Outsourcing with clouds

Peter Love
Lancaster University
HEPSYSMAN June 2015 - Rutherford Appleton Lab

Introduction

- The deployment of academic and commercial cloud facilities is becoming more widespread
- What does it take to use these opportunistic resources?
- HEP experiments are in a good position to take advantage of these resources, a variety of tools are being used
- Described here are some experiences when commissioning the site and compare a few metrics with our grid site
- Will mention a recent Helix-Nebula experience.
- What can we say from a site's point of view?

Datacentred

- Relatively new commercial operation providing co-lo facilities and a cloud hosting service
- Built on Openstack and Ceph
- Large facility 1800sq.m2, enough for 850 30kW racks with UPS backup
- Involvment in Helix Nebula becoming member in early 2015. Engaging with big science projects.

Datacentred

- Lancaster collaboration started at early stages of commissioning the facility, we have a small tenancy

- Objective from our side is to exploit opportunistic resources in as simple way as possible. This is not AWS or GCE.

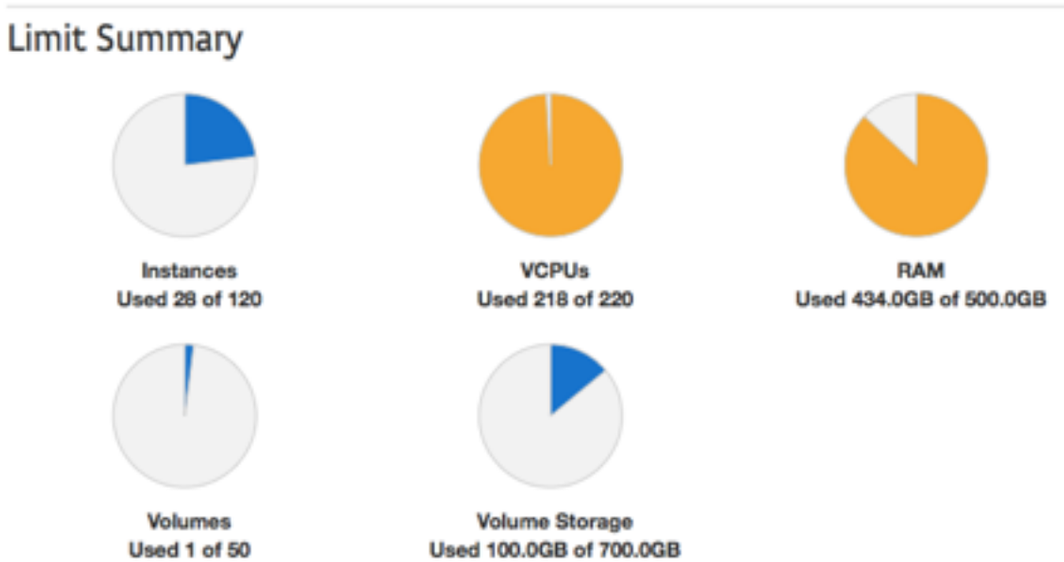

What approach do we use to make use of this opportunity?

