London Tier 2 Status Report

Daniela Bauer
on behalf of the
LT2 site admins

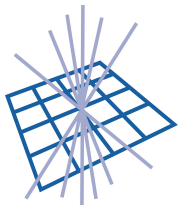

GridPP
UK Computing for Particle Physics

The London Tier 2 Grid

Sites:
Imperial College
Queen Mary
Royal Holloway
Brunel
UCL

Supported VOs
(not every site supports every VO):

LHC: Atlas, CMS, LHCb, Alice (?)

Other HEP: D0, calice, ilc, H1, t2k, supernemo, super-b, mice, pheno, zeus, geant4, totalep

Other: biomed, cedar, fusion, camont, dteam, ops, ngs, gridpp, londongrid

GridPP
UK Computing for Particle Physics

Top 10* Supported VOs

*According to CESGA accounting portal

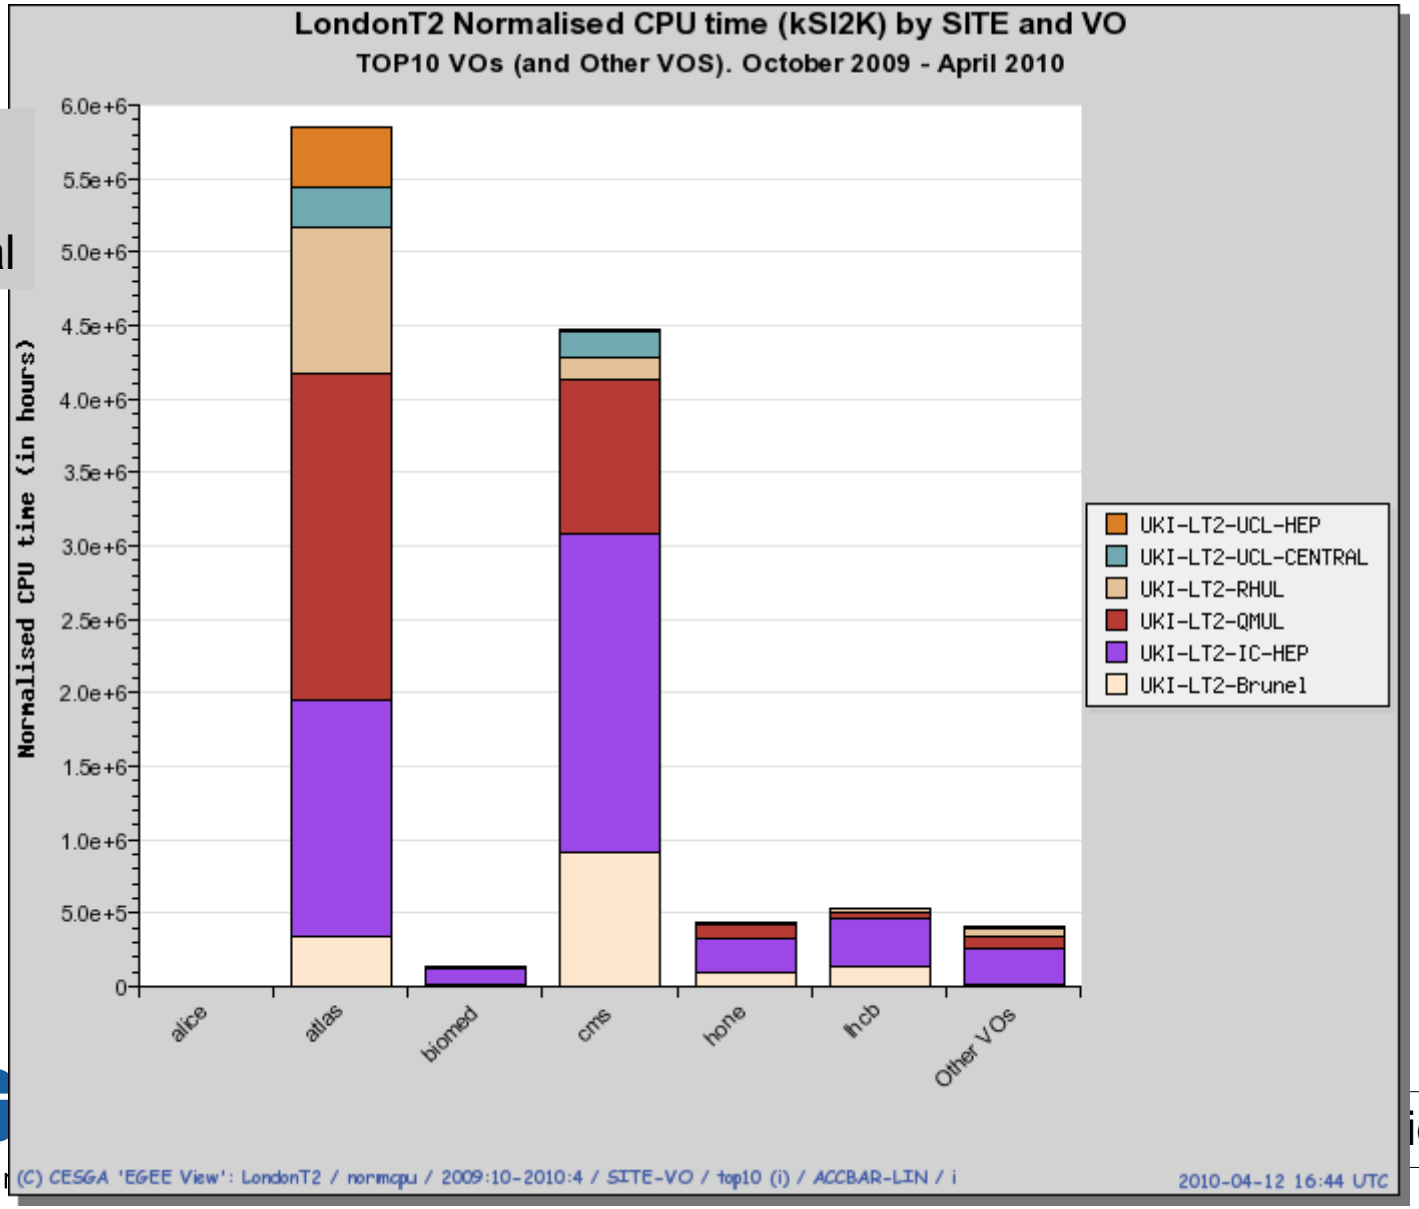

Availability

Past 6 months, for UCL-CENTRAL and RHUL 4th Quarter 09

Site	Ops	CMS	Atlas
Brunel	97%	96%	98 %
IC-HEP	90%*	91%	95 %
QMUL	91%	86%	99 %
RHUL	81%	N/A	N/A
UCL-HEP	93%	N/A	94%
UCL-CENTRAL	43%	N/A	N/A

* on the up, since I turned off all automatic glite updates.

GridPP
UK Computing for Particle Physics

Imperial College

(Mona Aggarwal), Daniela Bauer, Dave Colling, Simon Fayer (dedicated grid sysadmin), Janusz Martyniak (RTM), (Barry MacEvoy), Duncan Rand, Stuart Wakefield (CMS)

Farm:
125 Dell Poweredges x 8 Cores
~40 ageing Viglen x 4 Cores
= 1160 job slots

Storage:
605TiB (or to be exact 634984935MB!) in dCache

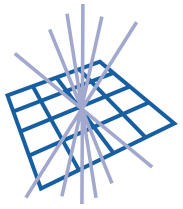

GridPP
UK Computing for Particle Physics

