

**GridPP**

UK Computing for Particle Physics

# London Tier 2 Status

Daniela Bauer  
Imperial College


**Imperial College  
London**

## The London Tier2 Grid

- Imperial College
- Queen Mary
- Royal Holloway
- Brunel
- UCL
  - HEP
  - Central



### LondonT2 Normalised CPU time per SITE ALL VOs. April 2009 - September 2009



## Imperial College

- New switches - this should enable us to integrate 300 TB of storage seamlessly into the cluster.
- dCache - Lots of small issues:  
Hardware tweaks/headnode transferred to more powerful machine.  
Suspended files not automatically retried.
- dCache has been the main cause of our downtimes (hotly followed by College turning off the power on Saturdays: no weekend coverage of Tier 2s)

- new CE (SL5) - the old CE seems to hold up fairly well\*
- new mon box
- WNs upgraded
- Security fixes deployed on all machines

**IC-LESC decommissioned (no manpower)**

\*once we gave it some memory and force-cleaned the home dirs on a regular basis

## The Imperial WMS: [wms00.hep.ph.ic.ac.uk](http://wms00.hep.ph.ic.ac.uk)

- This is a PPS install - i.e you get the latest and greatest software.
- But: It's not the latest and greatest hardware and WMS still need relatively (too?) high spec hardware
- If you have a VO we will support it, just ask. (But please fill out that VO card.)

Currently we support the following VOs:

alice, atlas, babar, biomed, calice, camont, cedar, cms, dteam, dzero, geant, gin, H1, ilc, lhcb, londongrid, minos, na48, ngs, northgrid, ops, pheno, sedi, sixt, southgrid, supernemo, t2k, zeus

Questions ? [lcg-site-admin@imperial.ac.uk](mailto:lcg-site-admin@imperial.ac.uk)

## UCL-HEP:

new 3.1. MON box

new DPM node (1.7) + new pools

Progress in establishing an Atlas production site:

ATLASPRODDISK/ATLASDATADISK token deployed

- site needs to be added to Tiers of Atlas Cache

**In Progress:**

Buying 50 TB of storage and a new NFS server for grid use.