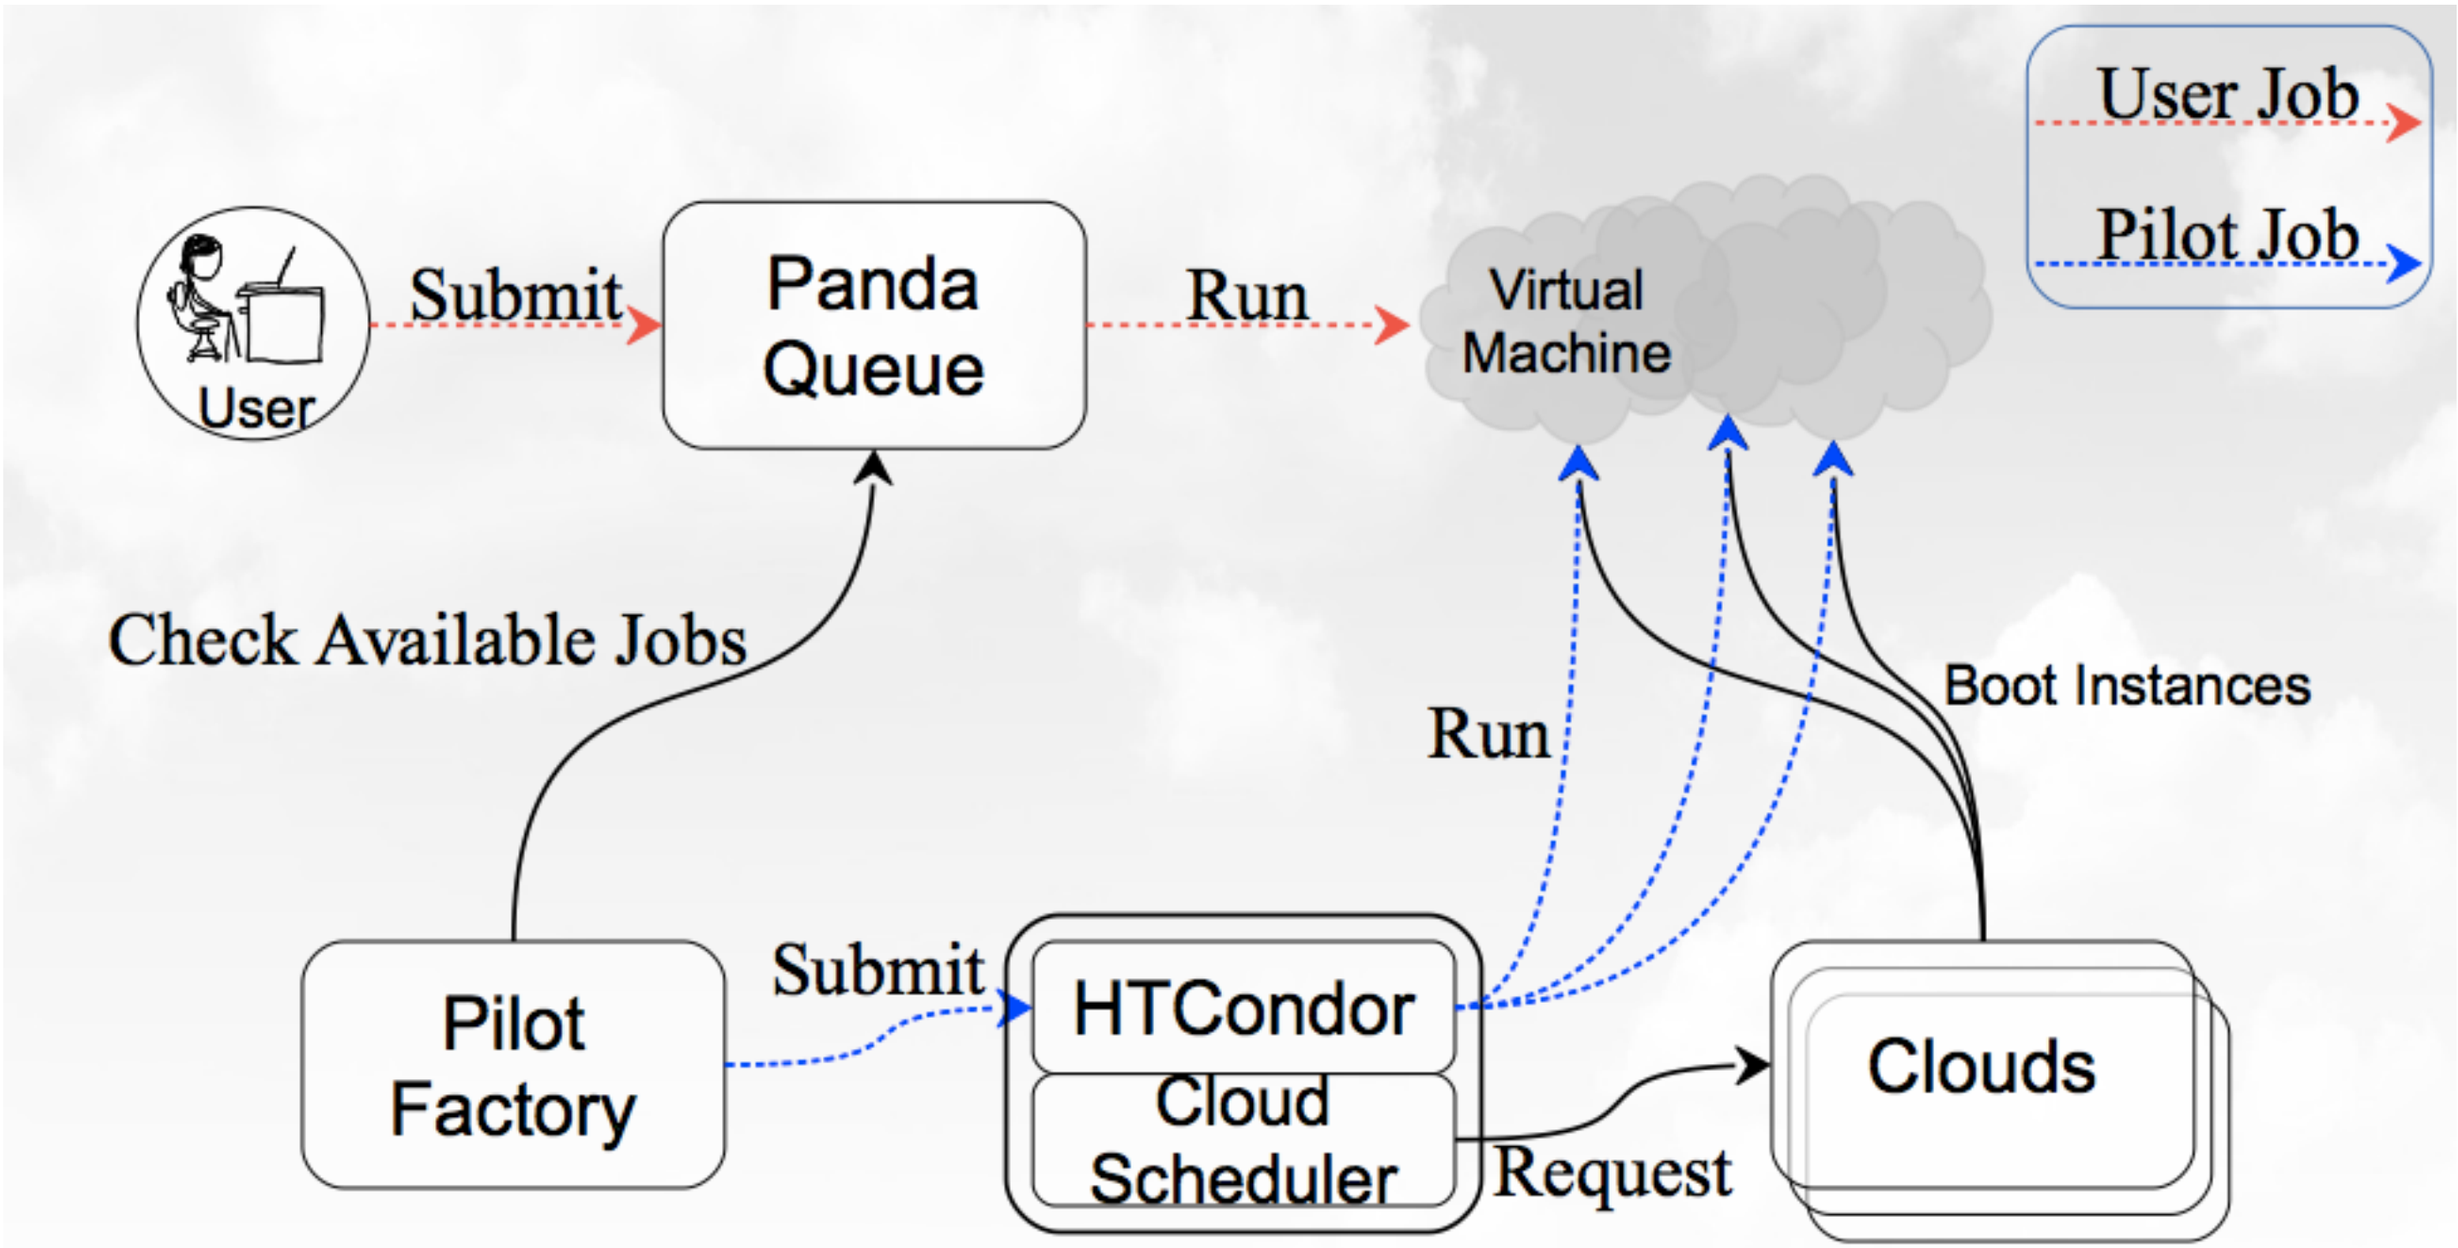

Slide: Frank Berghaus

VMs are uCernVM using Shoal for squid discovery

List of Active Squids

14 active in the last 180 seconds

#	Hostname	Public IP	Private IP	Bytes Out	City	Region	Country	Latitude	Longitude	Last Received	Alive	Verified	Access Level
1	ca01.cern.ch	188.185.163.63		35998 kB/s	Cern		Switzerland	46.2324	6.0502	2s	78h11m24s	✘	Same Domain Only
2	ca19.cern.ch	188.185.163.104		86046 kB/s	Cern		Switzerland	46.2324	6.0502	6s	54h38m19s	✘	Same Domain Only
3	ca08.cern.ch	188.185.164.64		0 kB/s	Cern		Switzerland	46.2324	6.0502	12s	0h14m8s	✘	Same Domain Only
4	t2software03.physics.ox.ac.uk	163.1.5.175		10569 kB/s	Oxford		United Kingdom	51.75	-1.25	14s	123h55m5s	✔	Global
5	kraken01.westgrid.ca	206.12.48.249	172.22.2.25	2785 kB/s	Vancouver		Canada	49.2836	-123.1041	15s	123h55m28s	✔	Global
6	ca02.cern.ch	188.185.165.173		31381 kB/s	Cern		Switzerland	46.2324	6.0502	15s	123h55m30s	✘	Same Domain Only
7	squid-test01.gridpp.rl.ac.uk	130.246.183.249		0 kB/s	Appleton		United Kingdom	51.7	-1.35	16s	123h55m23s	✘	Global
8	ip-172-31-36-99.us-west-2.compute.internal	52.11.50.231	172.31.36.99	1 kB/s	Boardman		United States	45.8399	-119.7006	17s	123h55m11s	✘	Same Domain Only
9	ca17.cern.ch	128.142.163.110		113821 kB/s	Geneva		Switzerland	46.1956	6.1481	17s	54h50m55s	✘	Same Domain Only
10	ca03.cern.ch	188.185.165.93		4368 kB/s	Cern		Switzerland	46.2324	6.0502	17s	123h55m15s	✘	Same Domain Only
11	ca00.cern.ch	128.142.135.59		98890 kB/s	Geneva		Switzerland	46.1956	6.1481	20s	123h55m13s	✘	Same Domain Only
12	atlascaq3.triumf.ca	142.90.110.68		25 kB/s	Vancouver		Canada	49.2765	-123.2177	20s	123h55m38s	✔	Global
13	pygrid-kraken.hec.lancs.ac.uk	194.80.35.16	10.41.52.16	1528 kB/s	Lancaster		United Kingdom	54.0667	-2.8333	27s	20h28m33s	✔	Global
14	ca06.cern.ch	188.184.148.164		82142 kB/s	Cern		Switzerland	46.2324	6.0502	27s	123h55m7s	✘	Same Domain Only

Early days - 2014

- First incarnation was Havana with nova networking
- Hardware was testbed quality
- Difficult to find stability, metadata service had hardware limits and was generally unreliable
- At this point adding workarounds was hard due to contextualization being buried in puppet modules located on a private repo - an organisational issue
- Debug cycle was slow

Later on

- Later upgraded to Icehouse and Neutron with production quality hardware, HA etc.

- The performance and stability was fixed but tweaking things was still cumbersome.

- Workaround was to spin-up persistent VMs via ‘nova boot’.

- Eventually moved to specific cloud-init yaml contextualization, hosted on github. Things were much more transparent.

- Flexibility was needed to workaround issues with uCernVM ganglia cloudinit module and also ganglia app version to deal with override_hostname.

