapt existing grid workload and storage management services to Cloud Computing technologies. In the following sections we will explain the different foreseen use cases and will focus on the results of workload management activities in the cloud and the orchestrators developed so far to provide dynamic provisioning of cloud resources.

2. Related Cloud Computing panorama

Cloud Computing [2] is the delivery of computing and storage capacity as a service to a community of end-recipients. Cloud computing entrusts services with a user's data, software and computation over a network. Although there are multiple service models, in the context of this paper we will refer to Cloud Computing as the Infrastructure as a Service (IaaS) model, where the providers rent computing, storage and network capacity to their customers.

Amazon played a key role in the development of Cloud Computing, since this company was using as little as 10% of their capacity in order to leave room for occasional spikes. Therefore Amazon was interested in renting out the unused capacity to external customers, which resulted in the launching of the Amazon Web Services in 2006 and since then several other providers have joined the Cloud Computing evolution. Given Amazon's leading market position, the Amazon EC2 [3] and S3 [4] models and interfaces are considered by some as the "de facto" standards for computing and storage respectively.

However, as Cloud Computing is a very recent technology emerging providers wish to offer services and models that distinguish them from the rest. Agreed standards and interfaces have not settled yet, but different working groups, such as the Open Cloud Computing Interface (OCCI) community [5] are working in order to provide interoperable and portable cloud interfaces.

In Europe, the Helix Nebula Science Cloud [6] is an initiative that aims to set up a Cloud Computing infrastructure for the European Research Area and identify and adopt policies for trust, security and privacy on a European-level (see general timeline in Figure 1). A consortium of leading IT providers (e.g. ATOS, CloudSigma, T-Systems), research institutes (CERN, EMBL, ESA) and existing projects (e.g. the EGI Federated Cloud Task Force [7]) form the Helix Nebula partnership and aim to build a federation of interoperable clouds.

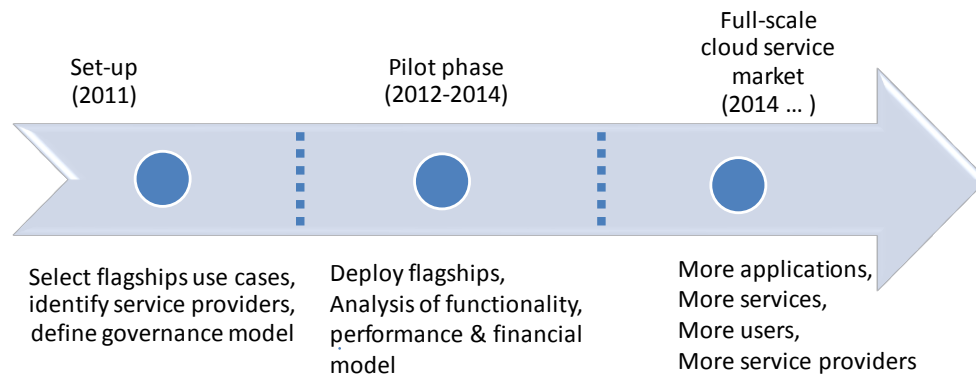


Figure 1 - Timeline of the Helix Nebula Science Cloud project

In the scope of CERN's participation, one important task of the Helix Nebula project is to push the integration of Grid and Cloud Computing. This is also one of the main purposes of the EGI Federated Clouds Task Force, since many of EGI's current and new user communities would like to access the flexibility provided by virtualization across the infrastructure on demand.

3.Cloud Computing scenarios for ATLAS

In mid-2011 the ATLAS experiment formed a Virtualization and Cloud Computing R&D [8] project to evaluate the new capabilities offered by these standards and to evaluate which of the existing grid workflows can best make use of them; this effort is being coordinated by the CERN IT Experiment Support group. In parallel, many existing grid sites have begun 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.

3.1 Data processing and workload management

In the area of Data processing and workload management, the ultimate goal is to understand how to run jobs on cloud CPUs in a stable way and test different data access strategies. We distinguish following use cases:

- i. **PanDA (ATLAS Production and Distributed Analysis framework [9]) queues in the cloud:** The idea is that usual grid sites make themselves available through cloud interfaces and are centrally managed. The deployment is not trivial but scalable and benefits both ATLAS and the sites, while it is transparent to users.
- ii. **Tier3 analysis clusters:** For local analysis, Tier3 institutes can manage internally their resources in the cloud way. The deployment should have low to medium complexity.
- iii. **Personal analysis queue:** User managed group of Virtual Machines (VMs) on a cloud provider with low complexity setup to run urgent analysis.

3.2 Data management and storage

In the area of data management, we see the following two use cases:

- i. Short term data caching to accelerate the above data processing use cases. Aggregating the disks that come with the VMs, we can build a data cache (e.g. with XRootD [10] or the Hadoop File System (HDFS) [11]) and store data in a transient way so that it can be reused.
- ii. Long-term data storage in the cloud by integrating cloud storage solutions into the ATLAS Distributed Data Management system [12].