**Future:**

64 bit Worker Nodes

General Site Expansion planned

## UCL-CENTRAL

New CE commissioned at the end of July

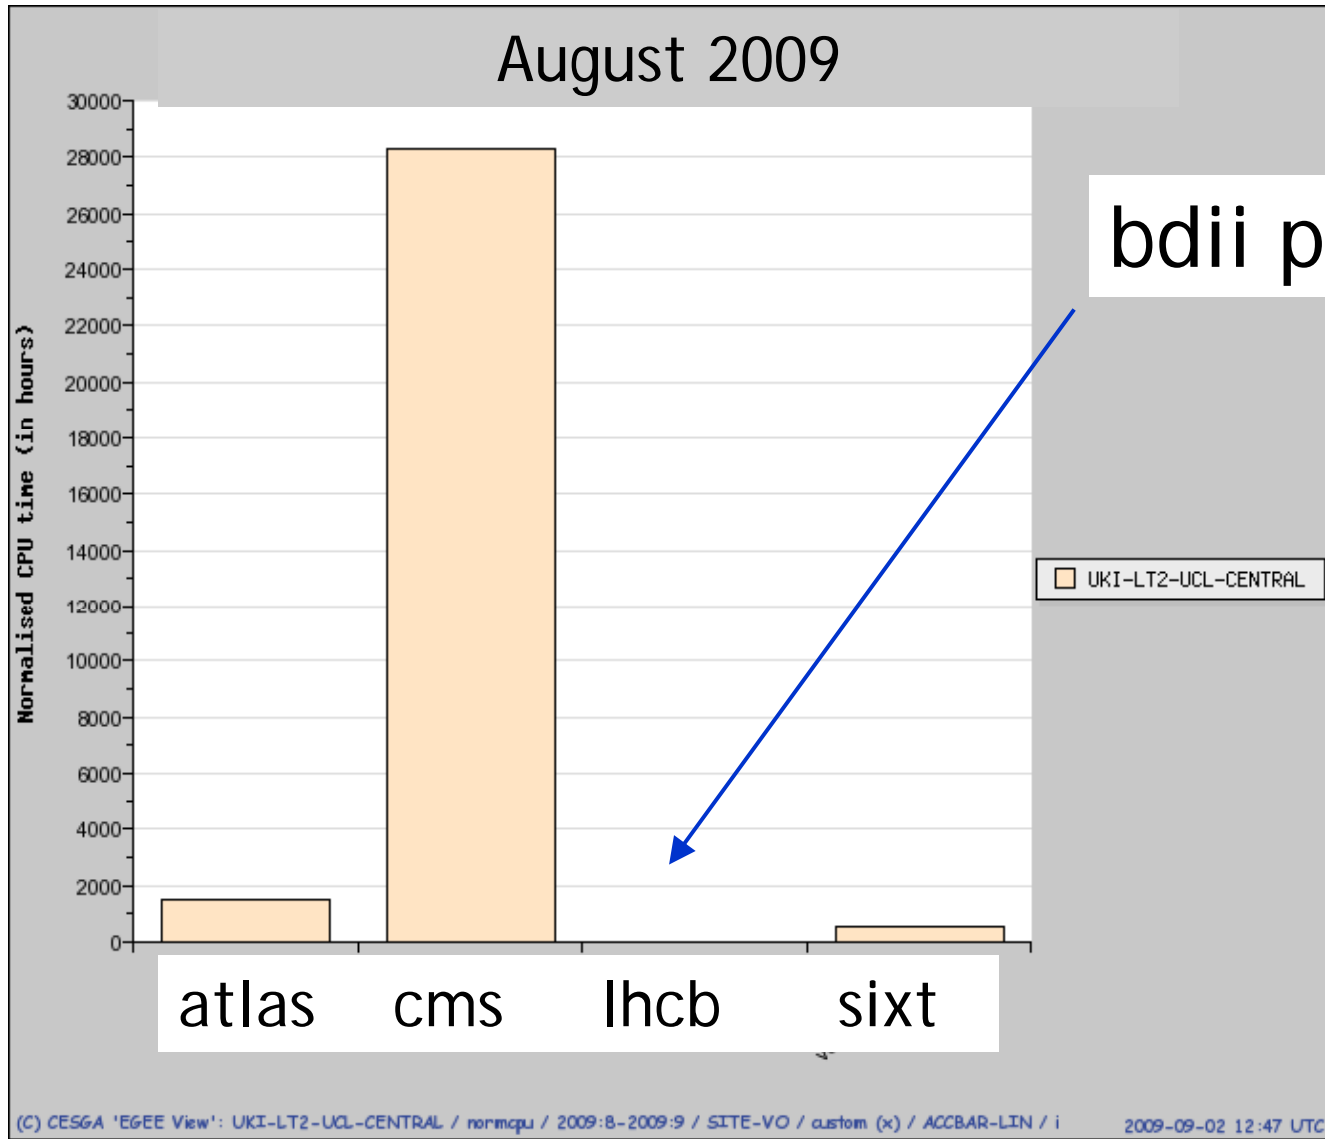
➡ Much improved Availability/Reliability

128 job slots

Some teething troubles with Lustre, now fixed.



# Some small kinks to iron out.....



## Brunel:

92 TB of storage online

600 TB of new storage has been delivered and is being installed

Move to new datacenter starting after GridPP23:  
new SE, new CE with 248 new (+ 152 old, i.e. already running) cores

## Final setup:

Production: 2 CEs, 2SEs

Separate CE (+ WN) and SE for testing

## Royal Holloway

**Main cluster** (at Imperial - next attempt to move in 6 months\*):

Hardware issues:

\* provided network issues have been solved

Brief downtime due to power cut

discovered wrong controller cabling in some nodes, in discussion with vendor - major intervention

Extended downtime for glite updates.

Good ClusterVision support for security updates.