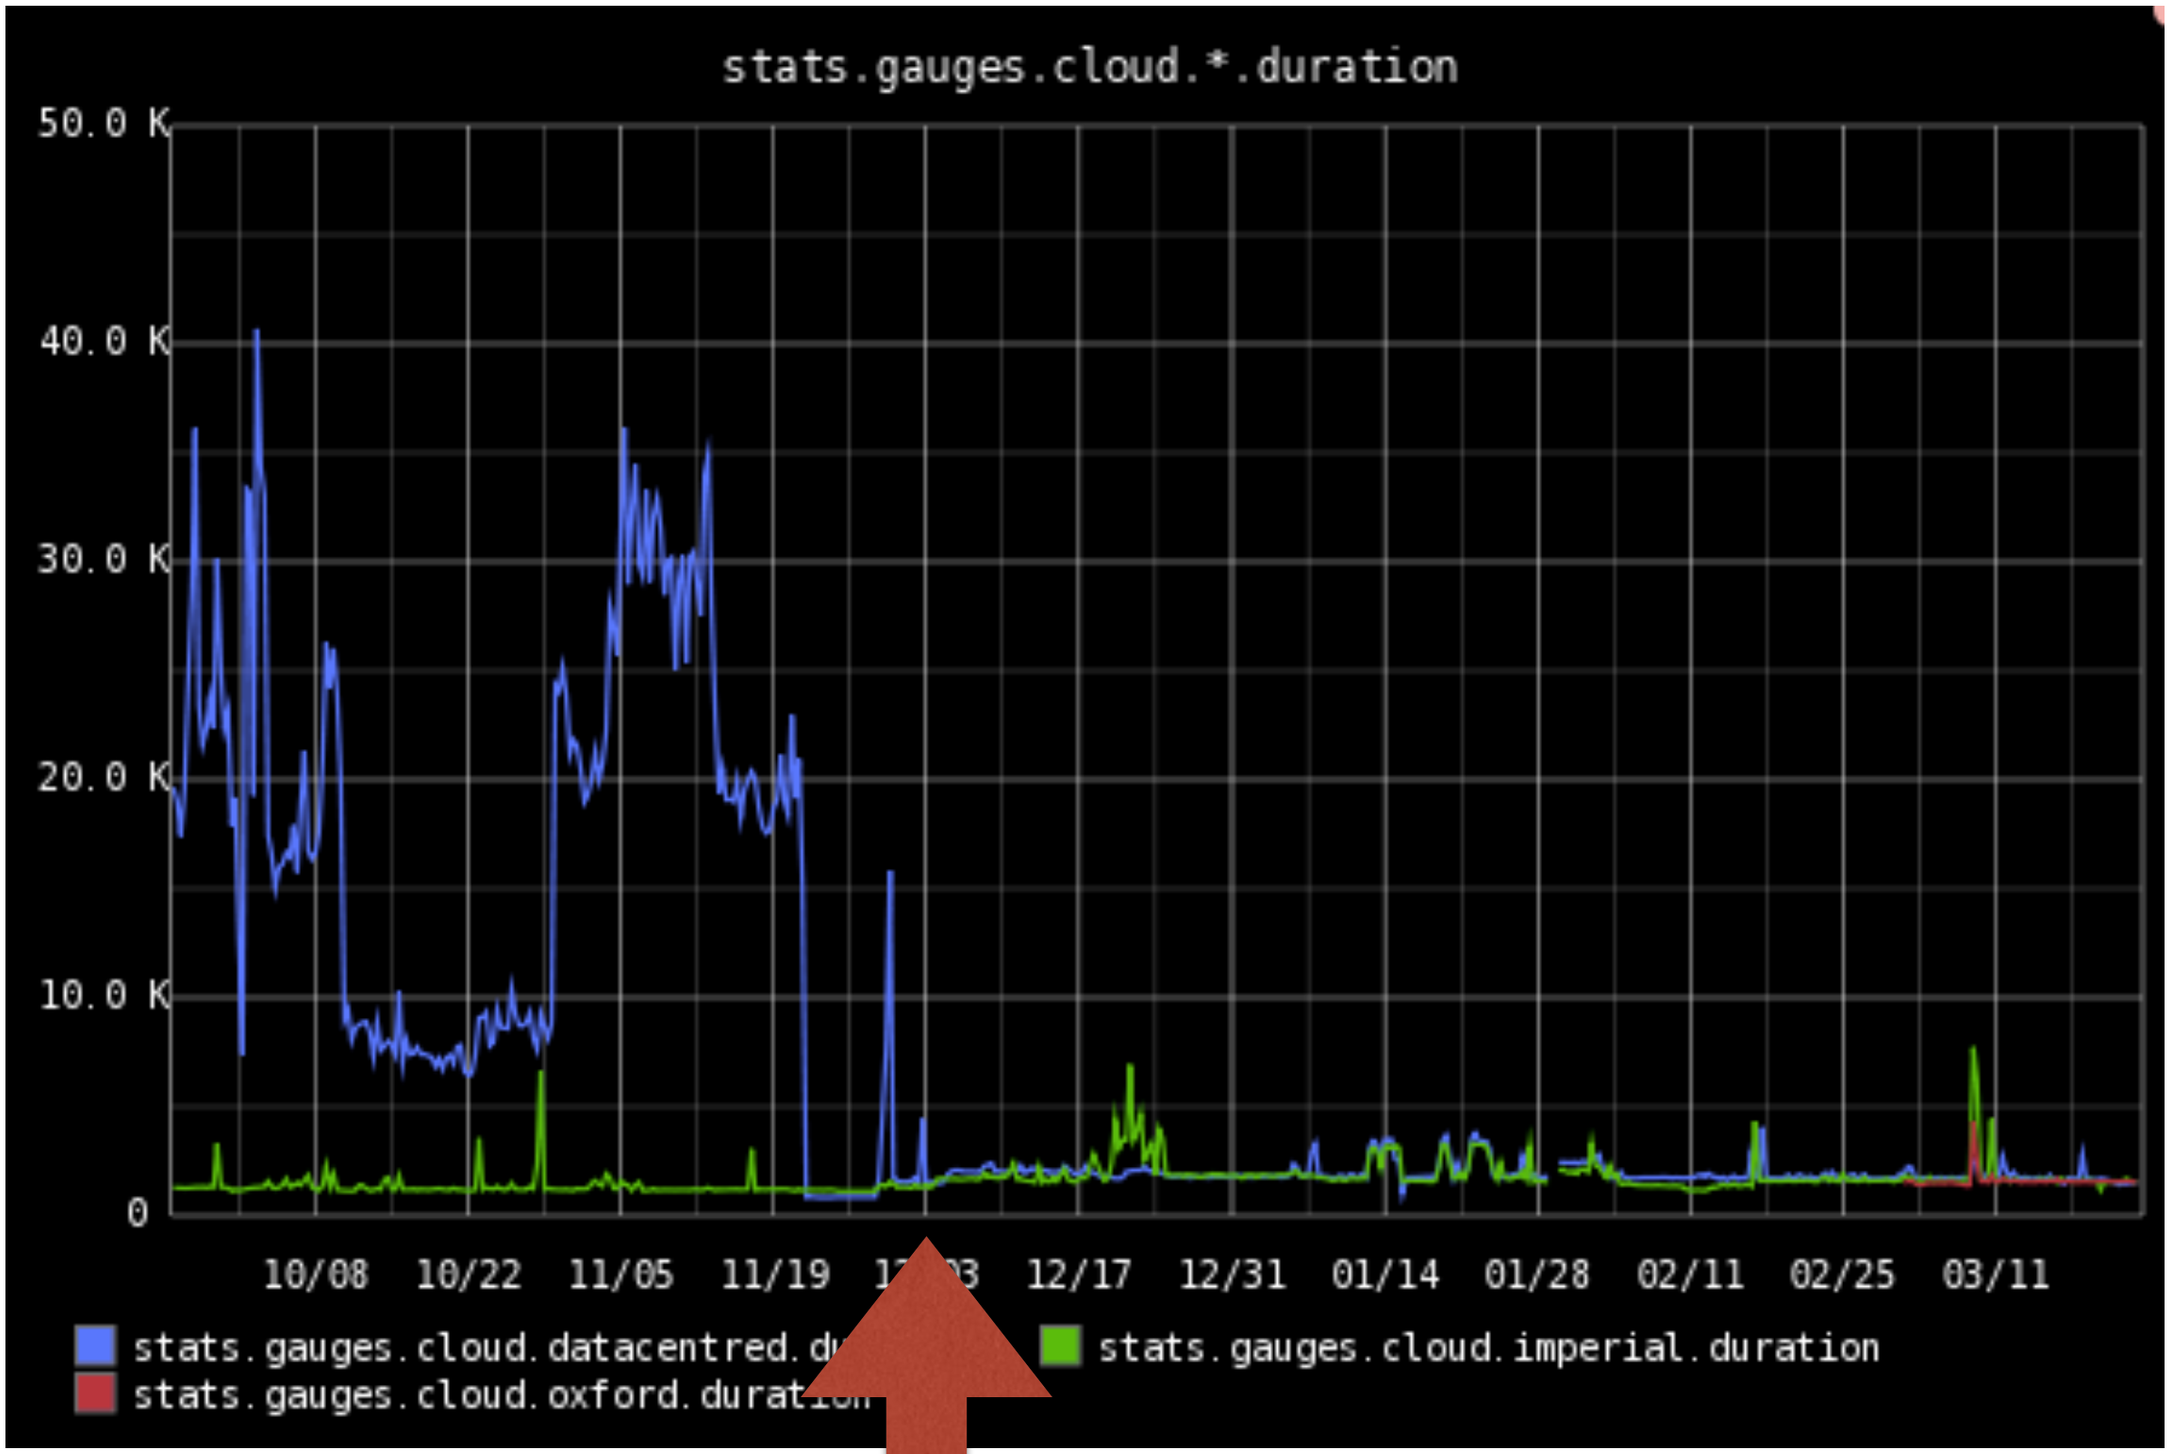

Hardware and Icehouse upgrade
Nova API response time

Comparison of a Grid and Openstack site for ATLAS production

- UKI-NORTHGRID-LANCS-HEP_SL6 (~2000 cores on Lancs grid site)
- UKI-NORTHGRID-LANCS-HEP_CLOUD (~200 cores on commercial Openstack)
- Several ATLAS Hammercloud Stress tests were run on both sites and metrics compared
- Each stress test ran for 24 hours and consisted of a continuous stream of jobs
- Jobs were mc12 AtlasG4_trf 17.2.2.2 using a single input dataset located on the grid storage ~100MB
- Metrics are compared on following slides with results as one may expect

