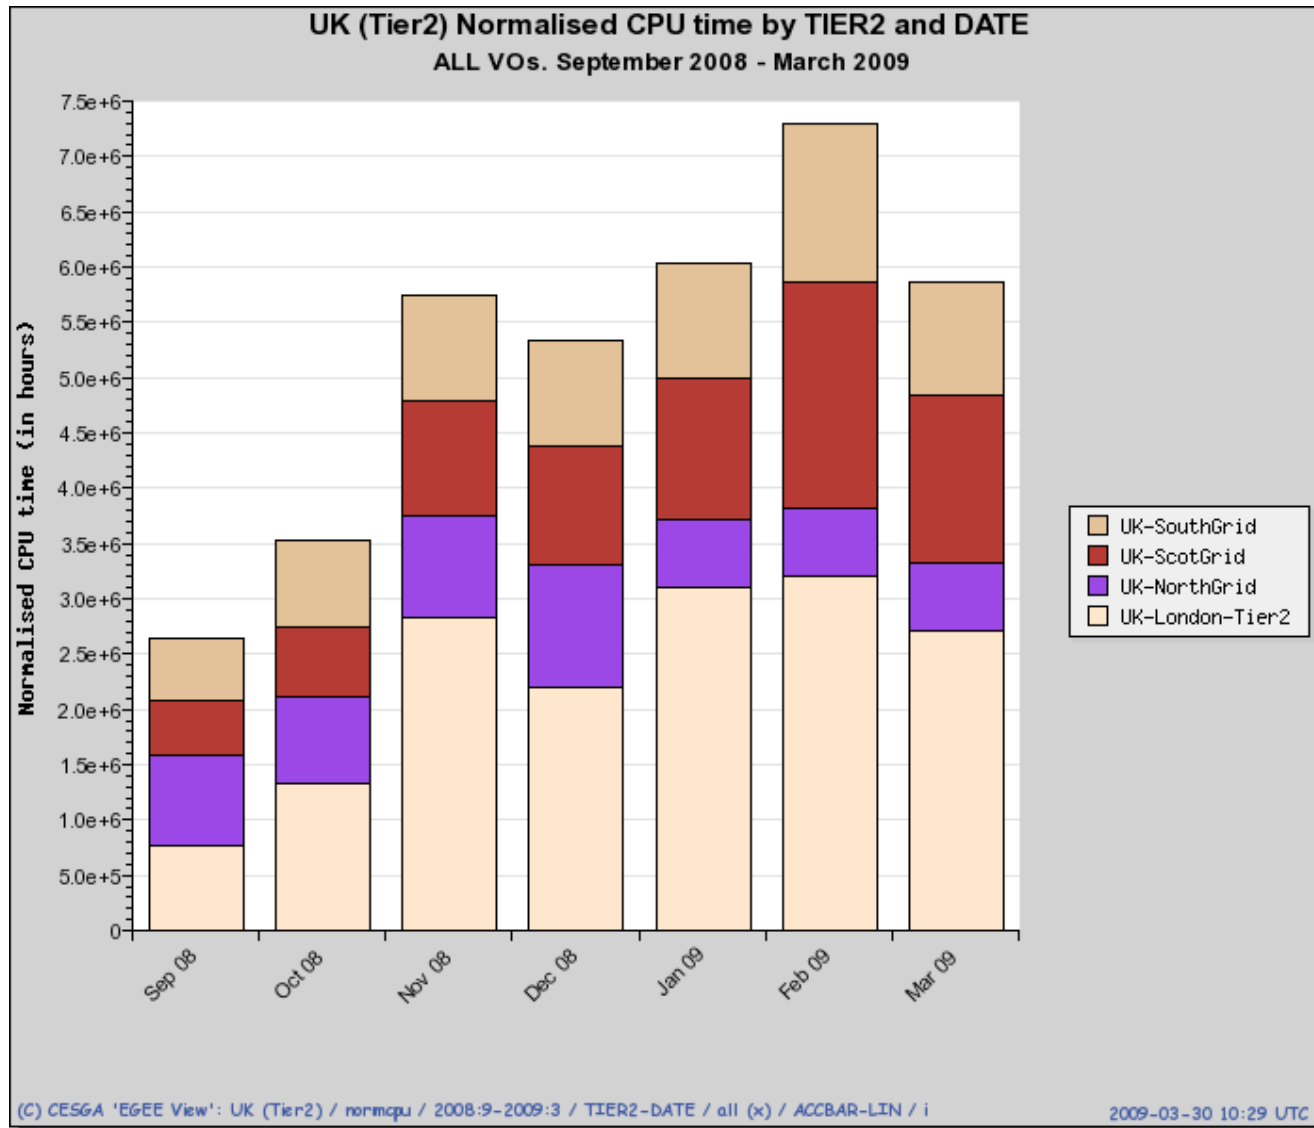


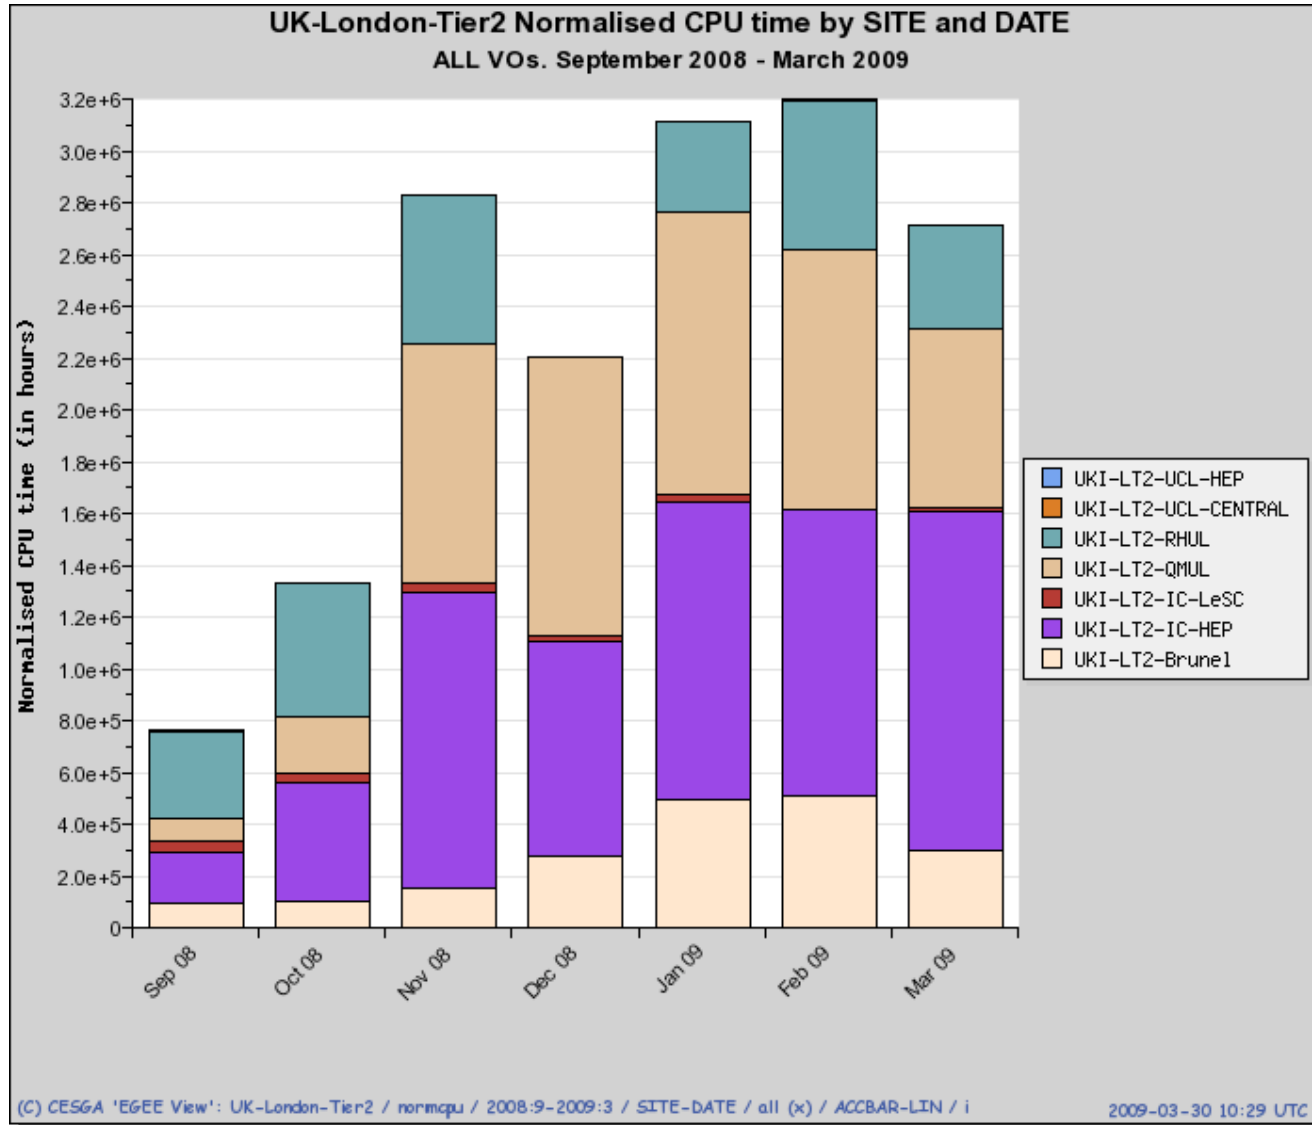
LondonGrid Site Status and Resilience Issues

