



# CMS in the Cloud

GridPP Cloud Meeting  
30 November 2012

Andrew Lahiff, Chris Brew



Science & Technology Facilities Council  
Rutherford Appleton Laboratory



# Why clouds?

- An interesting way to get additional resources when we need them
- An opportunity to develop components and services that we can float onto generic infrastructure
  - The images required for the cloud are the same as we could be developing for easy deployment of things
  - Being compliant with clouds may make our framework more modular/flexible/extendable...
- We may contribute to the development of cloud technologies so that they don't diverge too much from our needs
- A good way to stay relevant in high exposure development areas





# What's the end result?

- The high level goal is to make use of additional resources from the cloud at a level that is noticeable to the experiment
  - Doesn't make sense to get 10% additional resources for a month. It's a lot of work, maybe some money, and in the end no one will notice
  - Goal should be an infrastructure that could double the simulation capacity of CMS for a month
    - This would be gaining 25-30k cores for a month
    - Enable high priority samples to be completed a month early





# Requirements to run a CMS job

- CMS software framework (CMSSW)
  - CVMFS
- Frontier
  - Squid server in the cloud?
- Data (read access)
  - xrootd + sufficient WAN
- Data (write access)
  - SRM, or xrootd?
- Resource provisioning
  - glideinWMS starts VMs with pilots which register with central queue
  - Eliminates the need for a CE





# Activities

- StratusLab
- LxCloud
- Amazon spot instances
- HLT cloud





# StratusLab

- First tests carried out by CERN IT-ES
  - Dedicated Condor manager setup in the cloud; CRAB2 configured to submit to this
  - Software via CVMFS
  - Input data via xrootd
  - External Frontier squids
- Jobs ran successfully
  - MC
  - Analysis





# Trial of Amazon spot instances

- Spot instances
  - EC2 opportunistic access
  - Low-cost (~10x less than on-demand instances)
  - Actual cost variable via spot market (bid x dollars per hour)
- Trial runs
  - Attached EC2 VMs to T2\_US\_Wisconsin
    - 3 cores for 1 month, 100 cores for 1 week
  - 55% of cost was for data transfer
- Ideal use case
  - Low output to CPU ratio
  - Short jobs (less work lost due to instance termination)
  - Distant deadline (more flexibility to use cheaper options)





# HLT cloud

- The CMS High Level Trigger:
  - 13312 cores, 211280 HS06 → large resource
  - No storage, some computers have small disks
- Why not use these computers during the shutdown?
  - Actually, why not use them whenever they're not in use?
- Conceived as an Overlay Infrastructure
  - Try to minimize the impact on the existing production HLT cluster configuration (software & networking)
- Technology used:
  - Openvswitch
  - OpenStack (with as KVM underlying hypervisor)







# HLT cloud: status

- Work is ongoing
- Will be first major production use of cloud infrastructure by CMS
- The exercise of deploying an overlay cloud on the CMS online cluster is fruitful for the future
  - Can share knowledge of how to deploy such an overlay layer to existing sites that may transition this way to a cloud infrastructure while keeping existing services running
  - Gaining experience on how to contextualize and deploy VMs on the cloud infrastructures that are becoming commonplace





# Storage

- All tests have used storage outside the cloud
- Alternatives?
  - Setup a virtualized SE in the cloud
  - Cloud storage





# Looking forward

- Cloud provided computing tends to be factors more expensive than providing the resources in house for resources that are heavily used
  - Costs for commercial facilities is coming down
    - Interesting to cover peak periods
- Need to prepare for a time when this could be the norm
- Even if we never use large scale commercial resources there could be a benefit to making the infrastructure work
  - Simplify the packing and deployment of infrastructure and services
  - May increase ability to share resources across disciplines

