

HEP Data-Intensive Distributed Cloud Computing

- Agenda
- Project team
- Project motivation
- Current use in ATLAS and BelleII

Agenda

- Introduction and project overview
- Project Milestones and Deliverables (Colin)
- Subprojects:
 - VM Distribution. (Ron)
 - Storage Federation (Frank, Ryan)
 - Shoal (Ian)
- Documents review
 - Design document
- Review of project work plan
- Discussion

Project team

- R. Sobie Project leader
- F. Burghaus Applications
- A.Charbonneau Developer
- R.Desmarais Developer
- I.Gable Developer
- C. Leavett-Brown Developer
- M.Paterson Developer
- R.Taylor Applications
- A.Lam Coop

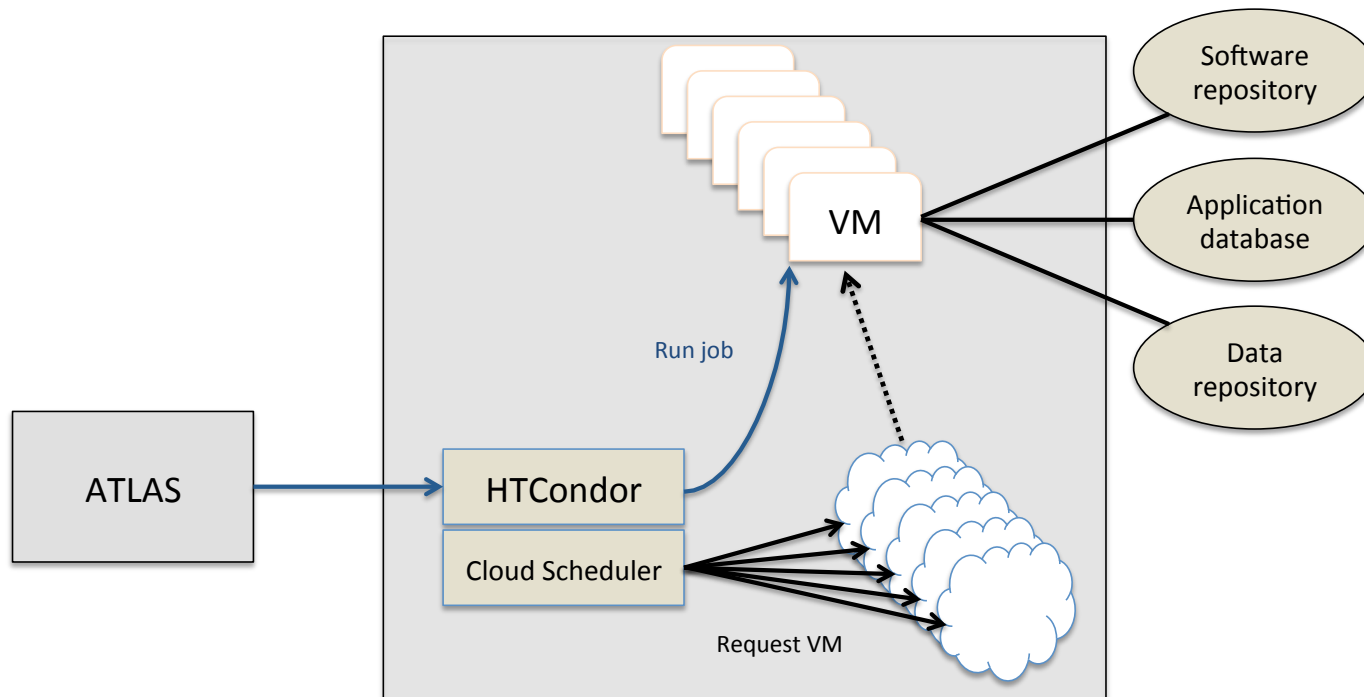
Motivation I


- ATLAS and other HEP experiments operate a distributed computing system (Grid) that is now evolving to cloud technology
 - CERN resources are soon to be provided through virtual machines (VMs) either transparently to the user or through a cloud interface
- We have been leading the initiative in ATLAS to use clouds (using our CS software developed in NEP52)
 - Our initial system is the longest running cloud system in ATLAS
 - CS is used by NECTAR (Australia) for ATLAS
 - CS is being installed at CERN to manage the CERN-ATLAS cloud
 - Testing on Google Compute Engine
- We have deployed a CS for the BelleII project in Japan
 - Amazon EC2 grant