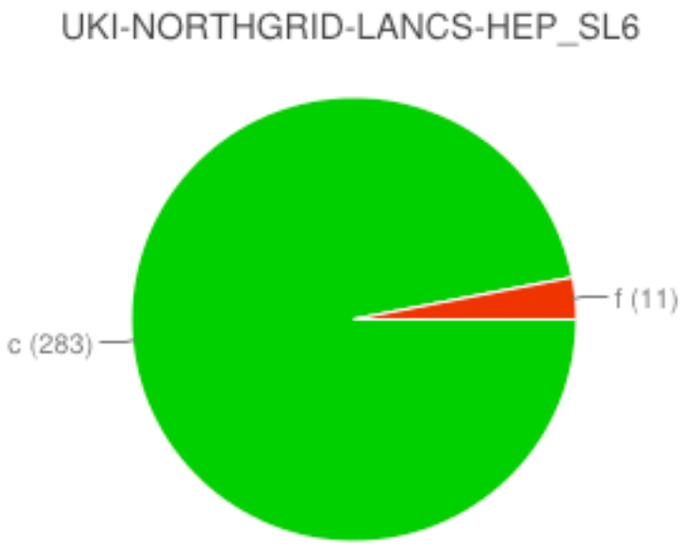

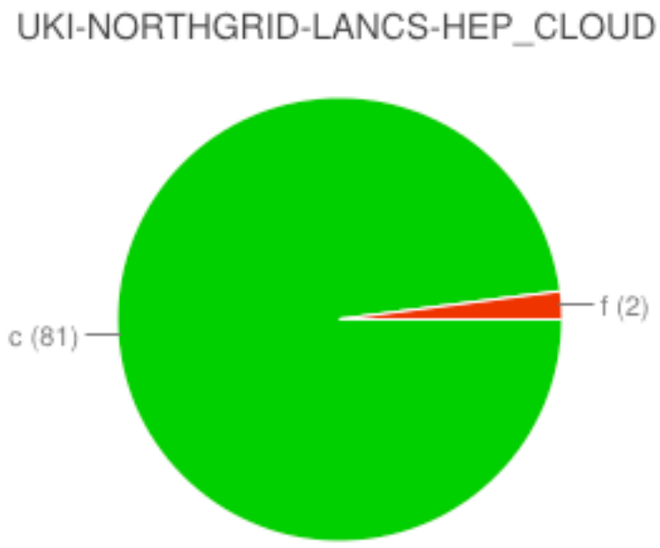

Success rate similar, grid site processed four times more jobs
283 vs. 81 jobs

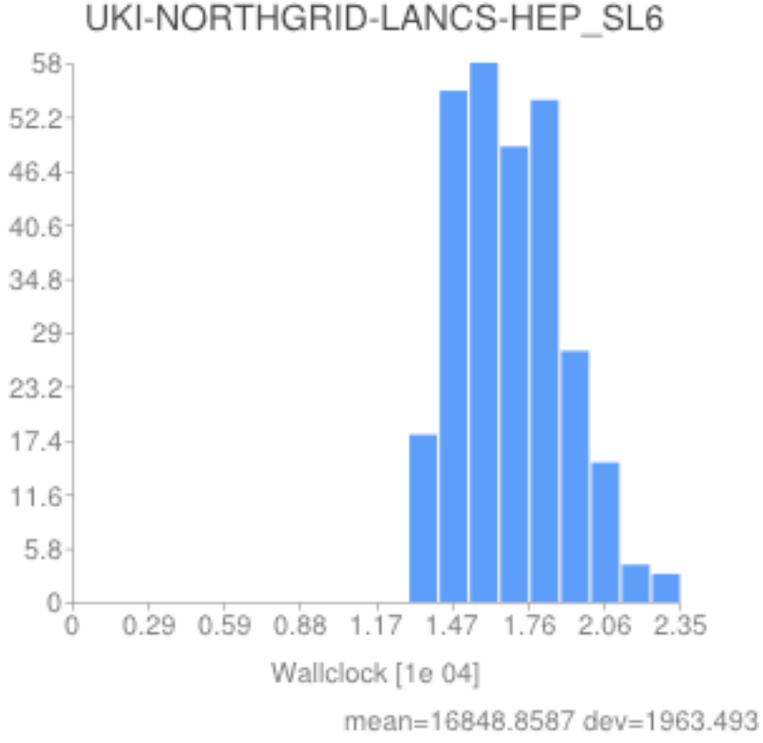

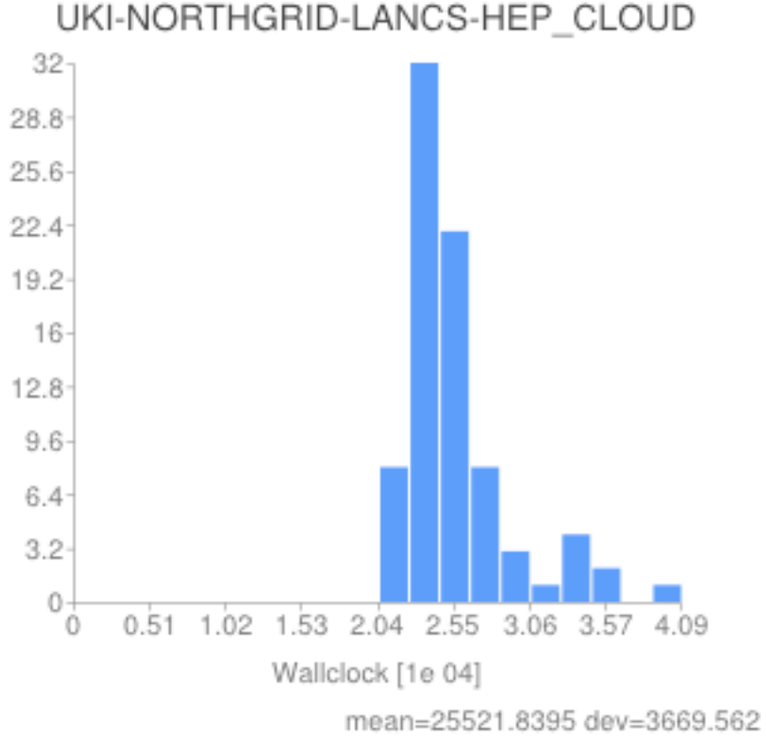

Wallclock twice as long on cloud site with greater spread in runtimes

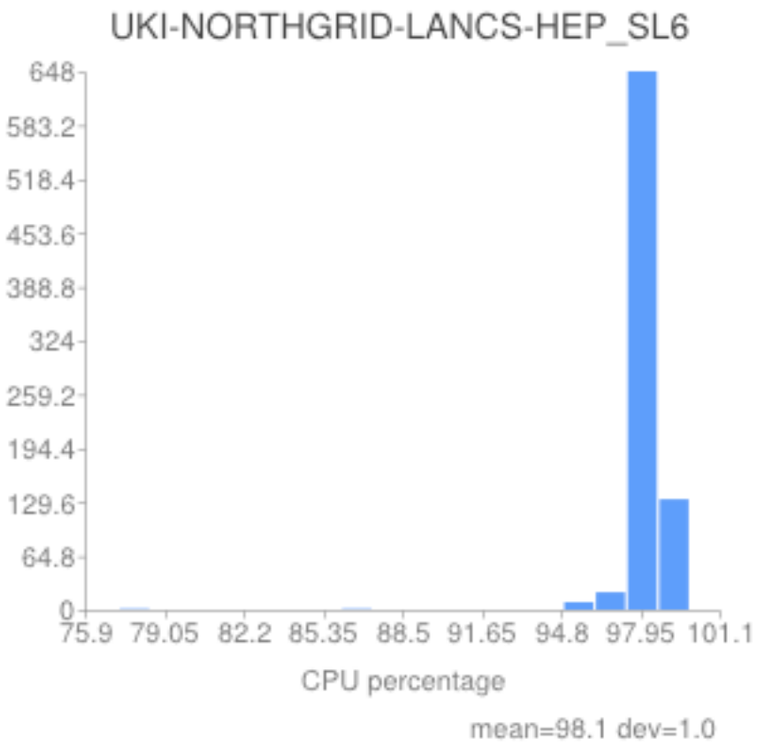

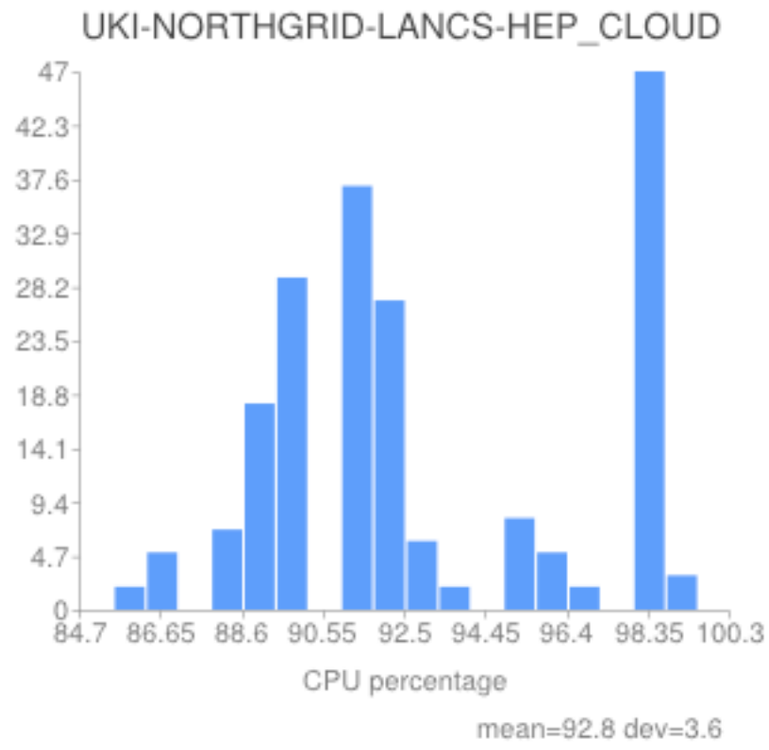