Activities on the data management and storage part have just started and there is less experience than on the workload management part. Therefore this contribution will focus on the ATLAS data processing and workload.

4. Data processing and workload management on the cloud

4.1 PanDA workflow

PanDA is the framework for the distributed analysis and production of ATLAS data on the grid. Figure 2 shows the basic workflow and how we can attach any cloud resource to a Condor [13] pool and therefore integrate cloud resources into the PanDA framework.

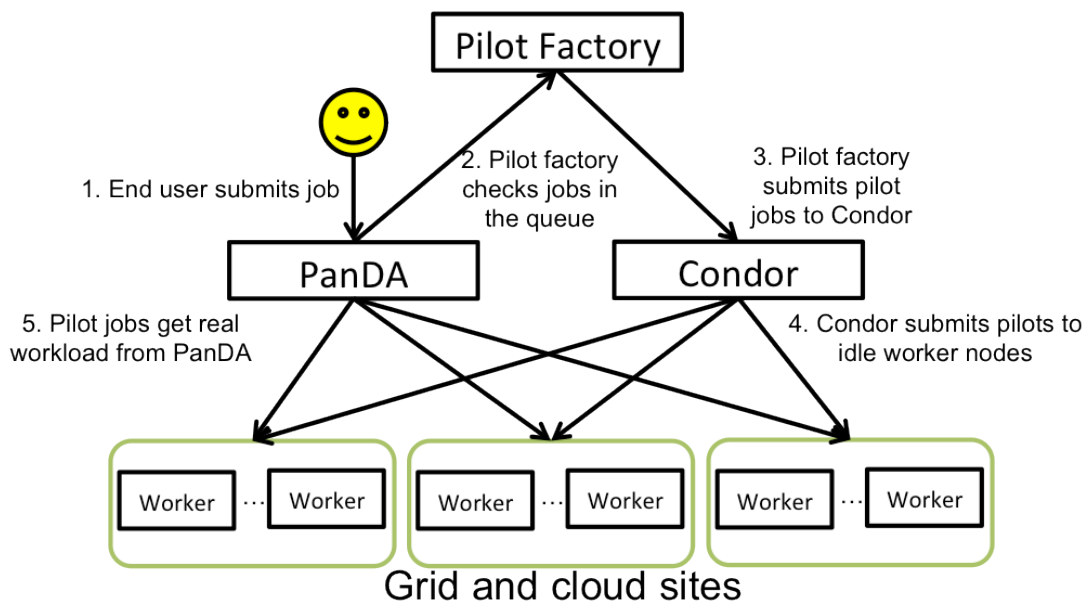


Figure 2 - ATLAS distributed analysis workflow

4.2 Running production jobs in the cloud

During the pilot phase of the Helix Nebula project introduced in Section 2, three laboratories (CERN, EMBL and ESA) will evaluate Cloud Computing resources from different European IT providers. ATLAS is committed to run real Monte Carlo simulation jobs for a few weeks on each commercial provider.

The first provider, Cloud Sigma AG, has been successfully tested with the following basic idea:

1. Upload the CernVM [14] image to Cloud Sigma. This image has the necessary software preinstalled
2. Boot a VM and configure it manually to join a Condor pool with the master (one of the pilot factories) at CERN.
3. Clone more instances of the VM
4. Create a new PanDA site *HELIX*
 - a. The site is in *brokeroff* status so that jobs are not assigned automatically to it
 - b. Real Monte Carlo simulation tasks are assigned manually. These tasks have a very high CPU consumption, while they have almost no I/O requirements.
5. The data input and output is copied over the WAN from the CERN Storage Element (*lcg-cp/lcg-cr* for input/output)

As a result we ran simulation tasks on 100 VMs (200 CPUs) and were able to finish 6 x 1000 job tasks in around 2 weeks. We ran one identical task at CERN to get reference numbers. The results are summarized in Table 1.

Table 1 - Production statistics summary

	<i>HELIX</i>	<i>CERN</i>
Failure rate	265 failures 6000 succeeded 4% failure rate	36 failures 1000 succeeded 3.5% failure rate
Job running time	16267s +- 7038s	8137s +- 765s

It is noticeable that there is a high variance in the job running time at HELIX, as we do not know what else is running on the bare metal our VMs are shared on. In respect of the mean running time, it is not fair to compare the running times at HELIX and CERN since there are different CPUs behind. The best comparison would be to have an estimation of CHF/event, which is presently unknown.

Similar production jobs have also been ran at a smaller scale with success on the Stratuslab [15] test cloud in GRNet [16].

