Scientific Computing Cloud

- Rod Walker, LMU Munich , 28th April 2014

1

Disclaimer: my own views, not endorsed by LMU or ATLAS

Acknowledge: dangers of linking HEP computing funding to HPC or wild new schemes

Local T2/3 cluster vs large facility

- Choose own hw & config
- Dedicated use
- Local extra funds
- Free room, cooling, electricity, network, manpower

- More HS06 per Euro
- Green IT
 infrastructure
- Reduced manpower
- Better reliability
- Shared, for bursts
- Network connectivity

University Computer Centre maybe large, can be green but not serving all DE science. No IAAS.

More a vision than a prediction

- National science funding buys single large facility
 - HW choice and design agreed by stakeholders
 - IAAS cloud interface custom VMs per user group
 - anyone with a grant: VO, university, chair
- ATLAS funding \rightarrow guaranteed share of facility
 - fully configurable by ATLAS(DE)
 - serial batch queue, MCORE slots, 6GB RAM,...
- extra Uni/regional funding \rightarrow guaranteed share
 - fully configurable by HEP group
 - proof cluster, batch, mylxplus, prod/analy, ..
- But what about our Tier2s?
 - T2 ATLAS-engaged manpower configuring/operating Cloud resources
- At least 3 arguments against

Get over it!

1) Freebies: University machine room, network, manpower, cooling, electricity are NOT free.

- total cost of ownership, capital and running costs, must be considered as one
 - eg. LRZ T2 run 2007 hw 5W/HS06 c.f. 2012 0.9W/HS06 turn it off! 3KEuros/mth
- 2) T2 empire: HEP people are not sysadmins or IT innovators
 - HEP success in sw and services experimentation
 - xrootd, PoD, ARC CE, cvmfs, hdfs, www(but not internet)
 - IAAS great tool for this
 - no benefit from site hw choice/setup experimentation
 - many bad decisions

3) Extra resources: Local money need not buy local hw

- get less for your money, with hassle and zero advantage
 - ssh works fine and RTT negligible
- sticker on the rack? Publicity shot? No problem.

Good for HEP

- More guaranteed resources
 - attractive, new funding request: multi-disciplinary, green, Cloud
 - economies of scale, purchase/manpower, better PUE for power
 - PUE 1.4-> 1.1 saves ~30% power, or 500kEuros per MWyr
- Plus opportunistic scale out and bursts
- Fewer sites with better admin & reliability
 - operations people very much appreciate this
 - less recycled HEP manpower playing sysadmin (include myself)
 - networking easier: SDN potential c.f hassle LHCONE deployment, PerfSonar
- Larger sites
 - less fragmentation of data and cpu, co-location easier



Big Science Unified IT



- More bang for the buck of science project computing funding
- Separate core competencies
 - IT admin, and no per user group manpower
- Gives EGI follow-on some meat
 - science gateways
 - standard recipes/patterns for cloud service deployment
 - Batch system, S3, Db, hdfs, web server, Proof,...
- D-HEP really takes D-BiScUIT