CPU percentage slightly down on the cloud site with a greater spread in efficiency

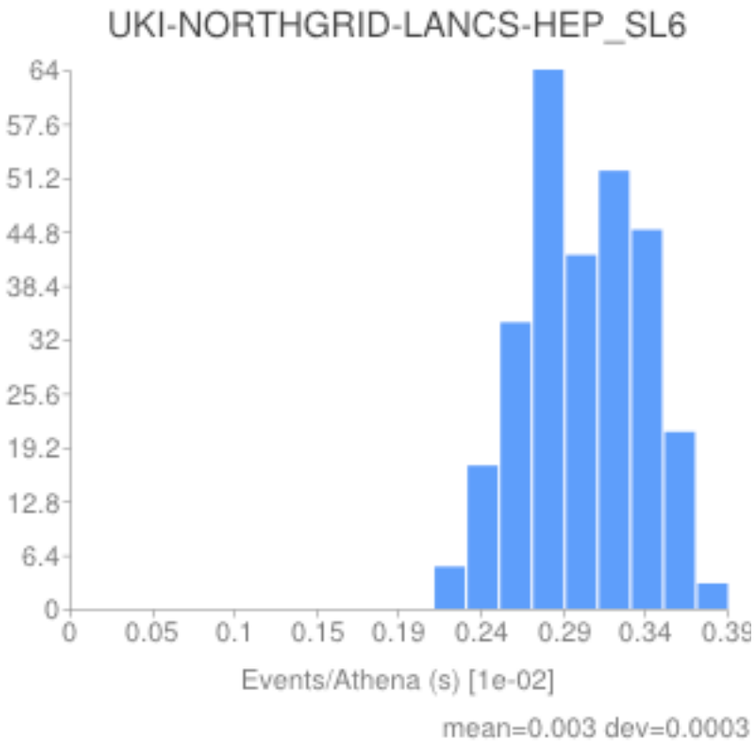

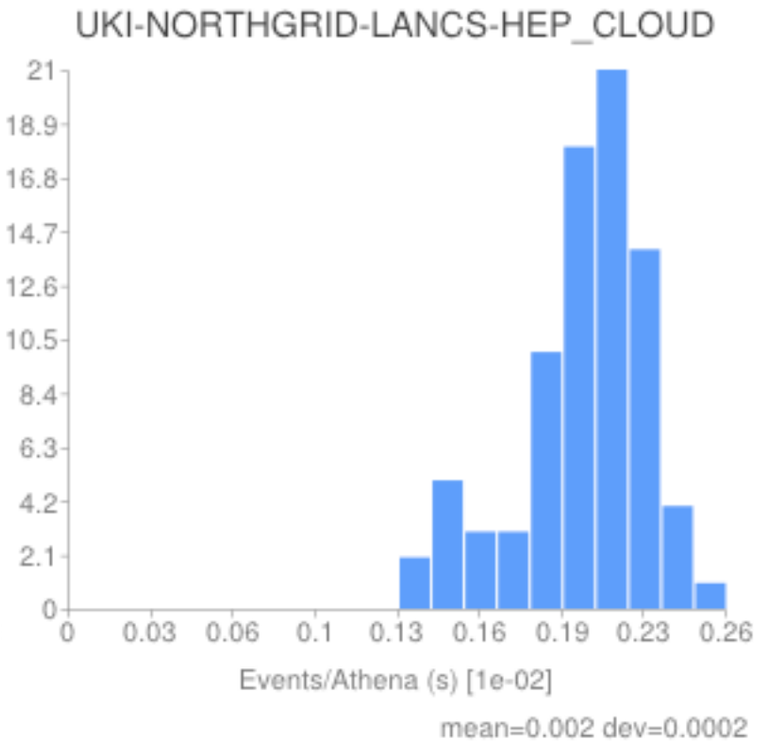

Events/Athena (s) although slower on cloud resource, the spread is similar to the grid site

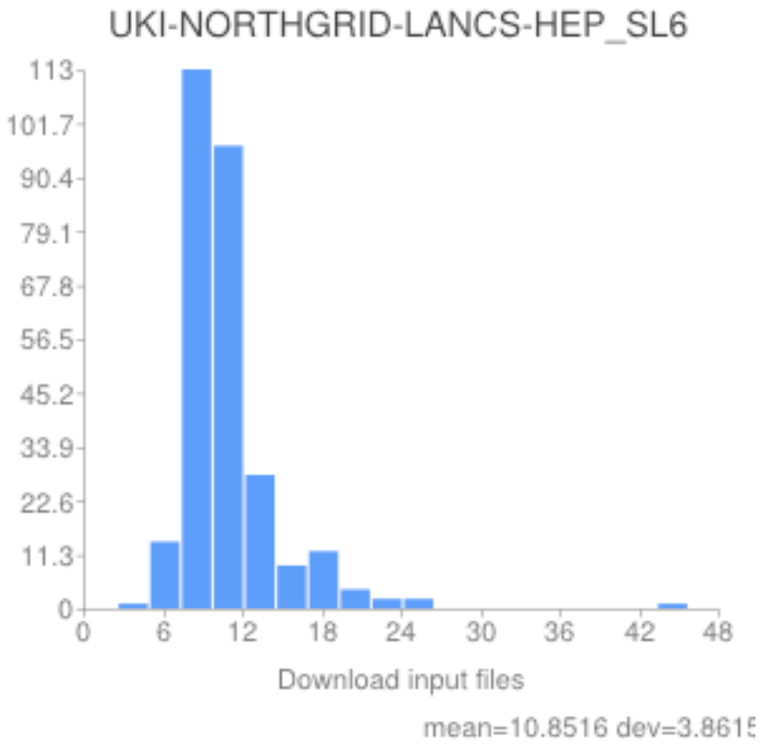

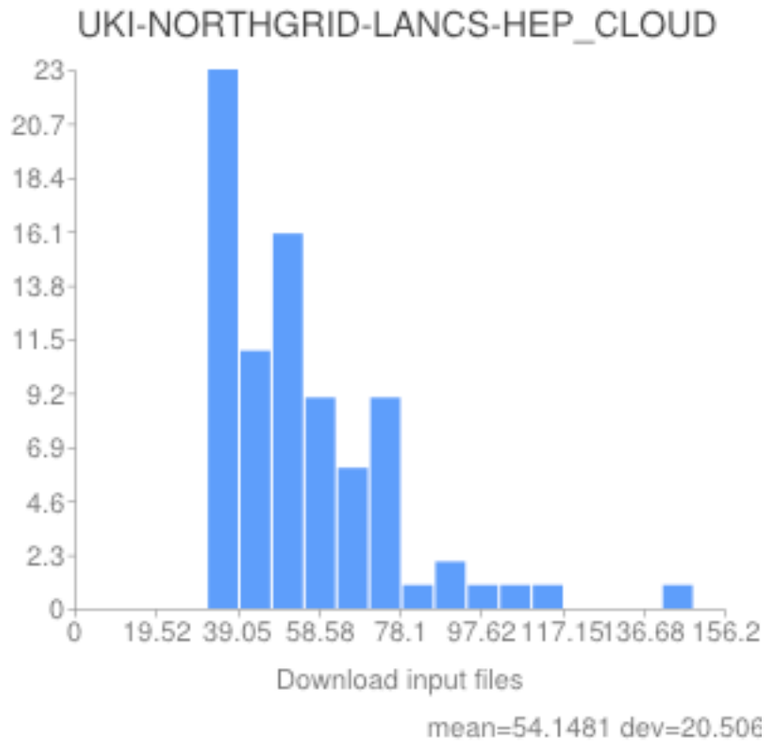

Input data stage-in time.
Remote vs. local storage.
(5x slower)

