Proposal for Cloud Computing R&D in ATLAS Distributed Computing

1. Introduction

The ATLAS Computing Model was designed around the concepts of grid computing to combine the resources of more than one hundred Worldwide LHC Computing Grid (WLCG) sites for ATLAS offline data storage, distribution, processing and analysis. After more than a year in full operation, the system has demonstrated that it can successfully cope with the experiment's computing requirements during data-taking.

However, a new and emerging paradigm in the delivery of IT services – cloud computing – presents an improved approach to managing and provisioning resources, thereby allowing applications to easily adapt and scale to varied usage demands. By providing an "Infrastructure as a Service" (IaaS), clouds aim to efficiently share the bare hardware resources for storage and processing without sacrificing flexibility in services offered to applications. A new and increasingly competitive market has emerged to offer these cost effective computing resources, and companies small and large already make extensive use of them.

It is therefore of interest for an ATLAS Distributed Computing (ADC) R&D project to investigate opportunities for adapting its existing services to make use of cloud computing resources and target an upgrade of the ATLAS distributed computing technologies during the LHC shutdown in 2013.

The present document gives an initial proposal to research cloud computing for ADC, and then to 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.

2. Goals and Scope of the Project

The primary goals of this R&D project are

- to evaluate the available cloud technologies in relation to the use-cases presented by ATLAS data management, processing and analysis,
- to design a model for transparently integrating cloud computing resources with the ADC software and service stack
- to implement the ATLAS cloud computing model into DDM, PanDA and related tools and services.

Throughout this activity, collaboration with related groups in WLCG, CERN-IT, tier centers, other experiments and existing research clouds is desired in order to consolidate efforts and leverage commonalities between the interested parties.

Some of the potential ADC use-cases which will be investigated include:

- 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 paragraph)
- 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)

It is important to note that the relative attractiveness of cloud computing for each use-case varies when considering the resources which are provided by research and academic institutions versus commercial entities. In the former case, the usage of cloud APIs can be seen as an alternative to grid middleware for providing remote access to a site while better managing its local resources; one could say that ADC should be able to use a WLCG site that chooses to make itself available via a cloud API. In contrast, commercial cloud computing presents a mechanism to rapidly scale up the overall ADC computing capacity; however, the present project must consider that all operations with commercial clouds would be chargeable and therefore the resources must be used in a way that optimizes cost effectiveness. Working toward a model for the cost effective use of commercial cloud resources will be a main goal of this project.

3. Roadmap

Here we present a roadmap outline for this R&D project. The general timing of the roadmap is that exploratory work will occur during summer 2011, concluding with a cloud computing model document in early fall. Development will follow. Also note that some items below (notably 1,2, and 3) can happen in parallel.

1. Basic Research
    a. Review the work already carried out within ATLAS, CERN IT, WLCG, sites, EGI and OSG
    b. 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
    a. Virtual machines
        i. CERNVM as a platform, which gives access to ATLAS tools and software
    b. Evaluate potential resources
        i. Various sites (e.g. Magellan at ANL, lxcloud at CERN, BNL cloud, other cloud infrastructures related to WLCG (e.g. in Canada), commercial clouds, etc…)
        ii. Various cloud APIs (Amazon (EC2, S3, etc.), OpenStack, Nimbus, etc…). Need to understand the long-term sustainability of the APIs.
    c. Implement primitive functionalities
        i. Move data in and out of the cloud
        ii. Execute basic jobs on the cloud

3. Use-cases study
    a. Evaluate the ADC use-cases in relation to both commercial cloud computing. Estimate costs for various models.
    b. Explore new use-cases presented by cloud computing.
4. Design of the Cloud Computing Model
    a. The goal of this document will be to:
        i. describe how ATLAS can make use of both pledged and chargeable cloud resources
        ii. present strategies to minimally impact the existing services so that the usage of cloud resources would be transparent to end-user physicists
        iii. present cost-effective models for the various ADC use-cases on commercial clouds
        iv. incorporate legal and security considerations
5. Development
    a. Initial ideas are detailed in the section below.
6. Testing
    a. Integration of the cloud with existing monitoring services, functional tests (DDM, analysis, production) for stability and reliability evaluations
    b. Perform stress tests of the cloud solutions to study performance

4. Development Areas

4.1. Changes in ATLAS DDM

ATLAS Distributed Data Management needs to be adapted to the cloud computing infrastructure. This development process will include the implementation of plugin libraries that support file transfers into and out of the cloud and that will be used in DDM Site Services and the DQ2 clients. Since there is not a unified cloud standard, this step might require the evaluation/usage of different cloud APIs (Amazon's S3, Open Cloud Computing Interface etc.) depending on the chosen technologies in the available resources. This work may create requirements for the FTS development team.

A different point to address is the bookkeeping of location information for ATLAS datasets stored in the cloud. It has to be seen if the cloud offers reliable file catalogues and how to integrate these with the existing DDM catalogues.

On a different note, the usage of existing commercial cloud services for content delivery may be investigated. These services allow data located in the cloud to be automatically replicated around the world; this may have applications in ATLAS in relation to data distribution for analysis and other "hot" data use-cases.

4.2. Changes in the workload management

We envisage that some development will be necessary in the PanDA pilot and server. The PanDA pilot will need a mover module which can stage-in (e.g. wget) cloud resident data and stage-out and register output files. The PanDA server should not require major developments in the early phases of this project; one could start with a set of cloud-based PanDA sites/queues, which would be

transparent to the server. In later phases of the project we envisage opportunities for

- Cost-aware job brokerage: PanDA could consider the estimated financial cost when selecting sites for a production task or analysis jobset.
- Automatic cloud resource provisioning: Panda would request new or larger cloud capacity subject to the global workload and cost considerations.

Further to this last point, a tool which enables automatic resource provisioning would need to be developed. This tool could be used by services such as PanDA or also human ADC operators to provision resources on-demand and automatically configure them for ADC applications.

4.3. Integration with the monitoring and information systems used in ADC

It will be necessary to add metadata about the cloud-based resources into the monitoring and information systems that ADC depends on (Dashboards, BDII, AGIS, PanDA schedconfig, TiersOfATLAS). This point will depend fully on the results of the design of the Cloud Computing Model.

5. Conclusions

This proposal presents the goals and a roadmap for the introduction of cloud computing into ADC in collaboration with CERN IT. After an initial exploration and prototyping phase, this project will introduce a Cloud Computing Model for ATLAS and deliver an initial implementation of the work. Development requirements will be placed on many areas in ADC, and collaboration with the larger team of ADC developers will be necessary. Finally, collaboration with other groups in the WLCG will be emphasized in order to find solutions which can be maintained and therefore sustained by a larger community.