

# Cloud Computing R&D Proposal

**F. H. Barreiro Megino, K. De.  
D. van der Ster, R. Walker  
*on behalf of ATLAS ADC***

ATLAS Software and Computing  
Workshop, 5 April 2011

- Introduction
- Goals and scope of the project
- Project Roadmap
- Proposed development areas
- Conclusions

This proposal is meant to start the discussion, and will evolve rapidly as the group starts working (after all this is R&D project)

- ATLAS Computing Model designed around grid computing with >100 WLCG sites
  - Over a year in full operation coping with data-taking requirements
- Cloud Computing:
  - efficiently share bare hardware resources without sacrificing flexibility in services offered to applications
  - new and increasingly competitive market has emerged to offer these cost effective computing resources
- It is of interesting for ATLAS Distributed Computing to investigate opportunities for adapting its existing services to make use of cloud computing resources.
  - Design and implement cloud-awareness in the Distributed Data Management (DDM) system, the Production and Distributed Analysis (PanDA) system, and in related tools and services.

- Project description is available now:
  - “Proposal for Cloud Computing R&D in ATLAS Distributed Computing”

[https://twiki.cern.ch/twiki/bin/viewauth/Atlas/TaskForcesAndRnD#Cloud\\_Computing\\_RD](https://twiki.cern.ch/twiki/bin/viewauth/Atlas/TaskForcesAndRnD#Cloud_Computing_RD)
- This talk summarizes the proposal document.

- Primary goals:
  1. to evaluate available cloud technologies in relation to the use-cases presented by ATLAS data management, processing and analysis
  2. to design a model for transparently integrating cloud computing resources with the ADC software and service stack
  3. to implement the ATLAS cloud computing model into DDM, PanDA and related tools and services.
- Collaborate with related groups in WLCG, CERN-IT and other LHC experiments:
  - Consolidate efforts and leverage commonalities
  - Profit from existing experiences

- Potential ADC use-cases:
  - Monte Carlo simulation on the cloud with stage-out to traditional grid storage or long-term storage on the cloud
  - Data reprocessing in the cloud (with a strong caveat related to cost; see next slide)
  - Distributed analysis on the cloud using data which is accessed remotely to the grid sites, or analysis of data which is located in the cloud
  - Resource capacity bursting which is managed centrally (e.g. to handle urgent reprocessing tasks) or regionally (e.g. to handle urgent local analysis requests)

In general, all ATLAS use cases should be possible, making cloud computing sites equivalent to grid sites

- Need to pay special attention to the cost of cloud resources:
  - Some will be provided by our collaborators and other resources will be commercial
- “Free” cloud resources at academic and research institutions:
  - A cloud API (e.g. OpenStack, etc...) would be an alternative to grid middleware and enables sites to better manage local resources
  - ADC should be able to use a WLCG site that makes resources available via a cloud API
- Commercial cloud resources:
  - These resources (e.g. Amazon, Rackspace, etc...) present a mechanism to rapidly scale up the overall ADC computing capacity, but come with financial costs
  - The ADC cloud computing model will incorporate strategies to optimize cost effectiveness

- Overall roadmap:
  - Exploratory phase to start now and finish around end of the summer
    - Conclude with a Cloud Computing Model document
  - Development phase to start in fall



## 1. Basic Research

- Review the work already carried out within ATLAS, CERN IT, WLCG, sites, EGI and OSG
- If possible, organize workshop to collect information about existing cloud computing activities in collaborating organisations

## 2. Implement primitive data management and job execution on the cloud

- Virtual machines
  - CERNVM as a platform, which gives access to ATLAS tools and software
- Evaluate potential resources
  - Various sites (e.g. Magellan at ANL, Ixcloud at CERN, BNL cloud, other cloud infrastructures related to WLCG (e.g. in Canada), commercial clouds, etc...)
  - Various cloud APIs (Amazon (EC2, S3, etc.), OpenStack, Nimbus, etc...). Need to understand the long-term sustainability of the APIs.
- Implement primitive functionalities
  - Move data in and out of the cloud
  - Execute basic jobs on the cloud

## 3. Use-cases study

- Evaluate the ADC use-cases in relation to both commercial cloud computing. Estimate costs for various models.
- Explore new use-cases presented by cloud computing.

## 4. Design of the Cloud Computing Model

- The goal of this document will be to:
  - describe how ATLAS can make use of both pledged and chargeable cloud resources
  - present strategies to minimally impact the existing services so that the usage of cloud resources would be transparent to end-user physicists
  - present cost-effective models for the various ADC use-cases on commercial clouds
  - incorporate legal and security considerations

## 5. Development

- Initial ideas are detailed in the section below.

## 6. Testing

- Integration of the cloud with existing monitoring services, functional tests (DDM, analysis, production) for stability and reliability evaluations
- Perform stress tests of the cloud solutions to study performance

- Need to implement plugin libraries that support cloud I/O for Site Services and DQ2 clients
  - Probably need to support different cloud APIs depending on the available resources
  - May create requirements for FTS
- Bookkeeping of location information for ATLAS datasets stored in the cloud
- Investigate the usage of cloud-based content delivery solutions to automatically address “hot” data use-cases.

- PanDA pilot needs a mover module which can
  - stage-in (e.g. wget) cloud resident data
  - stage-out and register output files
- Investigate Condor glide-in technology for pilot factory (send one glide-in – get n pilots)
- PanDA server:
  - start with a set of cloud-based PanDA sites/queues
  - in later phases evaluate cost-aware job brokerage
- Automatic cloud resource provisioning:
  - Panda could request new/larger cloud capacity subject to global workload and cost considerations
  - A tool which enables automatic resource provisioning would need to be developed

- Add metadata about cloud-based resources into the ADC monitoring and information systems
  - Dashboards, BDII, AGIS, PanDA schedconfig, TiersOfATLAS...
- This point depends fully on the design of the ADC Cloud Computing Model

- Comments about the proposal document are welcome
- We expect development requirements will be placed on many areas in and outside ADC
  - collaboration with the larger team of ADC developers will be necessary
  - collaboration with other groups in WLCG will be emphasized in order to find solutions which can be maintained and therefore sustained by a larger community
- Next step: Wait for Cloud Computing Workshop announcement.