**Old cluster at RHUL:**

SE (se1.pp.rhul.ac.uk) and CE (ce1.pp.rhul.ac.uk)

decommissioned, keep a few WNs for training/testing

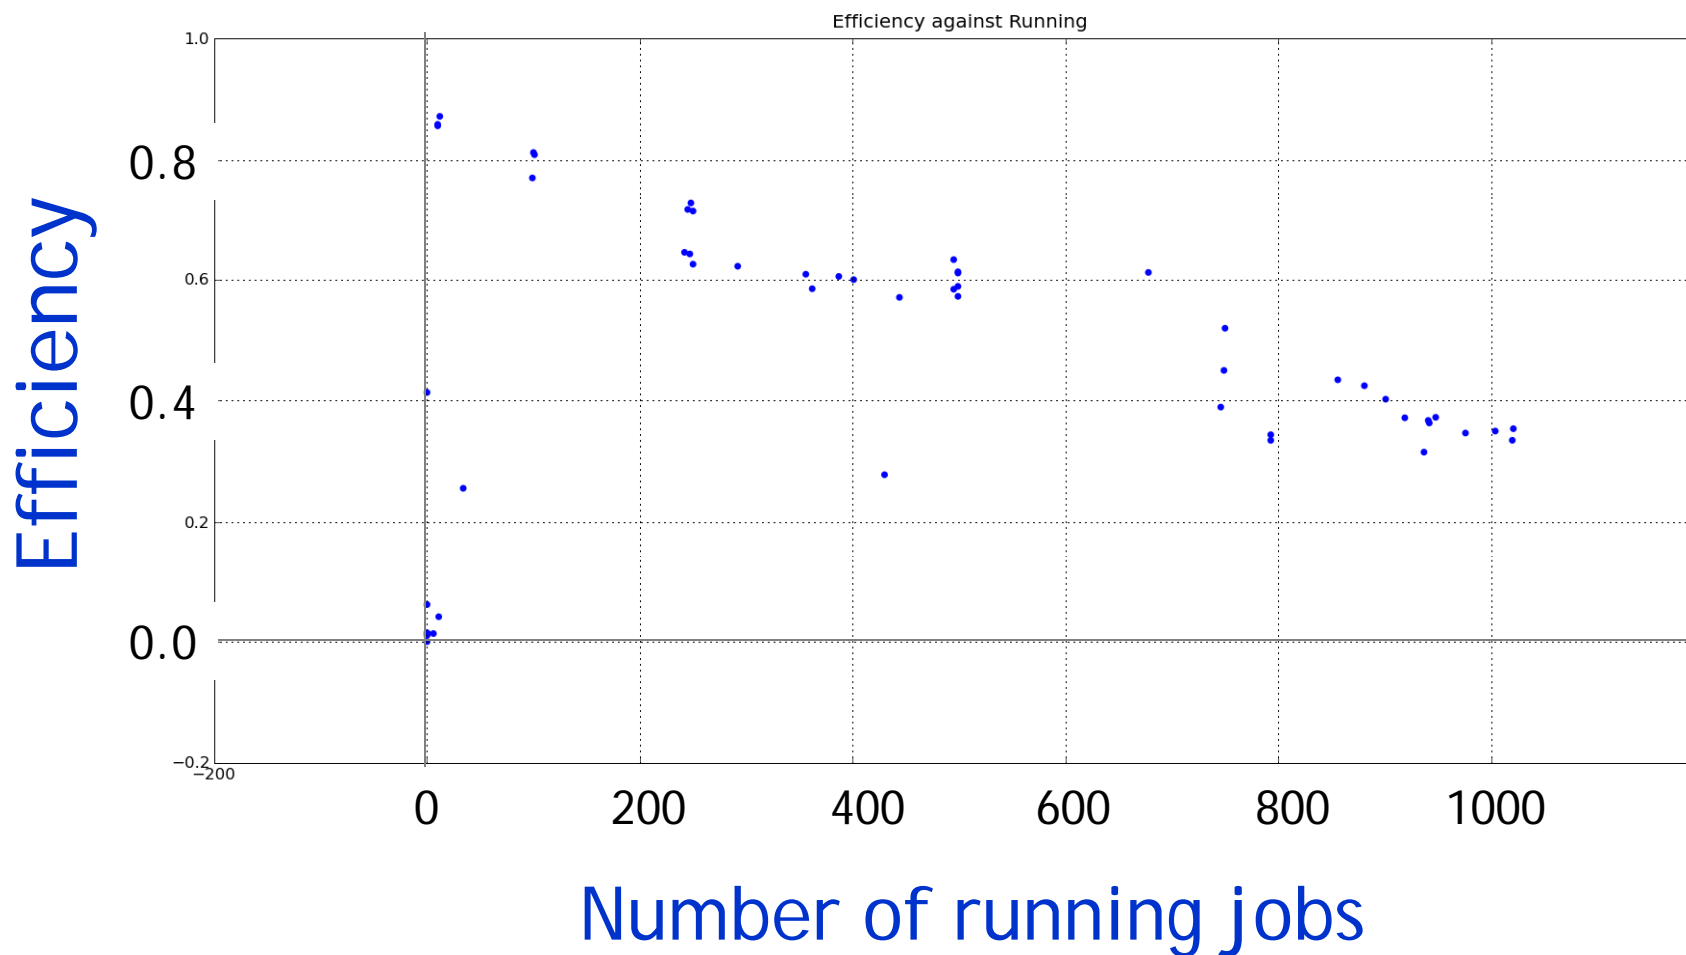
cluster to remain for Tier3/testing after hardware upgrades, not Tier2