Imperial College

Since Gridpp23:

- Migration to SL5 completed.

- Stable running until recently:

Jobs constantly accessing the home directories triggered an NFS/gass_cache bug that caused the home directories to become corrupted on a regular basis. Currently working on a fix.

- All Storage commissioned
- dCache tuning
- Started enabling space tokens for Atlas

Plans:

Hardware:

9 HP with 2xIntel Xeon Westmere quad-core processors 24 GB RAM

Software:

cream-ce and sge (work in progress)
new WMS (working) and LB
Soon: SCAS or Argus

UCL

Gianfranco Sciacca (main admin), Ben Waugh (coordination),
IS-Research Computing systems team at UCL

UCL HEP – UCL Central 'Merger' (1)

Before merger:

UCL-HEP:

- 112 cores on SLC5 since March 2010 (hep cluster) fronted by a lcg-CE and run by HEP at UCL. New purchases in the near future.
- Few old cores on SLC4 still in use serving misc minor VOs
- 50TB on DPM run by hep at UCL (mainly for ATLAS)

UCL-CENTRAL:

- A nominal 5% fair-share (not a hard limit) of the 2560-core central cluster on SL4 fronted by a lcg-CE. This is run by central IS at UCL.
- 50TB on DPM run by IS at UCL (various VOs)

UCL

UCL HEP – UCL Central 'Merger' (2)

Current status:
UCL-CENTRAL lcg-CE transferred to UCL-HEP site (Nov 09). Still run by IS at UCL, with HEP acting as intermediary for WLCG requirements and operations.

