

A Cloud Project Proposal

- Introduction and motivation
- Current ATLAS cloud activities
- Possible projects

Introduction and motivation

- IaaS clouds gaining popularity in industry
 - elastic scaling of resources, pay as you go, reliability
 - concentrate on core competence
- Does this apply to us? Maybe not public clouds.
 - >80% occupancy – then buy your own
 - special performance needs, e.g. analysis.io
- Profit aside, are genuine cost saving of fewer larger resources
 - scale, infrastructure, manpower, ...
 - quickly loses advantage if need N people per VO
 - operate large site as cloud, then it needs no VO specific support
 - rather than hand over cash to Amazon – investigate possibility of community cloud
- ATLAS and CMS-only sites: lack opportunistic sharing due to manpower, e.g. CMS validation, or ATLAS space token obsession.
 - Cloud techniques will give back this Grid ideal

ATLAS Cloud Activities

- A few summer student level pilot projects
- 1FTE @ LBNL, cloud provisioning and control
 - CloudCRV to start/stop various node types
 - uses Puppet to configure them
 - not yet tried for production.
 - small group of interested parties attend meetings
 - strong US bias. BNL will hire 1 FTE for Cloud.
- Future computing topic of interest
 - funding bodies will ask if we can benefit from cloud
 - working group/task force formed
 - early days
- CMS has some projects
 - I know of Brian Bockleman, Karlsruhe Virtualized WN

Public Cloud Evaluation

- Principally Amazon EC2, but there are others
- Rent or buy?
 - Helpful EC2 spreadsheet gives one answer!
 - \$0.72/core/day. 1000 cores for a year ~260k\$
 - scales with occupancy
 - buy 1000 cores - use for 3/4 years
 - electricity, cooling, network, rack space, manpower
 - balance is still with buying if occupancy > 70%
 - but this may change. Academic price?
- Funding for $O(10)$ servers constant, plus short $O(1000)$ bursts for scaling tests, or high priority work.
 - deliverable is large scale simulation

