Cloud Computing Perspectives in HEP

@3rd DPLTA Workshop

@ CERN

7 December 2009

by

Homer Neal (SLAC NATIONAL ACCELERATOR LABORATORY)

BaBar Computing Coordinato

Defining the Cloud

- cloud computing - your programs are run at single site providing virtualized services ELSEWHERE
  - In contrast to a GRID which is a « set of distributed and highly integrated facilities » - R. Sobie
- Typically taking advantage of spare computing cycles of large companies like Amazon (Amazon Elastic Compute Cloud, **EC2**)
- “Commercial cloud providers such as EC2 allow users to deploy groups of unconnected virtual machines, whereas scientists typically need a ready-to-use cluster whose nodes share a common configuration and security context. The Nimbus Context Broker bridges that gap,” said Kate Keahey, a computer scientist at Argonne and head of the Nimbus project. - http://insidehpc.com/2009/03/26/cern-and-argonne-use-science-clouds-for-computing/

- You don't know who the other users are
- You pay as you go

Using EC2

http://aws.amazon.com/ec2/#functionality

Amazon EC2 Functionality

Amazon EC2 presents a true virtual computing environment, allowing you to use web service interfaces to launch instances with a variety of operating systems, load them with your custom application environment, manage your network's access permissions, and run your image using as many or few systems as you desire.

To use Amazon EC2, you simply:

- Select a pre-configured, templated image to get up and running immediately. Or create an Amazon Machine Image (AMI) containing your applications, libraries, data, and associated configuration settings.
- Configure security and network access on your Amazon EC2 instance.
- Choose which instance type(s) and operating system you want, then start, terminate, and monitor as many instances of your AMI as needed, using the web service APIs or the variety of management tools provided.
- Determine whether you want to run in multiple locations, utilize static IP endpoints, or attach persistent block storage to your instances.
- Pay only for the resources that you actually consume, like instance-hours or data transfer.

symmetrybreaking
extra dimensions of particle physics
A joint Fermilab/SLAC publication

VOLUME 06 ISSUE 05
OCTOBER 09

VIEW CURRENT ISSUE

Belle

A side of cloud with your grid, ma'am?

May 21, 2009 | 10:34 am

As the cloud becomes a more popular computing solution in the commercial world, it is starting to pique the interest of the academic research community. Collaborators of the experiment at the High Energy Accelerator Research Organization, known as KEK, in Japan are considering supplementing their computing with Amazon's Elastic Compute Cloud (EC2), which provides on-demand, virtual computing resources over the Internet.

Belle II co-spokesman Masanori Yamauchi (left) with 2008 Nobel prize winner in physics Makoto Kobayashi. Image courtesy of Belle II.

The KEKB particle accelerator produces a more densely-packed, intense beam than any currently operating collider, and an upgrade to the machine in 2013 is expected to increase the intensity by a factor of 50. The new Belle II detector will come online at the same time equipped to handle the estimated hundred-fold increase in data—expected to total 40 petabytes per year. This load will require more than 100,000 CPU cores, leading the collaboration to consider new computing options, said Martin Sevior, a KEK collaborator based at the University of Melbourne in Australia.

Sevior and his colleagues ran the complete Belle simulated data analysis chain on EC2 to test it, and found it easy to deploy jobs. They created an Amazon Machine Image (AMI), a computing environment customized to contain both the Scientific Linux operating system and applications for the Belle analysis system. Each AMI contains eight CPU cores—imagine eight PCs—and can be duplicated 20 times, creating a virtual 160-core cluster. To lower costs, the team set up an automation system that duplicates the AMI on demand and shuts down instances of it as need drops.

BaBar Legacy Analysis

Can clouds provide a solution? - R. Sobie (UVic)

- One model is to build a cloud-like facility at SLAC
  - Users run old-OS platforms in a virtual environment
- We propose extending this model to utilize external research computing facilities
  - BaBar has access to many resources outside of SLAC
  - Access to non-HEP resources (e.g. FutureGrid)
  - Possible model for SuperB computing
- We (Victoria) have been funded to explore this option
  - Other communities in Canada (astronomy) developing a similar plan