4.3 Orchestrators

In the gaps between tasks we are being billed for idle CPUs and therefore we need orchestrators that take care of the automatic provisioning and grow or shrink the size of the cluster based on the existing demand. In this topic there are two products: the centrally managed Cloud Scheduler and the personally managed Cloud Factory.

4.3.1 Cloud Scheduler

Cloud Scheduler [8] is a Python package that boots VMs on Infrastructure-as-a-Service clouds and is based on Condor queues. The users create their usual batch jobs and submit them to e.g. PanDA and specify which VM type should run their images. The Cloud Scheduler will watch the depth of the queue and boot up/shut down VMs accordingly.

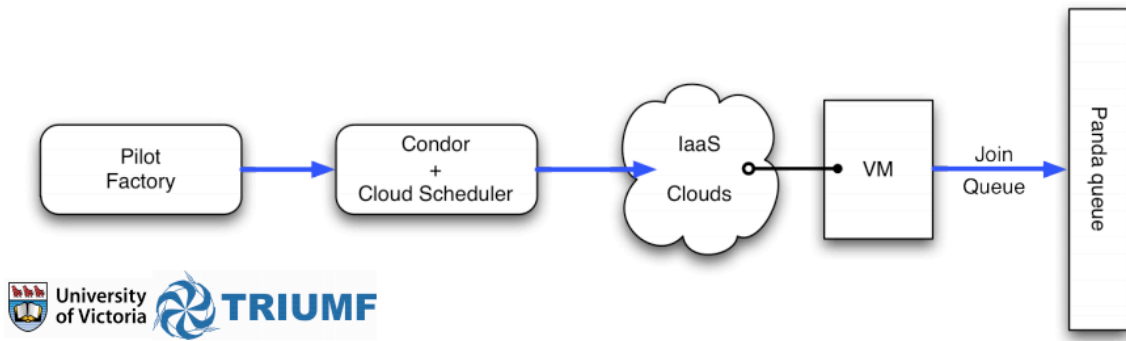


Figure 3 - Cloud Scheduler workflow

The Cloud Scheduler works with multiple cloud sites simultaneously and supports different cloud APIs. It is based on CernVM and relies on Condor for the credential delegation at job submission time. It keeps the VMs alive until there are no more jobs needing that VM type.

The Cloud Scheduler has been successfully used to run 1500 analysis jobs over 1 day. The worker nodes were part of FutureGrid (Chicago), which were reading the data from University of Victoria. Cloud Scheduler automatically manages 100 VMs. As opposed to the simulation jobs, these analysis jobs do have significant I/O requirements and during the exercise network traffics of up to 3.5 Gb/s were seen between Victoria and Chicago (see Figure 4).