Motivation II

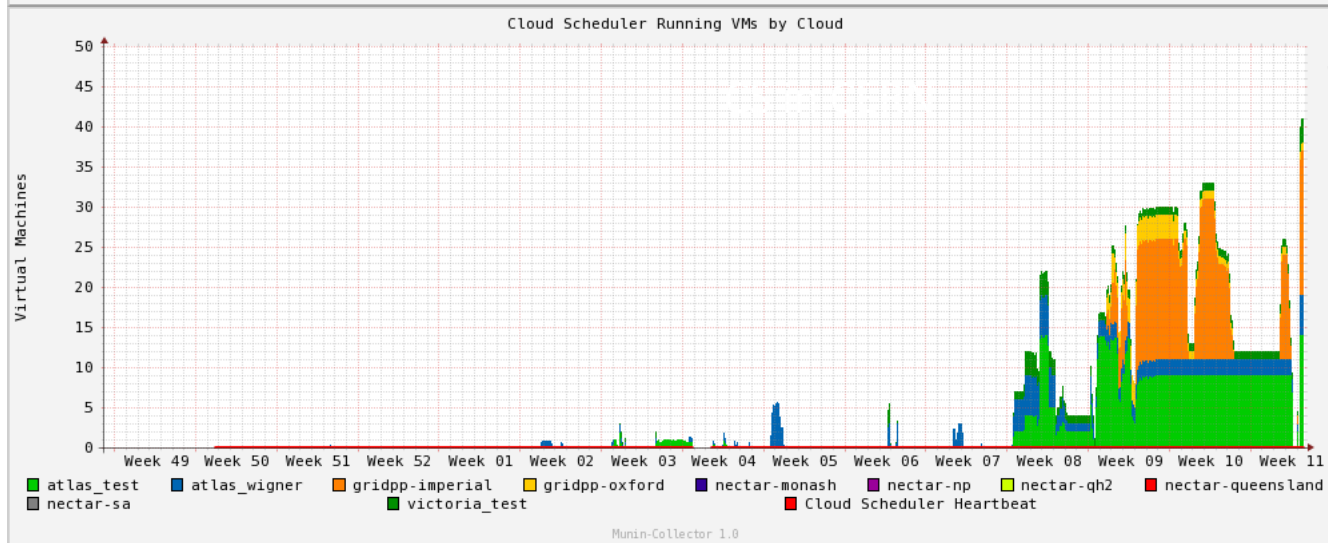
- We are running “production” jobs with no input data
(production means it is managed by a central team)
 - Simulation of the particle collisions
 - Simulation of the interaction of each particle with the detector
- We are planning to run “production” jobs with input data
 - “pile-up” jobs (overlying multiple particle collision data)
- We do not plan to run “analysis” jobs at the moment
 - ATLAS job scheduler will need significant changes