Randall Sobie   University of Victoria

Cloudy BaBar

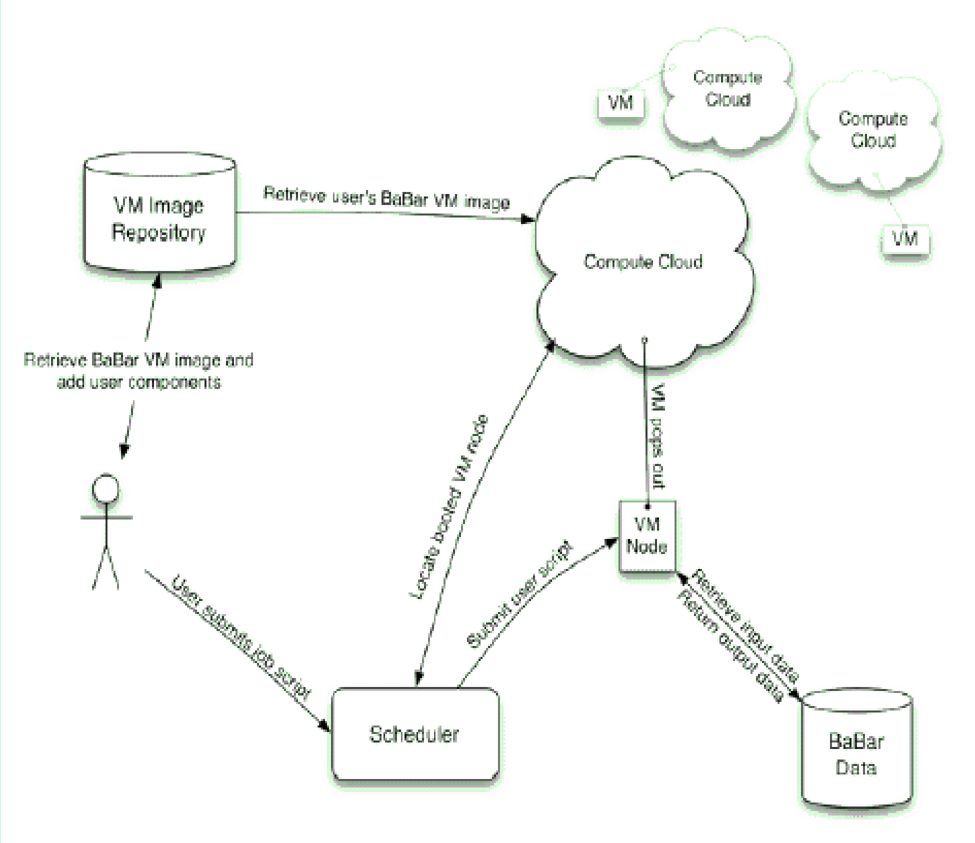

Based on Nimbus cloud software (Argonne NL project)

Ability to submit to Amazon EC2 and Eucalyptus clouds
(Eucalyptus is an open source and commercial provider of Amazon-like software)

Cloud scheduler based on Condor

Use Grid Certificates for authentication

Similar environment being built for astronomy community in Canada

- « In fact CERN and others are investigating the feasibility of layering cloud services from Amazon and others over the grid. »

http://www.computerweekly.com/Articles/2009/12/02/239433/CERN39s-LHC-pioneers-quantum-leap-in-cloud-computing.htm

,

ComputerWeekly.com

You are here Networks and Communications › Network Infrastructure

CERN's LHC pioneers quantum leap in cloud computing

David Binning
Monday 23 November 2009 04:21

The Large Hadron Collider at CERN **(the European Centre for Nuclear Research) in Geneva,** which was switched back on **on the 21 November** following an electrical fault which saw the particle collider switched off in September 2008, **is** designed to create hundreds of millions of collisions between subatomic particles every second**, in an attempt to advance knowledge of energy and matter, but most significantly to piece together the events immediately after the big bang.**

Four experiments spaced along an underground tunnel the size of the London city line will produce 15 petabytes of data per year, creating one of the greatest challenges ever in the history of computing, requiring that data be collected, formatted, stored, shared and secured on a massive scale.

On site is a staggering concentration of heavy-duty computing power, which includes 5,700 systems, 36,600 processing cores, 41,500 disk drives, 45,000 tape cartridges and 160 tape drives. But this only begins to touch on the effort needed to support the project.

When the Large Hadron Collider accelerator at CERN is running at full strength, access to experimental data needs to be provided for over 10,000 scientists in several hundred research institutes and universities worldwide participating in the Large Hadron Collider (LHC) experiment. All the data need to be available over the 15-year estimated lifetime of the LHC. Analysis of the data, including comparison with theoretical simulations, needs about 100,000 CPUs at current measures of processing power.

CERN's IT department head, Frédéric Hemmer is helping to drive an ambitious project to develop huge grid computing networks that would support the LHC and many other important research initiatives throughout Europe and rest of the world.

Frédéric Hemmer's former colleague, Tim Berners-Lee, of course invented the web. Now his and CERN's work on grids is itself leading to powerful and interesting innovations leading to major changes in computing and communications; most notably the move towards the cloud.

...

He adds: "As the technology stabilises through improved international standards and reliability, grids are gaining acceptance in mainstream business and science communities." CERN's Hemmer observes a trend towards industry-specific cloud platforms which operate much the same way as grids, only that the users still essentially rent the infrastructure and applications from a service provider. The benefit of this is that things like capacity and security can be more easily guaranteed.

However, it is expected that the pricing and business models currently in use in the cloud will look very different in years to come, as more industry- and application-specific services emerge and companies demand more flexibility.

Clouds and DPLTA

- Archived data will be used sporadically so the cloud pay-as-you-go seems very well suited
- Images/Data kept at cloud site
  - Funding needed for idle period (small cost) as well as production periods
- Still need a system to support the rest of the analysis infrastructure and allow users to test setup/analysis/production without having to immediately apply for a grant

Clouds and DPLTA Concerns

- Ability of virtual machines to fully isolate jobs from the unknown hardware
- Who governs access to the repositories
- Can they be used for more than simulation production?
  - Analysis tasks can require as much computing resources.
    - Who would manage the transfer of any newly simulated data to the cloud repository?

Summary

- DPHEP should recommend this as a viable means of procurring any extra CPU and or storage resources needed for production and analysis activities in the experiments archival period
- Guidance on management needs to be provided
  - Requirements on the « collaboration »
  - Regulation (or not) of access
- Guidance on funding needs to be provided
  - Especially for the idle periods
- When in the collaboration lifetime does the use of clouds needs to be tested?

Summary

- Recognize that preparation of cloud use is directly related to preparation for virtual machine usage.