Cluster migration to SL5 was scheduled to have completed by now, but unexpected problems with vendor (Clustervision) are holding this back: New schedule May - July.

UCL

UCL HEP – UCL Central 'Merger' (3)

Last phase of transition (in progress):

UCL-CENTRAL storage phased out.
UCL-HEP to commission an equivalent amount (or more)

New CE fronting central cluster will be run by HEP

Complete decommissioning of UCL-CENTRAL as site, Statement by UCL-CENTRAL: “UCL Research Computing remains committed to continuing support for GridPP as the current platform for UCL's HEP research, regardless of lack of direct funding from recent rounds. It is recognised that the lack of funding in the current round is a direct consequence of performance over the last 18 months.”

UCL HammerCloud Test

Overall | Sites | Metrics | Other

Summary

state	id	host	clouds	start time (CET)	end time (CET)	submitted jobs
completed	1110	voatlas73.cern.ch	UK	2010-02-17 12:00:00	2010-02-18 12:00:00	571

Input type: FILE_STAGER
Output DS: user09.JohannesElmsheuser.ganga.sitetest.UK.muon.20100217.1.[sitename]
Input DS Patterns: mc08.*merge.AOD.e*_s*_r6*tid*
Ganga Job Template: /data/gangarobot/hammercloud/inputfiles/muon1531/muon1531_lcg.tpl
Athena User Area: /data/gangarobot/hammercloud/inputfiles/muon1531/source.1531.tar.gz
Athena Option file: /data/gangarobot/hammercloud/inputfiles/muon1531/MuonTriggerAnalysis.py

View Test Directory (for debugging)

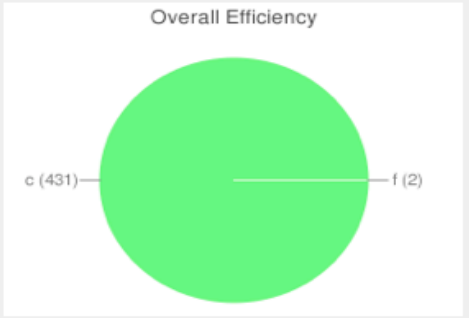

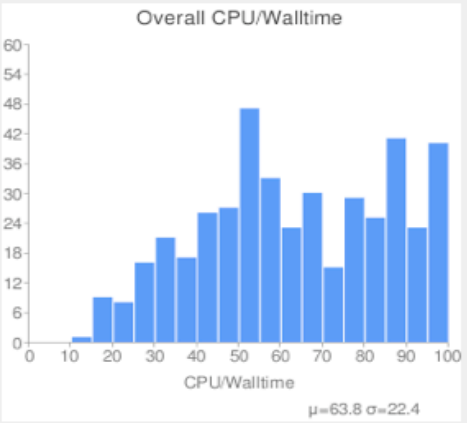

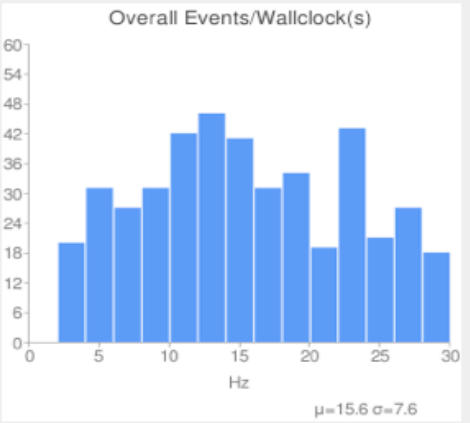

Sites

Site	Submitted jobs	Running jobs	Completed jobs	Failed jobs	Num datasets per bulk	Min queue depth	Max running jobs	Resubmit enabled	Resubmit force	View jobs
UKI-LT2-UCL-HEP_MCDISK	60	45	431	2	20	50	1000000	yes	no	»
Total	60	45	431	2						»