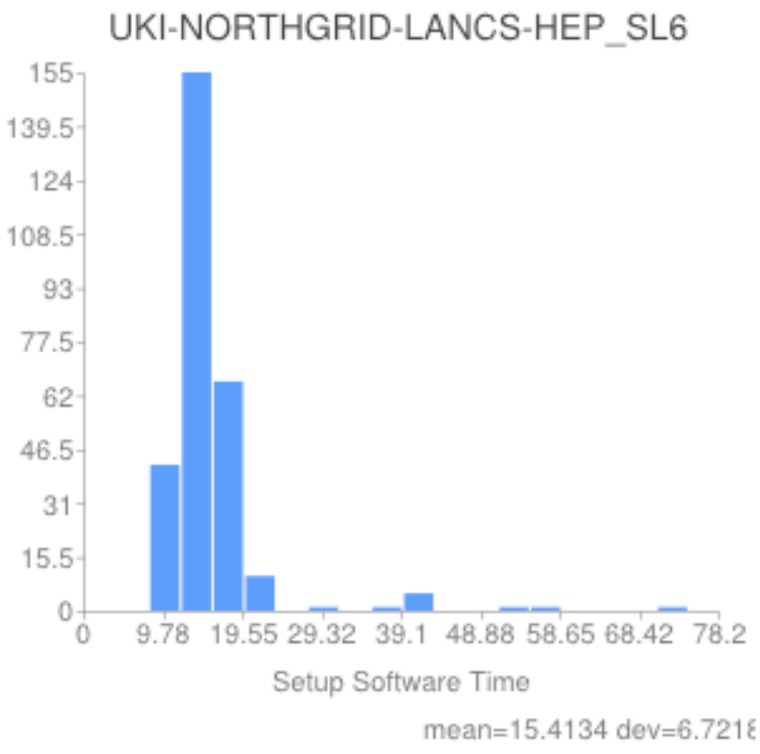

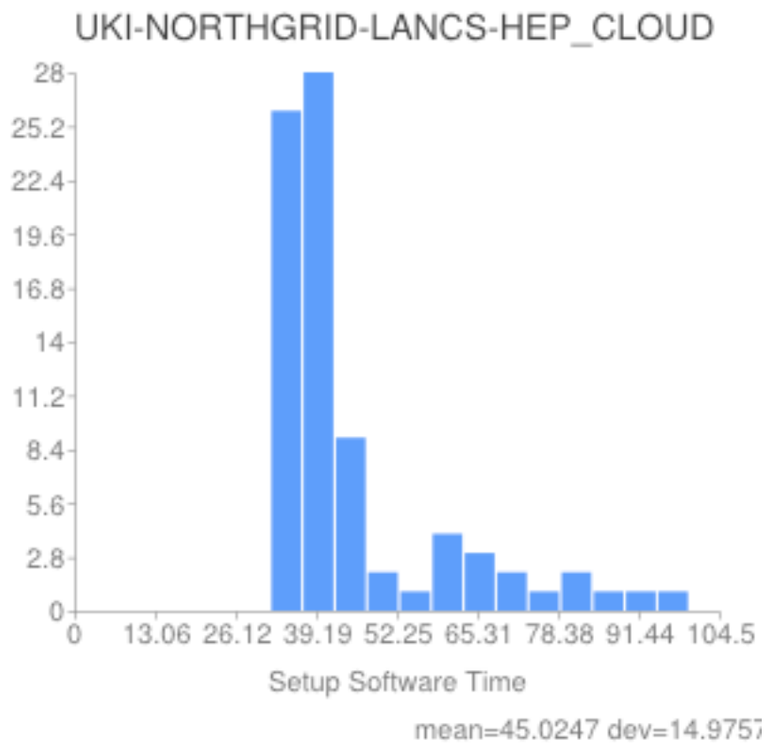

Software setup time. Via cvmfs
and influenced by squid.
(3x slower)

Summary of metric comparisons

- These are a few metrics of interest showing the differences in performance.

- The result are to be expected given the architectural differences in hardware and network between the grid and cloud site.

- No real worries although clear where more work is needed. Immediately:

    1. persistent local squid (no brainer, not solely WNs)

    2. persistent local ARC CE (simpler orchestration)

Cloud local object store

- This hosting service provides a Ceph object store
- Used as a backend for both Swift and S3 interfaces
- How can we (as users) exploit this facility?
- Various approaches are in development
  - ATLAS Event Service
  - Via FTS3 and special pilot settings
  - ARC-CE as a gateway to pre-staging data

dashboard

Running jobs

204 Days from Week 35 of 2014 to Week 12 of 2015

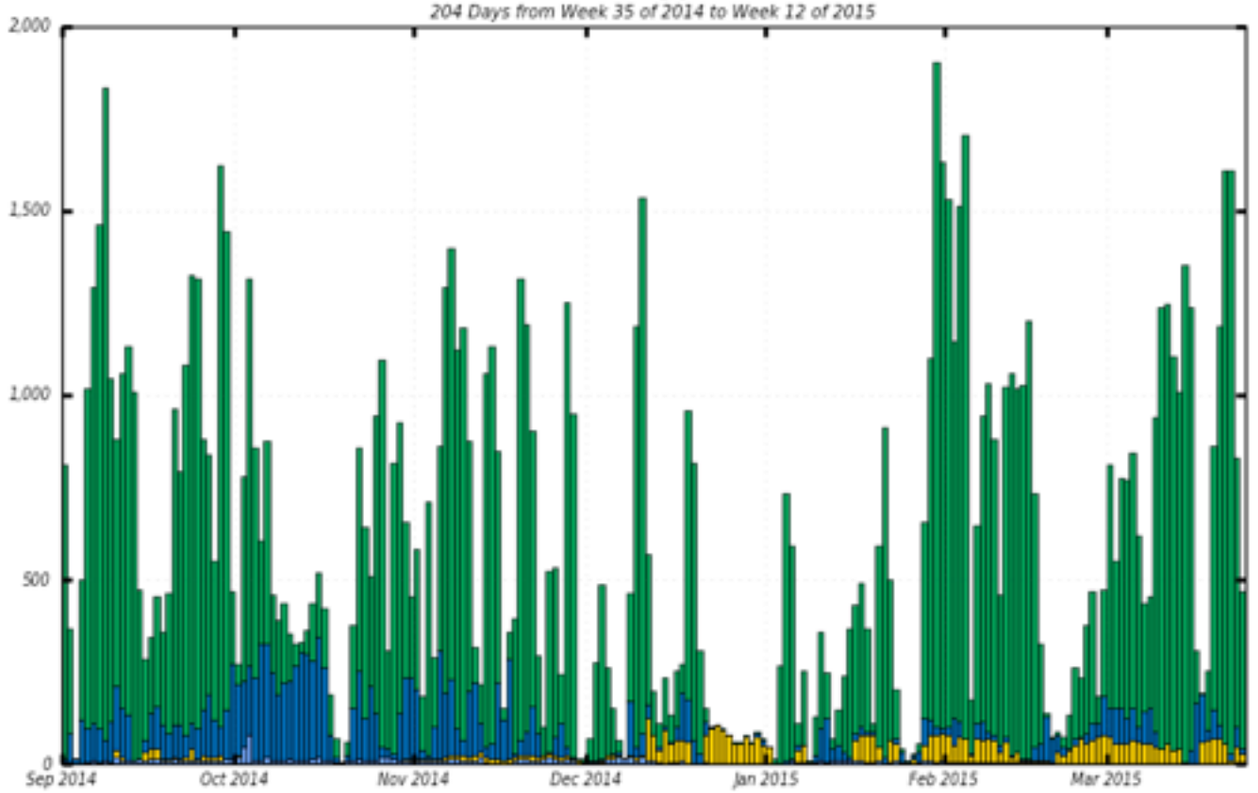

UKI-NORTHGRID-LANCS-HEP_SL6 UKI-NORTHGRID-LANCS-HEP_MCORE UKI-NORTHGRID-LANCS-HEP_CLOUD UKI-NORTHGRID-LANCS-HEP_VAC
UKI-NORTHGRID-LANCS-HEP