Future: SIF funding for site expansion.

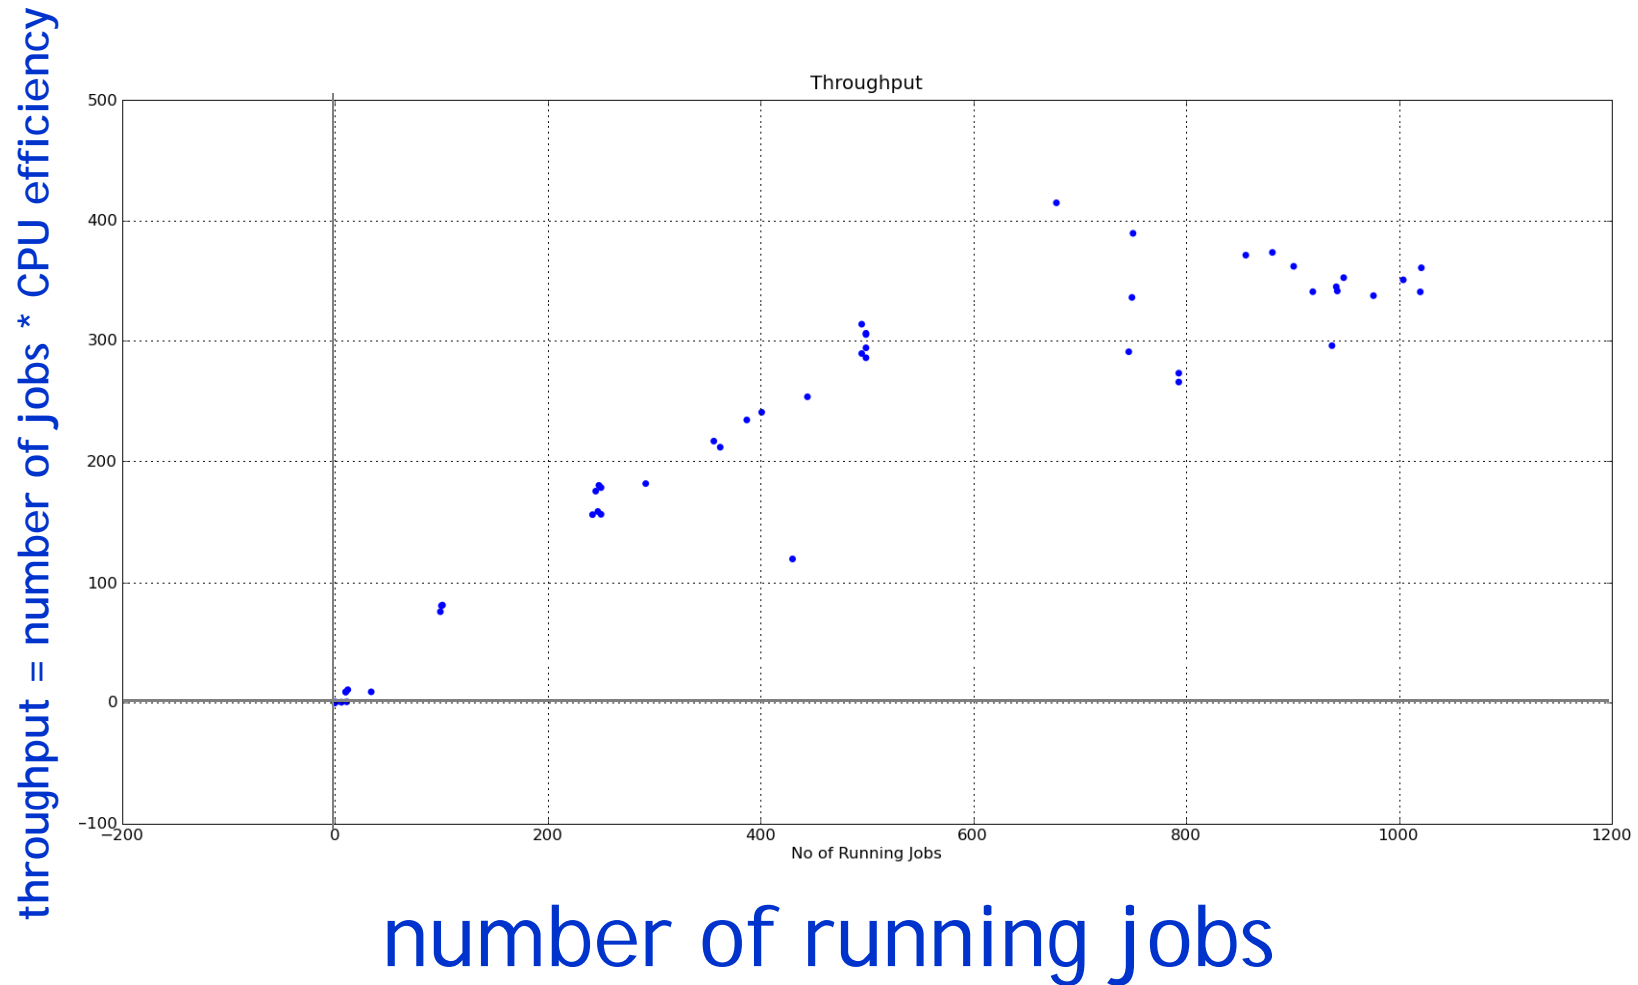
Major news is the deployment of the new Storm SE which has performed very well, provided the file protocol is used.

Lustre is also holding up well.

# Lustre performance in Atlas hammercloud tests



Plot by Chris Walker



Plot by Chris Walker

## Brunel:

One node in each cluster/further upgrades after feedback from experiments.

## QMUL:

Two SL5 test nodes, needs new CE to make use of them.

## RHUL:

long term upgrade in discussion with vendor

short term - ClusterVision has prepared SL5 WN image  
for testing

**UCL - HEP:** Test system with 16+ cores on a separate CE  
for mid-October. Working on deployment system.

**UCL- CENTRAL:** ---

## Imperial College: SL5 all the way.....

9 WNs with 36 cores running CentOS 5.3

32 bit WN tarball (3.1.33) - same as on SL4 nodes

behind a new CE (3.1.34) also on CentOS 5.3

Shares software area with old CE.

(Almost) ready: Accounting and bdii need tweaks.

**Buyer beware:**

CMS showed so much enthusiasm at discovering an SL5 cluster that they send 4000 jobs to 36 slots.

We are happy to upgrade the rest of the cluster, provided we get the OK from the experiments. (How ?)<sub>16</sub>

## Are we ready for data ?

Yes.