System Overview

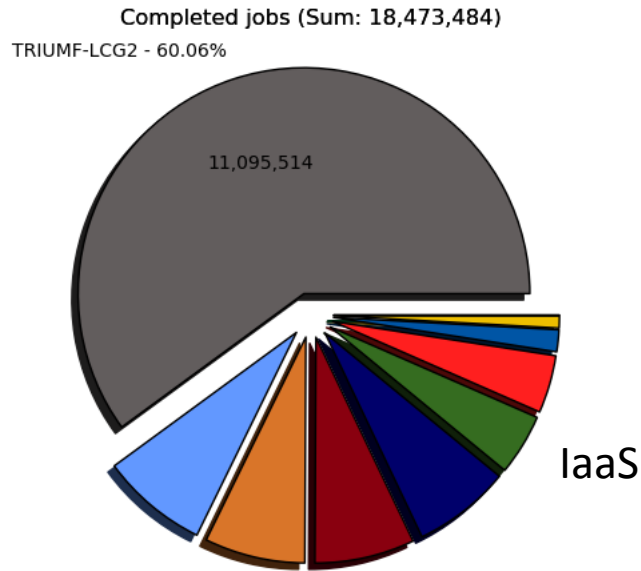




CS in Victoria has been in operation for 2 years



CS in CERN was installed in Feb 2014



- TRIUMF-LCG2 - 60.06% (11,095,514)
- CA-SCINET-T2 - 7.20% (1,329,671)
- SFU-LCG2 - 6.98% (1,288,997)
- AUSTRALIA-ATLAS - 4.13% (763,011)
- AUSTRALIA-NECTAR - 0.86% (159,755)
- CA-MCGILL-CLUMEQ-T2 - 7.67% (1,416,945)
- CA-VICTORIA-WESTGRID-T2 - 7.14% (1,319,909)
- IAAS - 4.38% (809,561)
- CA-ALBERTA-WESTGRID-T2 - 1.56% (288,553)
- CA-JADE - 0.01% (1,568)

IaaS cloud queue in ATLAS is equivalent to a Canadian Tier 2 site (even though we do not run all types of jobs)

Original IaaS cloud for ATLAS:

```
[rsobie@elephant belle2]$ cloud_status
```

Resource pool Resources:

NAME	CLOUD TYPE	NETWORK ADDRESS	
Alto	OpenStack	alto.cloud.nrc.ca	xen
Synnefo	Nimbus	synnefo.westgrid.ca	xen
nectar-qh2	OpenStack	nova.rc.nectar.org.au	kvm
FGFoxtrot	Nimbus	f1r.idp.ufl.futuregrid.org	xen
nectar-monash	OpenStack	nova.rc.nectar.org.au	kvm
Jadeite	Nimbus	nimbus.nic.ualberta.ca	kvm
Elephant	Nimbus	elephant.heprc.uvic.ca	kvm
FGHotel	Nimbus	svc.uc.futuregrid.org	xen
dair-ab	OpenStack	nova-ab.dair-atir.canarie.ca	xen
Mouse	OpenStack	mouse01.heprc.uvic.ca	xen
gridppcl	OpenStack	gridppcl02.grid.hep.ph.ic.ac.uk	xen
dair-qc	OpenStack	nova-qc.dair-atir.canarie.ca	xen
nectar-np	OpenStack	nova.rc.nectar.org.au	kvm

Belle II cloud

```
[root@bellecs ~]# cloud_status
```

Resource pool Resources:

NAME	CLOUD TYPE	NETWORK ADDRESS
alto	OpenStack	alto.cloud.nrc.ca
Hotel	OpenStack	c19.uc.futuregrid.org
amazon	AmazonEC2	ec2.us-west-2.amazonaws.com
Elephant	Nimbus	elephant.heprc.uvic.ca
cernopenstack	OpenStack	openstack.cern.ch
mouse	OpenStack	mouse01.heprc.uvic.ca
synnefo	Nimbus	synnefo.westgrid.ca

Cloud at CERN:

\$ cloud_status

Resource pool Resources:

NAME	CLOUD TYPE	NETWORK ADDRESS
gridpp-imperial	OpenStack	gridppcl02.grid.hep.ph.ic.ac.uk
nectar-qh2	OpenStack	nova.rc.nectar.org.au
atlas_wigner	OpenStack	openstack.cern.ch
nectar-queensland	OpenStack	nova.rc.nectar.org.au
atlas_test	OpenStack	openstack.cern.ch
nectar-monash	OpenStack	nova.rc.nectar.org.au
nectar-sa	OpenStack	nova.rc.nectar.org.au
gridpp-oxford	OpenStack	cloudhead.physics.ox.ac.uk
nectar-np	OpenStack	nova.rc.nectar.org.au
victoria_test	OpenStack	openstack.cern.ch

In addition, operate a test system that will be used to evaluate Google Compute Engine

Group web site: <http://heprc.phys.uvic.ca>

The screenshot shows the website for High Energy Physics Research Computing. The header is blue with the text "High Energy Physics Research Computing" in white. Below the header is a navigation bar with "Home" as a link. The main content area is titled "Publications" and lists several articles from 2013 and 2012. Each article entry includes a title, authors, and a link to the publication details.