Community Cloud Evaluation

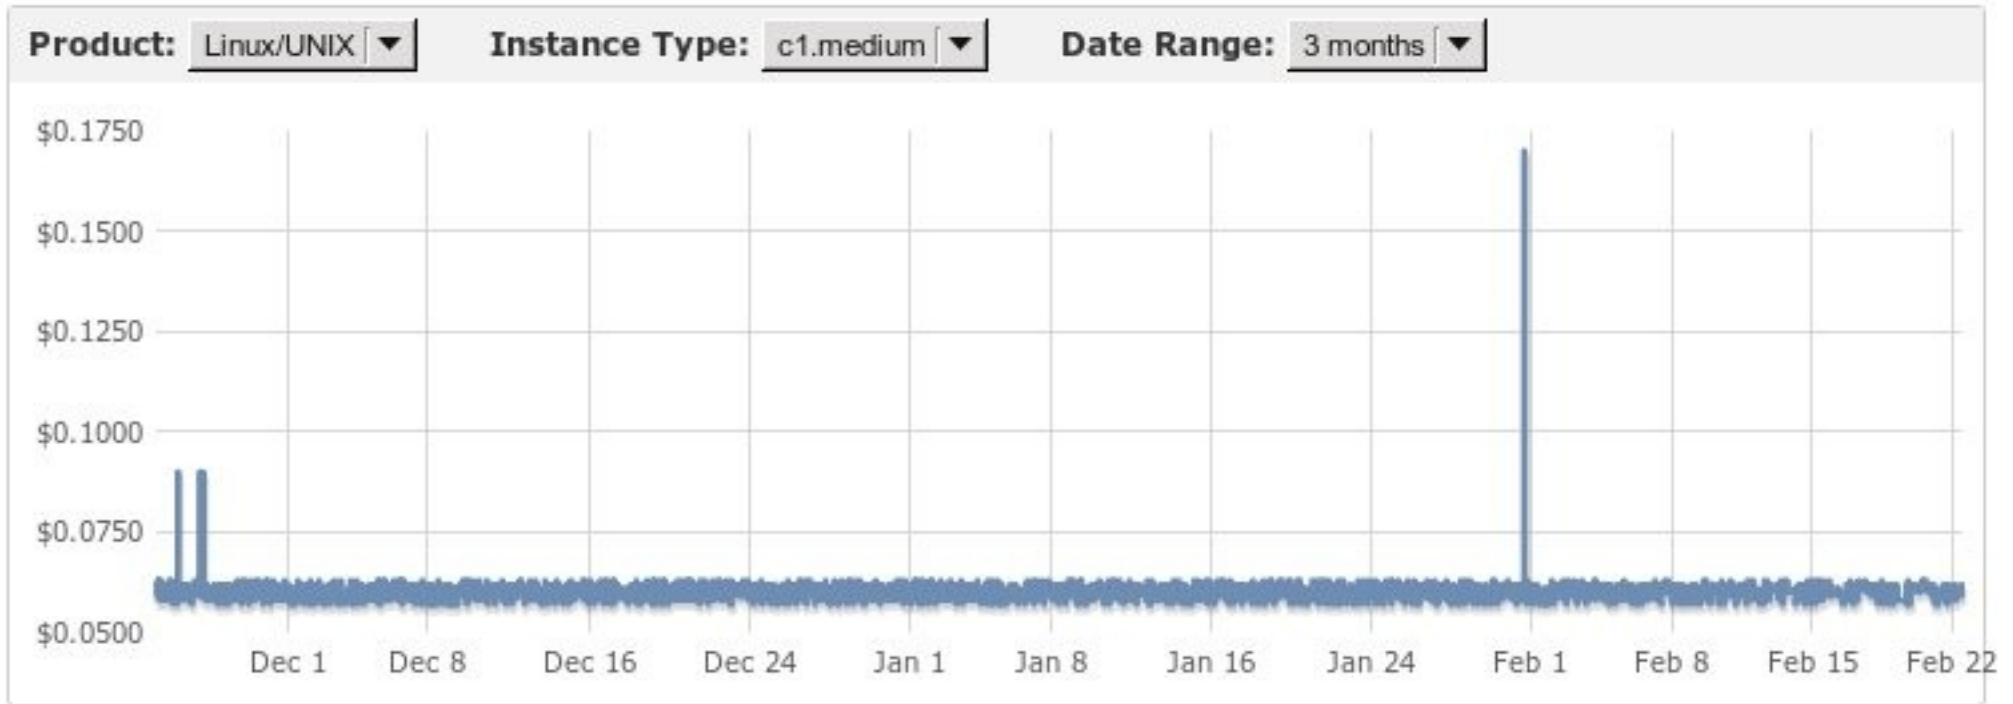
- More data-intensive friendly and cheaper
 - io is free (c.f. 10c/GB). Storage available.
- No large scale example but some demonstrators
 - Magellan(LBNL) and StratusLab
 - StratusLab – 3MEuros, 2yr EU project
- Possible projects
 - partially convert one or few existing sites to cloud
 - allow other VOs access
 - can we use clouds for analysis
 - io optimization
 - Private/Hybrid clouds – alternative Grid deployment
 - Dell & ibm like this concept – more of the same

Cloud Configure & Control

- Common work to support VOs in cloud
- Preparing VM image, starting and configuring
- Storage in cloud: dCache, HDFS, xrootd?
- Monitoring & VM life cycle
 - start more workers, retire hosts
- Wallet optimizations
 - EC2 spot price instances ...
- Support group recipes, documentation to avoid everyone doing this themselves.

Spot Instance Pricing History

Cancel



Close

Medium compute – 2 cores of 2.5 EC2 units (1EC2 unit = 2007 Xeon 1.2GHz)

On-Demand @ \$0.17/hr

Spot @ \$0.06/hr

Spot instances killed without mercy when price climbs above what you are prepared to pay.

Specify maximum, but only pay the spot price.

1 spike in 3months and no daily/weekly structure

Suggests that spot price market is not really active.

If this was smooth then could optimize usage to limit uncontrolled kills.

Longer term vision

- Bulk compute resources envisaged to grow with LHC data
 - Good reasons for them to be at a single site(per funding pot)
 - Economies of scale, manpower, co-locate with cheap/green electricity/cooling. Opportunistic use & sharing is built-in.
 - Less data movement (save on HEPPI). Move roottuples, not TBs.
 - Cloud paradigm keeps resource 'owner' in control
 - DE Scientific bulk computing provided by Science Cloud?
 - better value for existing VOs
 - removes overhead for new projects
- User analysis, PROOF cluster at T2
 - bulk computing needs can piggy back these
 - local funding, 'free' infrastructure and manpower
 - diversity in technology choices
 - good reasons to grow T2s too

Conclusions

- Proof of concept & operational models can be done in public cloud, at scale.
- Economics & Analysis prefers community cloud
 - StratusLab collaboration, volunteer site
- Clear opportunity for cooperation between LHC experiments – cloud control, storage
- Necessary work and evaluation for Science cloud
 - should the question arise.
- Private clouds opportunities
 - sustainable support, VO sharing