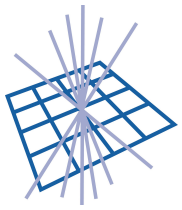

QMUL

Chris Walker, Alex Martin

Worker nodes:
64* 2*4core Xeon E5420 (2.5GHz) 8 GB RAM
2 * 2*4core Xeon E5420 (2.5GHz) 32 GB RAM
160 *4core Opteron 270 4GB RAM
128*2core(2 thread) Xeon (2.8GHz) 2GB RAM
Total 1424 cores

The old 2.8GHz Xeon machines run SL4, and as they are 32 bit.
The rest runs SL5
(10 of the E5420 machines are providing a residual service for local users on SL4).
Storage is SL5

Storage:
290 TiB storage on lustre_0 and 30TiB on Lustre_1 (for testing)

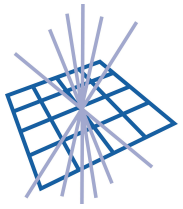

GridPP
UK Computing for Particle Physics

QMUL – SuperB support

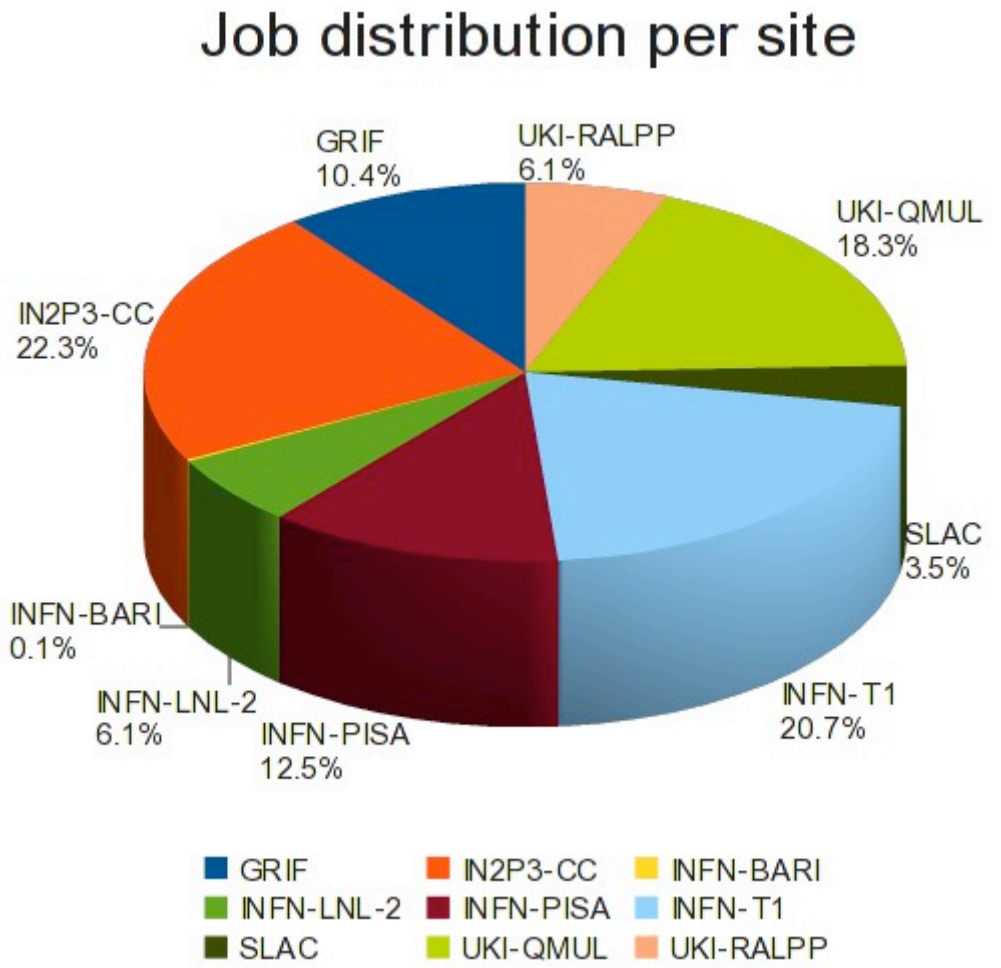

GridPP
UK Computing for Particle Physics

QMUL – improved results HC tests

October 2009

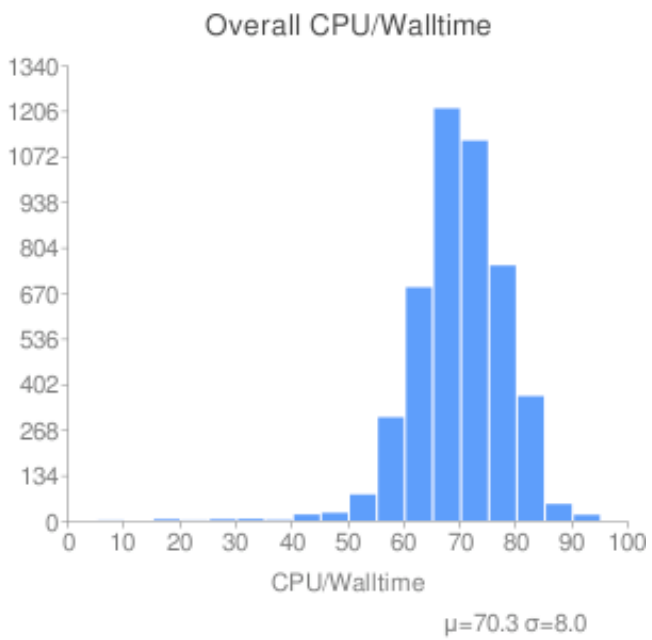

369 million events/24h

April 2010

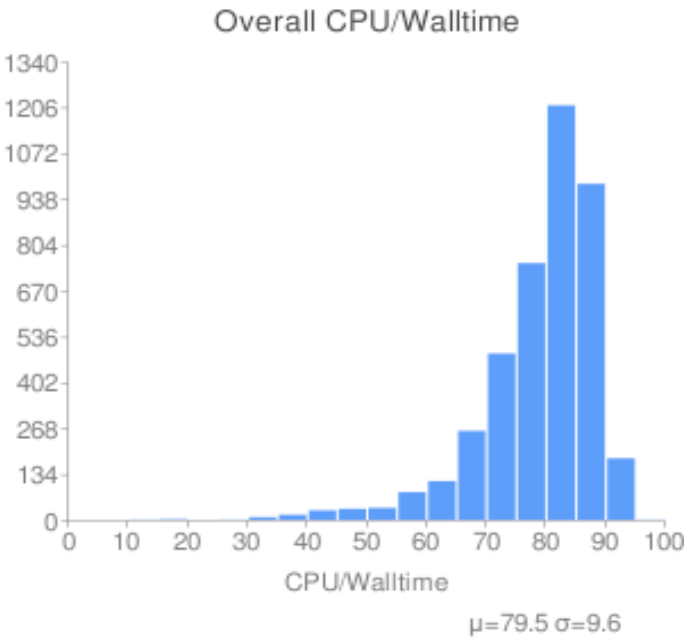

433 million events/24 h

GridPP
UK Computing for Particle Physics

QMUL Plans

Hardware purchases: CPU soon, Disk in the Autumn.
Inquiring about a network upgrade.

Software:

Storm 1.5 (ideally on SL5) - this will provide checksum support
gridFTP servers on separate hardware

Cream - after Storm - or when Imperial get it working.
SCAS/Argus and Glexec

Brunel

Raul Lopes, Henry Nebrensky

Since GridPP23:

- December 2009: moved new storage to new data center