Maximum: 1,903 , Minimum: 5.00 , Average: 607.82 , Current: 470.00

Completed Jobs per site

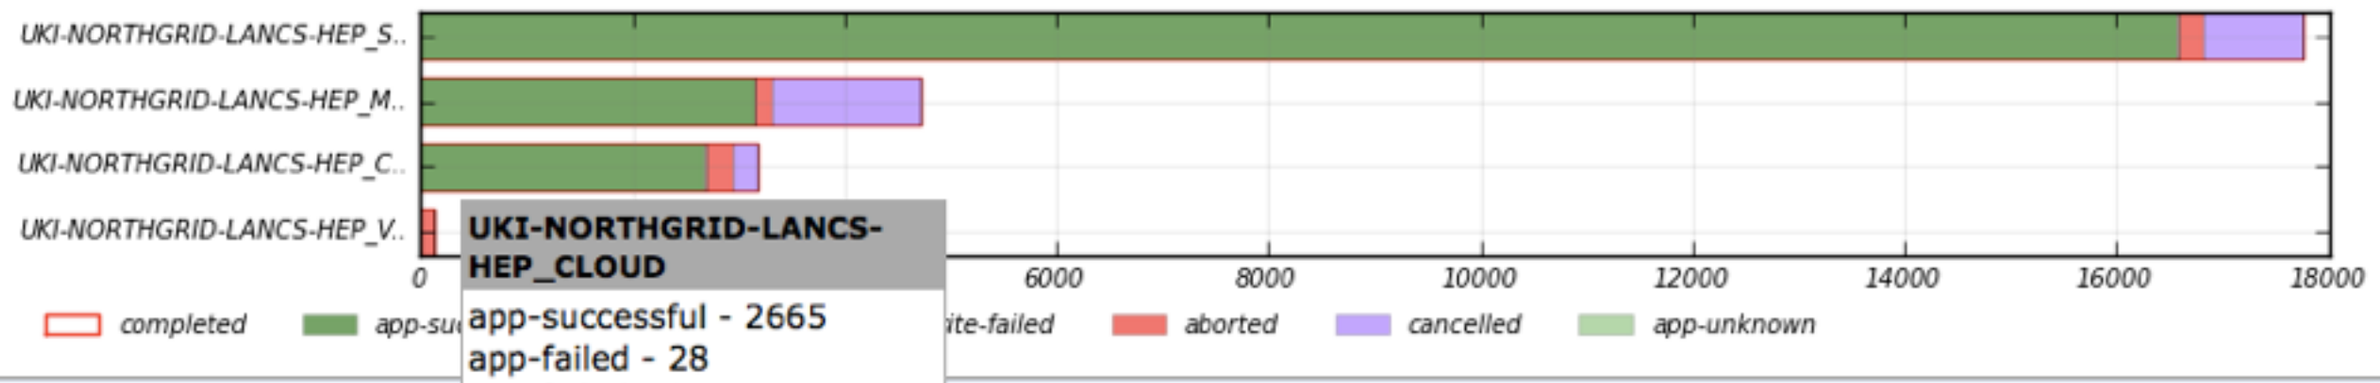

UKI-NORTHGRID-LANCS-HEP_CLOUD

app-successful - 2665
app-failed - 28
site-failed - 0
aborted - 250
cancelled - 229
app-unknown - 0
completed - 3172