High Energy Physics Research Computing

Home

Publications

2013

Dynamic web cache publishing for IaaS clouds using Shoal
I Gable, M Chester, P Armstrong, F Berghaus, A Charbonneau, C Leavett-Brown, M Paterson, R Prior, R Sobie, R Taylor
[arXiv:1311.0058](#) Proceedings of the CHEP 2013, Amsterdam 2013

ATLAS Cloud R&D
Panitkin, S ; Barreiro Megino, F ; Caballero Bejar, J ; Benjamin, D ; DiGirolamo, A ; Gable, I ; Hendrix, V ; Hover, J ; Kucharczuk, K ; Medrano LLamas, R ; Love, P ; Ohman, H ; Paterson, M ; Sobie, R ; Taylor, R ; Walker, R ; Zaytsev, A
To appear in the proceedings of the CHEP 2013 Amsterdam 2013

HTC Scientific Computing in a Distributed Cloud Environment
R.J. Sobie, A. Agarwal, I. Gable, C. Leavett-Brown, M. Paterson, R. Taylor, A. Charbonneau, R. Impey, W. Podiama
[Science Cloud '13 Proceedings of the 4th ACM workshop on Scientific cloud computing](#)

Efficient LHC Data Distribution across 100Gbps Networks
H. Newman, A. Barczyk, A. Mughal, S. Rozsa, R. Voicu, I. Legrand, S. Lo, D. Kcira, R. Sobie, I. Gable, C. Leavett-Brown, Y. Savard, T. Tam, M. Hay, S. Mckee, R. Hocket, B. Meekhof, S. Timoteo.
[SCC '12 Proceedings of the 2012 SC Companion: HPC, Networking Storage and Analysis IEEE Computer Society, May 2013](#)

2012

Disk-to-Disk network transfers at 100 Gb/s
Artur Barczyk, Ian Gable, Marilyn Hay, Colin Leavett-Brown, Iosif Legrand, Kim Lewall, Shawn McKee, Donald McWilliam, Azher Mughal, Harvey Newman, Sandor Rozsa, Yvan Savard, Randall J. Sobie, Thomas Tam, Ramiro Voicu. Proceedings of the 2012 CHEP Conference, Journal of Physics: Conference Series 396 (2012) 042006 New York City 2012 (doi: [10.1088/1742-6596/396/4/042006](#))

Exploiting Virtualization and Cloud Computing in ATLAS
Fernando Harald Barreiro Megino, Doug Benjamin, Kaushik De, Ian Gable, Val Hendrix, Sergey Panitkin , Michael Paterson, Asoka De Silva, Daniel van der Ster, Ryan Taylor, Roberto A. Vitillo, Rod Walker on behalf of the ATLAS Collaboration. Proceedings of the 2012 CHEP Conference, Journal of Physics: Conference Series 396 (2012) 032011 New York City 2012 (doi: [10.1088/1742-6596/396/3/032011](#))

Press

Videos: [HEP RC YouTube Videos](#)

2014

OpenStack Users Video

[Youtube - February 2014](#)

CANARIE invests in 9 innovative software projects and boosts Canadian researchers toward faster scientific discovery

[CANARIE News February 2014](#)

2013

International team demonstrates intercontinental OpenFlow network at Super Computing 2013

[BCNET News December 2013](#)

Enabling subatomic data sharing

[International Innovation July 2013](#)

Enterprises launch their own private clouds

[CNN Money, June 18 2013](#)

150,000 cloud virtual machines will help solve mysteries of the Universe.

[ars Technica \(UVIC-ATLAS listed as largest user in screen shot\)](#)

Physicists link OpenStack, Nimbus clouds around the world to share resources [IT World, April 2013](#)

OpenStack Users Step Into the Limelight [Innovation Insights](#)

OpenStack Summit keynote in High Energy Physics acknowledges Nimbus and FutureGrid [FutureGrid](#)

Moving at the speed of data

[University of Victoria Research Letter January 2013](#)

Summary

- Project is going well
 - Work plan is on schedule
 - Budget on track
- No issues (though timescales are extremely tight)
- Software developed in NEP52, CANFAR and NEP101 is being used around the world for distributed cloud computing