  460TB online + 140 TB available to go online

- February 2010: created new cluster in new data center

- April 2010: finished upgrading to SL5

- Brunel has now 592 cores (4400 HepSpec06)

Brunel - Plans

- preparing to buy another 240 cores

- looking at an upgrade of the internal network to 10 Gb/s

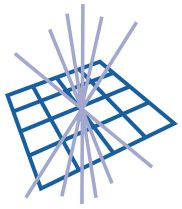

GridPP
UK Computing for Particle Physics

Royal Holloway

Govind Songara, Simon George

Royal Holloway cluster was moved from Imperial to RHUL.

* Cluster moved from IC to RHUL, hardware and data intact
* SL5 o/s upgrade complete (Clustervision).
* Currently working on complete middleware reinstall and learning new Cluster Manager.

* Middleware status:
  CE (SL4) Lcg-CE set up
  WN (SL5) set up
  SE (SL5) set up.
  DPM pool nodes (SL5) close to set up
  sBDII (SL5), MON (SL4) - not affected by upgrade

RHUL

Work in progress:

1. Complete SE + storage nodes setup.
2. network bandwidth limit test with computer centre & CMS
3. squid
4. backups
5. SAM testing -> site back into production
6. VO software install
7. ATLAS testing (production, analysis) -> back up as an ATLAS site

After that, on new hardware:

- CREAM CE (SL5)
- SCAS/ARGUS (SL5)
- MON (SL5)

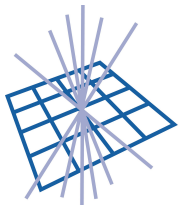

GridPP
UK Computing for Particle Physics

RHUL Hammer Cloud

Before site move.