All ATLAS workflows

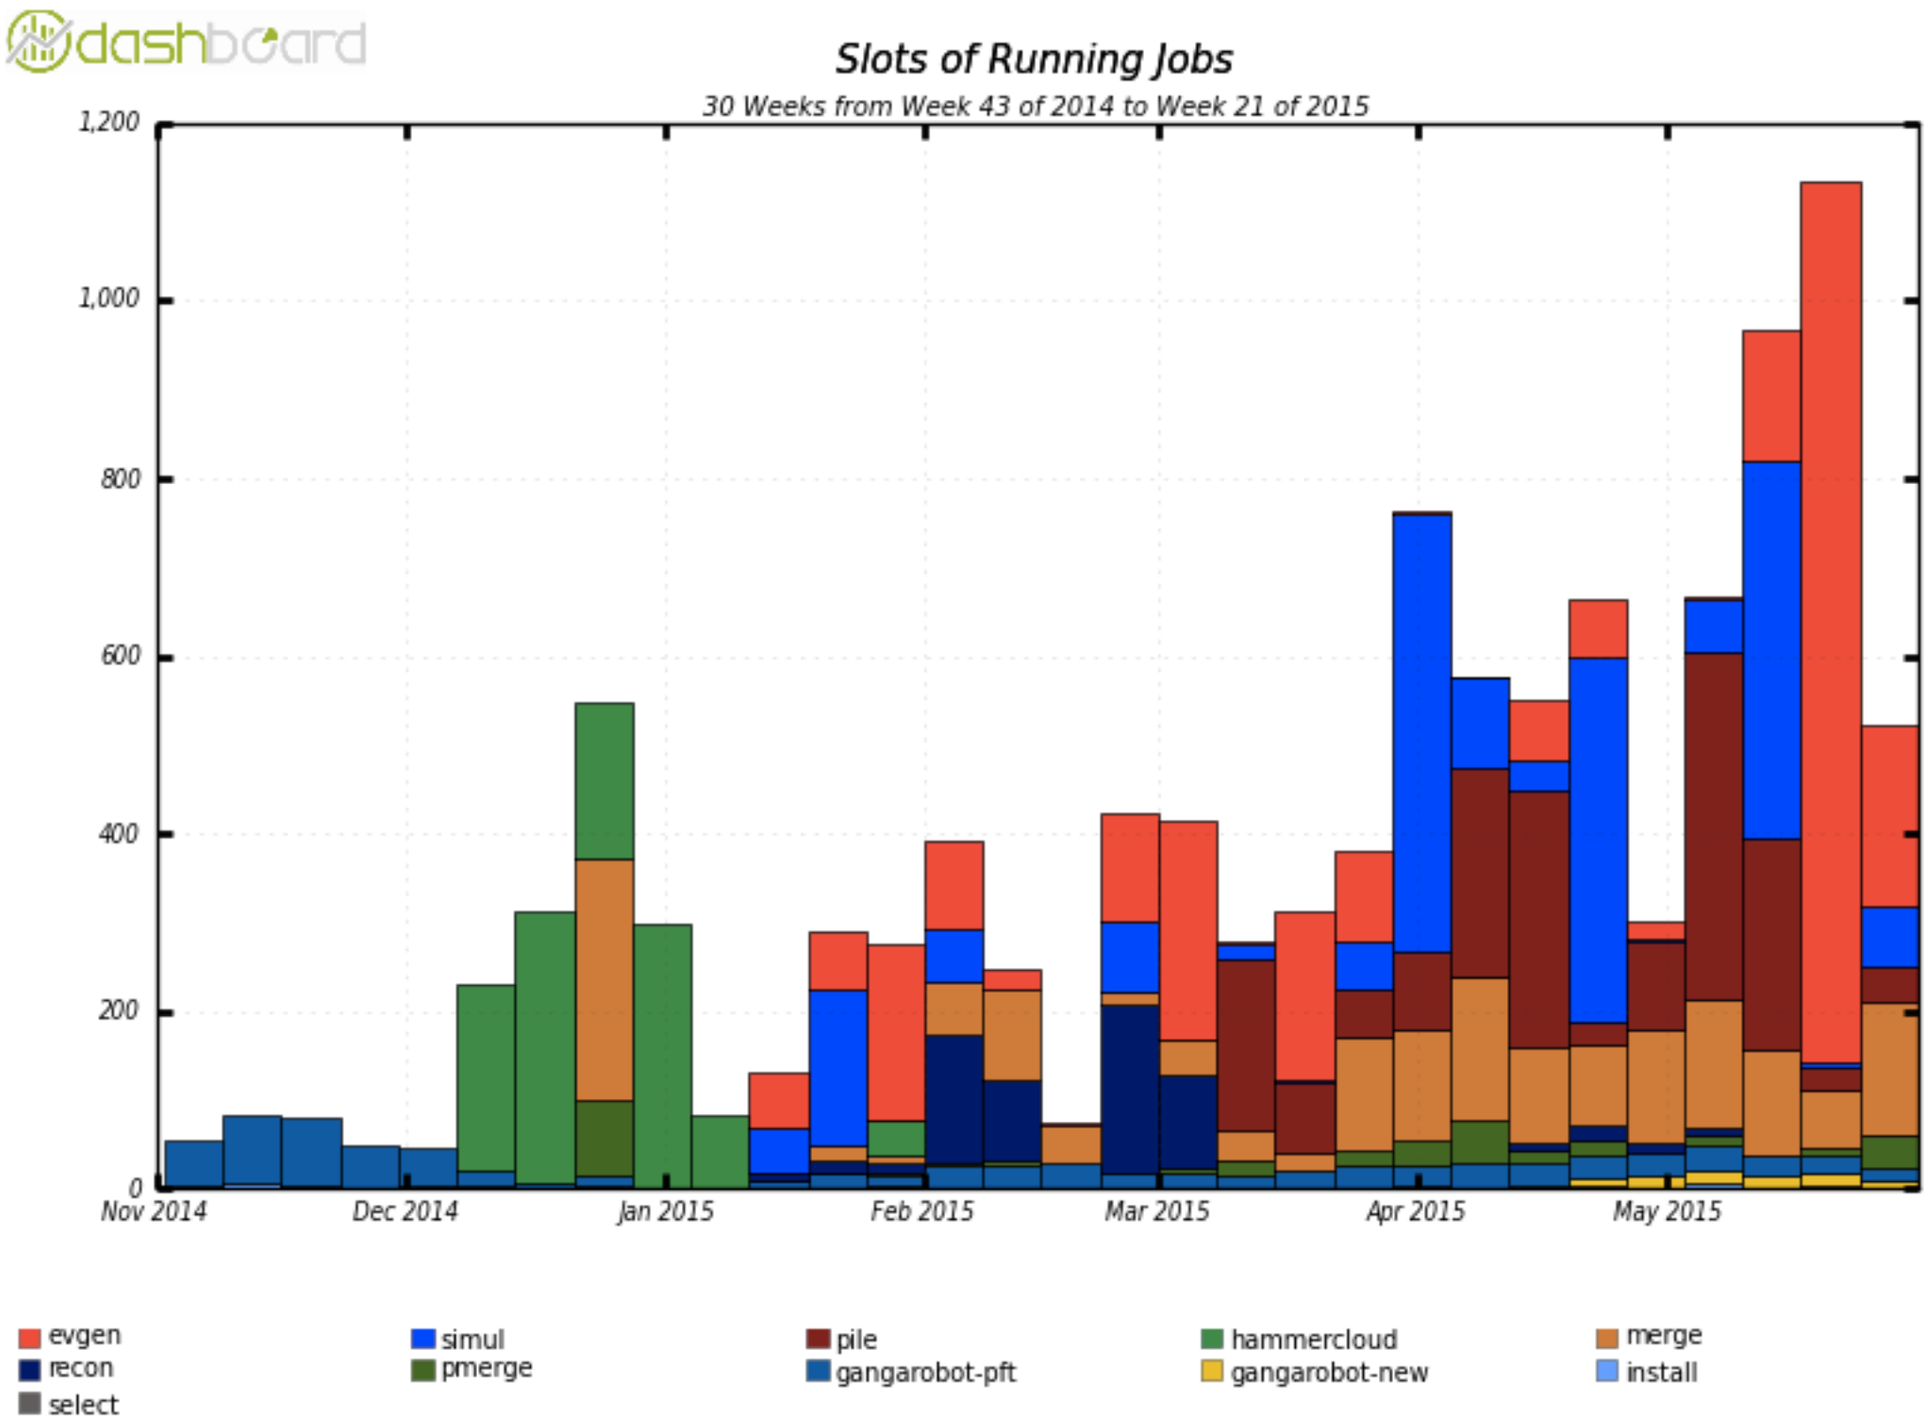

Slide: Domenico Giordano

Helix Nebula: Public & Private Partnership

Strategic Plan

- Establish multi-tenant, multi-provider cloud infrastructure
- Identify and adopt policies for trust, security and privacy
- Create governance structure
- Define funding schemes

To support the computing capacity needs for the LHC experiments

Setting up a new service to simplify analysis of large genomes, for a deeper insight into evolution and biodiversity

To create an Earth Observation platform, focusing on earthquake and volcano research

To improve the speed and quality of research for finding surrogate biomarkers based on brain images

Suppliers

Adopters

Project phases

2011

- Endorse the Common **Strategy**
- Agree on the **Partnership**
- Select **flagships** use cases
- Define **governance** model

2012-2013

- **Pilot** Phase
- **Deploy** flagships,
- **Analysis** of functionality, performance & financial model

2014 …

Towards an **open market for Science**

Slide: Domenico Giordano

The monthly activity: Ganglia source

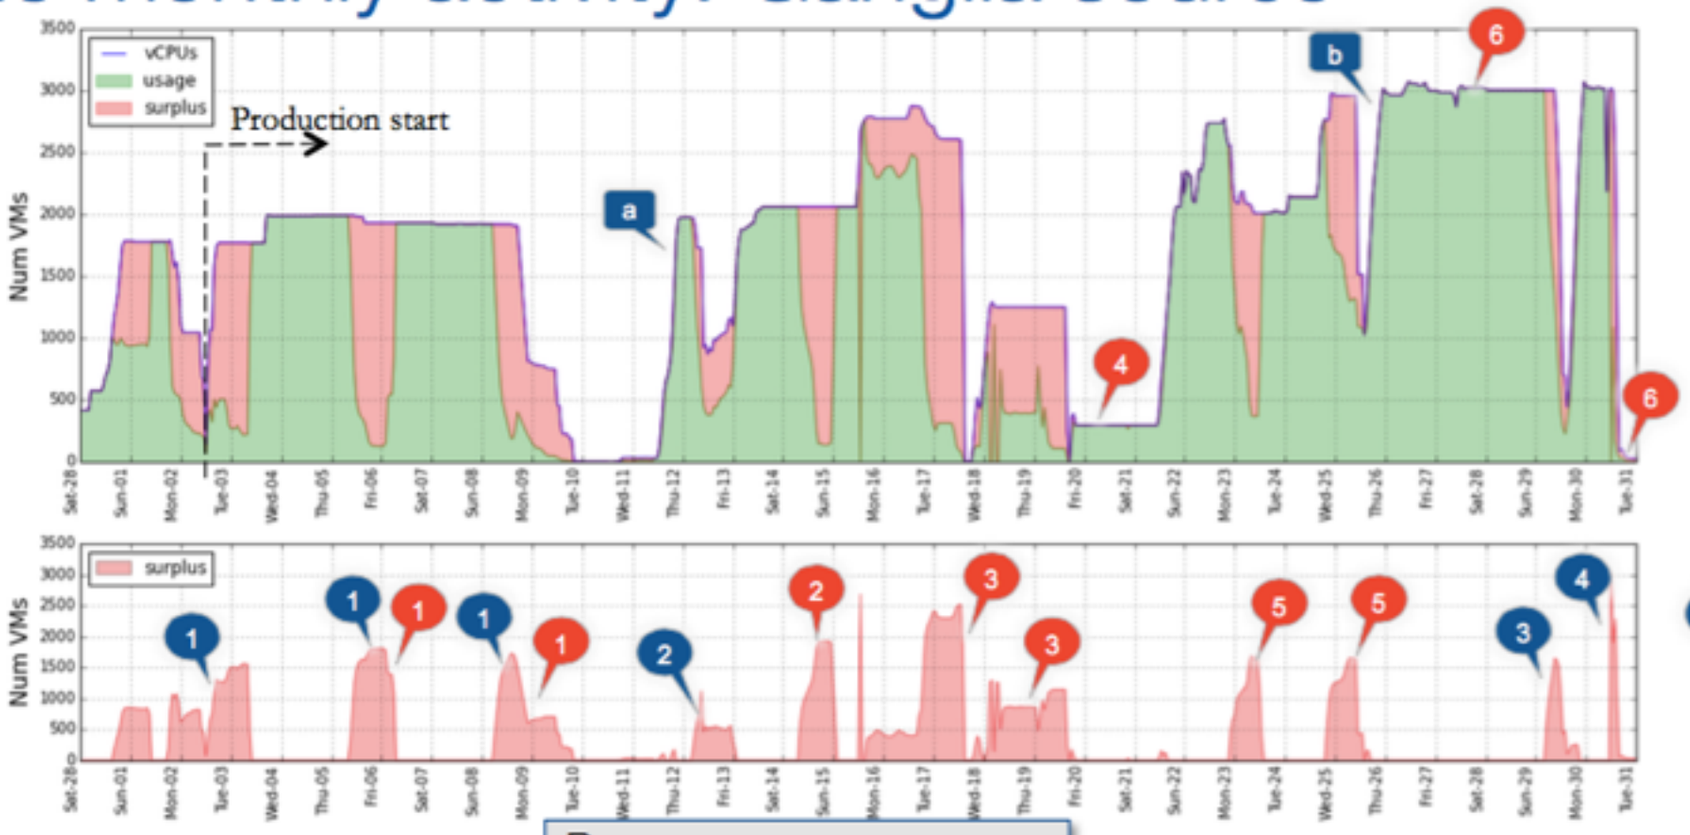

CERN

1. Task completed (no more jobs)
2. CERN network issues
3. Agent auth. cache not renewed
4. Task abruptly terminated

Resource usage:
- Effective: 77%
  - Surplus causes 10% CERN, 13% IaaS
- After improvement "b":
  - Effective: 93%

IaaS

1. VMs stuck in provisioning
2. Cloud layer reports zero VM running
3. Read-only file system
4. Stuck orchestrators: 4 faulty KVMs
5. Stuck runs: missing status from few VMs/run
6. VMs still alive after run deletion

- a Improvement: auto-scaling (up/down) based on load
- b Improvement: Orchestrator-less single-VM runs

Summary

- New cloud resources need commissioning and we have procedures to do this quickly.

- Development continues to optimize the performance and also create a recipe for a self-contained facility, relying less on outside services.

- Plenty of work required in terms of sysadmin development, operations, monitoring etc.

- More generally, HEPSYSMAN is probably interested in pushing local compute into clouds. Discuss…


Lancaster
University