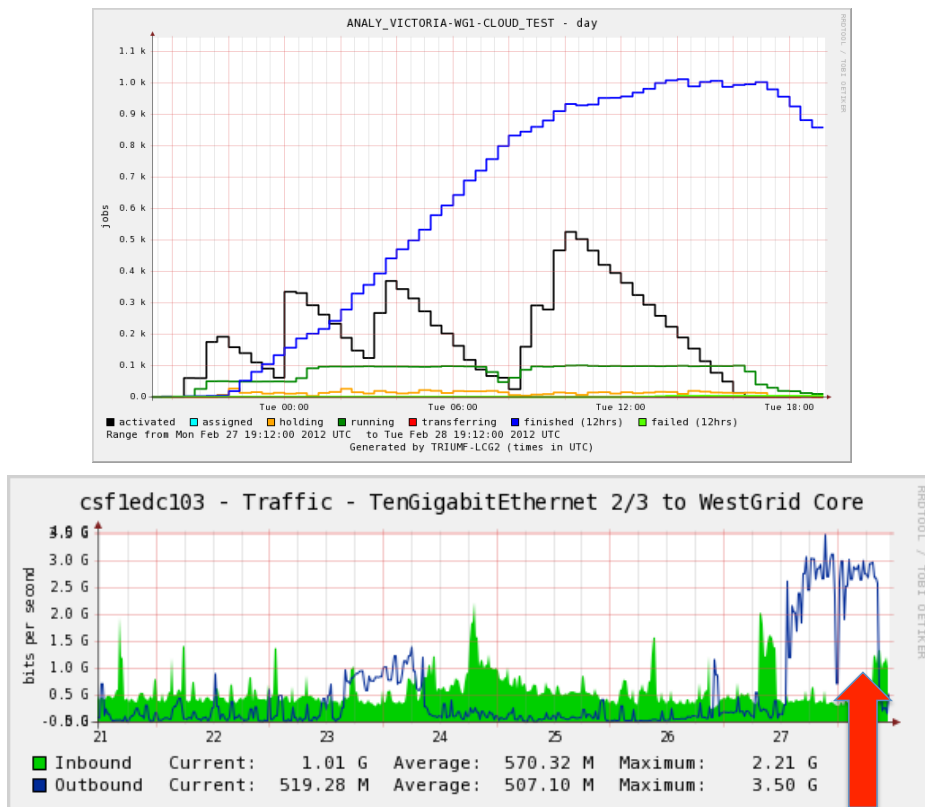


Figure 4 - VM provisioning by the Cloud Scheduler and the generated traffic

POS (EGICF12-EMITC2) 005

4.3.2 Cloud Factory

The other available orchestrator is the Cloud Factory [8], which is a proof of concept, personal orchestrator tested at CERN. The main difference is the simplicity to manage it: the Cloud Factory does not require the VMs to attach to any Condor pool. It gets the personal jobs from the PanDA queue and instantiates CernVM nodes to run them. For the moment it only speaks EC2. The Cloud Factory uses the Grid Messaging Service (MSG) [17] to let the VMs communicate their status and a heartbeat to the Cloud Factory.

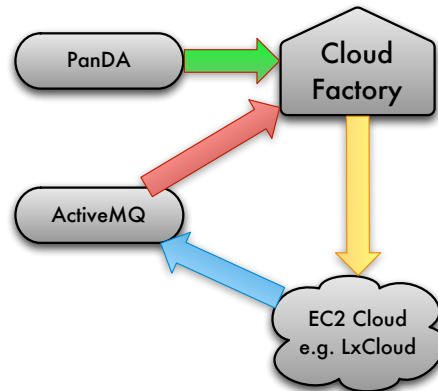


Figure 5 - Cloud Factory workflow

5. Conclusions

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, 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.

However Cloud Computing is at an early stage and it would be beneficial if the Cloud Computing industry and technology providers agree on a standards-based open interface layer that allows common ways to provision, access and account resources across suppliers.

The Virtualization and Cloud Computing R&D project in the ATLAS experiment is evaluating techniques to incorporate these technologies to the existing grid infrastructure. During the first year it has been successfully demonstrated that it is possible to run ATLAS payloads in the cloud. Given the lack of standards, the activity is still labour intensive and a lot of human effort has to be put in the configuration and preparation of the VMs for each cloud provider. Many of the activities are maturing and in the medium term we should determine what

we could deliver in production by automating the deployment procedures and improving the node monitoring.

It must be noted that the current set ups are not necessarily the final ones that we would like to use in future: we should understand the models of cloud storage that can allow us to make use of different storage implementations in the cloud. Some institutes are starting to provide S3 storage endpoints and the effective R&D will start soon. As well, we are looking forward to more progress in data caching in the cloud using XRootD or HDFS. This research should allow bringing the data closer to the worker nodes and avoid the use of remote data I/O across the WAN.

In addition, there are more topics that still need to be addressed, as can be security issues, privacy concerns, operational procedures or CPU accounting on the cloud.

References

- [1] The ATLAS Collaboration, The ATLAS Experiment at the CERN Large Hadron Collider, 2008 JINST 3 S08003 doi:10.1088/1748-0221/3/08/S08003
- [2] http://en.wikipedia.org/wiki/Cloud_computing
- [3] <http://aws.amazon.com/ec2/>
- [4] <http://aws.amazon.com/s3/>
- [5] <http://occi-wg.org/>
- [6] <http://helix-nebula.eu/>
- [7] <https://wiki.egi.eu/wiki/Fedcloud-tf:FederatedCloudsTaskForce>
- [8] Fernando Barreiro Megino et al, *Exploiting Virtualization & Cloud Computing in ATLAS*, CHEP, New York, Canada, May 21-25, 2012
- [9] Tadashi Maeno et al, *PanDA: distributed production and distributed analysis system for ATLAS*, CHEP, Victoria, Canada, September 02-07, 2007
- [10] Alvise Dorigo et al.; XROOTD- A highly scalable architecture for data access. WSEAS Transactions on Computers, 1(4.3), 2005.
- [11] <http://hadoop.apache.org/hdfs/>
- [12] M Branco et al. Managing ATLAS data on a petabyte-scale with DQ2. 2008 J. Phys.: Conf. Ser. 119 062017 doi:10.1088/1742-6596/119/6/062017.
- [13] Jim Basney, Miron Livny, and Todd Tannenbaum, "High Throughput Computing with Condor", *HPCU news*, Volume 1(2), June 1997.
- [14] P Buncic et al, *CernVM - a virtual appliance for LHC applications*, Advanced Computing and Analysis Techniques in Physics Research, Erice, Italy, November 03-07, 2008
- [15] <http://stratuslab.eu/index.php>
- [16] <http://www.grnet.gr/default.asp?pid=1&la=2>

[17] <https://tomtools.cern.ch/confluence/display/MIG/Messaging>

POS (EGICF12-EMITC2) 005