125 million events

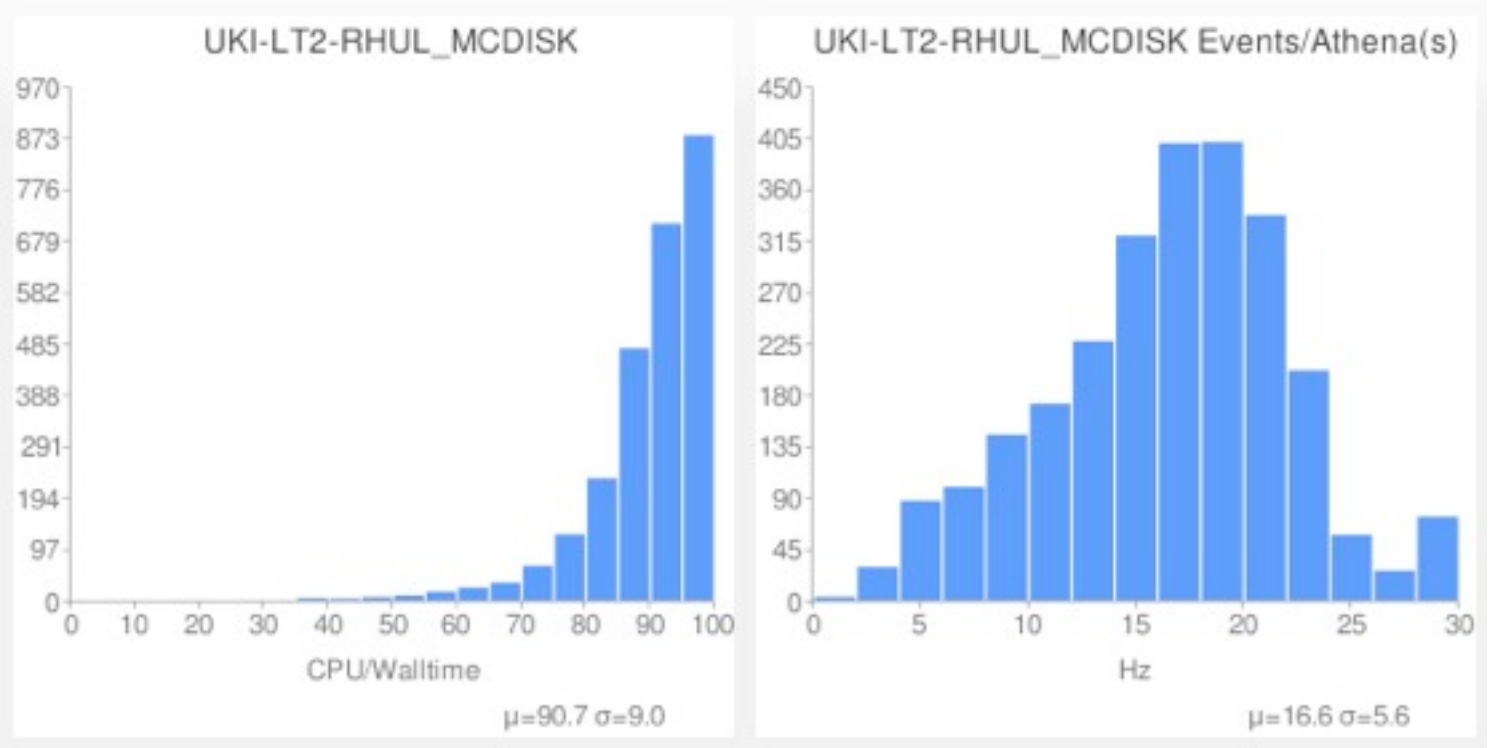

Conclusion

Brunel, Imperial, QMUL, UCL are ready for data and expanding.

RHUL move to be completed within the next couple of weeks to rejoin the London Tier 2.

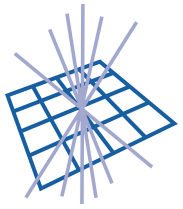

GridPP
UK Computing for Particle Physics

Clustervision's "Bright Cluster Manager" (CM)