Duncan Rand
GridPP22

Tier-2 contributions



Site usage



UKI-LT2-Brunel

- Raul Lopes started at beginning of Nov 08
- Move to new SE has improved reliability
- New disk servers have been brought on-line
 - now 75TB total
 - preparing to purchase 300-400TB more disk
- Some of the new Dell worker nodes on-line and working well (air conditioning)
- Moving to new data-centre 2Q09
 - Improved infrastructure and hopefully resilience
 - WAN bandwidth should increase to 1 Gbit/s
- Having two clusters can help resilience

UKI-LT2-IC-HEP

- Some delays bringing new disk on-line
 - now at 150 TB however
- Problems of reliability
 - CE gradually worsened - traced to excess of directories in LHC VO production accounts
 - CE/SE both ran out of memory - both have had RAM increased
- Large numbers of CMS analysis users are stressing dCache SE
 - data concentrated on few disk pool nodes
 - overweight CPU - new disk delivery (300 TB) expected next week should ameliorate this
- Starting to duplicate important services such as site BDII for resilience purposes
- New services now on virtual machines

UKI-LT2-LeSC

- Keith Sephton recently moved on
- Matei Vasile (PhD student) taking over part-time
- SAM availability also suffered from Imperial-HEP SE problems
- Future of LeSC as a grid site is unclear

UKI-LT2-QMUL

- Chris Walker started mid-February
- Site has run CMS, ATLAS and LHCb production in roughly equal amounts
 - recently voted "Most Improved ATLAS Site" by Graeme Stewart
- DPM/Lustre SE
 - providing adequate service for production
 - expectedly poor performance in Hammer tests
- Storm v14 just about released
 - need to install and get Storm running as main SE
- Resilience: dedicated 1 Gbit/s WAN link

UKI-LT2-RHUL

- Govind Songara started a fortnight ago
- Site generally running well
 - Imperial machine room has proved very reliable
 - often full of jobs and regularly gets 100 MB/s WAN rates (current with dedicated 1 Gbit/s WAN link)
 - blotted copybook with a disastrous period of downtime over December (caused by problematic local area network switch)
- Continuing to taking part in ATLAS production and analysis tests
 - showed good but not stunning analysis performance, tuning?
- Continuing to acting as CMS Tier-3 (56TB of CMS data on-site)
- Plans to
 - upgrade old cluster
 - increase existing participation within ATLAS UK
 - move newer cluster to Egham

UKI-LT2-UCL-CENTRAL

- Andreas Roussos and Thomas Jones have joined systems team
 - with four sysadmins should be well placed to provide resilient cover
- Legion cluster now accepted from supplier
- Rolling out service
 - SAM jobs not running in time - fairshare now fixed
- SAM availability improving (97% last week)
- Difficulties getting LHC VO software installed on Lustre so have moved to using NFS space
 - LHCb software now installed and site ready for production
 - CMS some versions of CMSSW installed should start production soon
 - ATLAS almost there - a few RPMs missing, space tokens in place
- Cluster does have admin node resilient pair but not in use yet

UKI-LT2-UCL-HEP

- Concentrating on installing new hardware
- New CE online and stable
- Next immediate steps:
 - Install new DPM head node
 - Upgrade MON to glite-3.1
 - Add new 64-bit worker nodes to cluster

CMS Tier-2 site readiness

<http://lhweb.pic.es/cms/SiteReadinessReports/SiteReadinessReport.html>


T2_UK_London_IC																						
Site Readiness Status: W R R R R R R R R R R R R R R R R R R SD																						
Daily Metric: O E O O O O O O E O O O O O O O O O O O O O O SD																						
Maintenance:	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	CE-SD	
Job Robot:	99%	98%	99%	100%	98%	90%	96%	73%	99%	97%	98%	99%	100%	99%	100%	100%	99%	100%	100%	100%	96%	
SAM Availability:	84%	52%	88%	96%	88%	92%	84%	80%	100%	96%	100%	96%	100%	100%	96%	92%	80%	84%	100%	100%	100%	88%
T2::uplinkT1s:	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	
T2::downlinkT1s:	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	
	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26

Mar

T2_UK_London_Brunel																						
Site Readiness Status: R																						
Daily Metric: O																						
Maintenance:	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up	
Job Robot:	99%	100%	99%	n/a	n/a	99%	100%	97%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	94%	99%	100%	99%
SAM Availability:	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	96%	100%	100%
T2::uplinkT1s:	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	
T2::downlinkT1s:	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	
	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26

Mar

Quality Ranking for T2 Site Readiness last 15 days



UK ATLAS production (last 6 months)

<i>site</i>	<i>success</i>	<i>failure</i>	<i>efficiency</i>
RAL-LCG2	461973	205388	69.2%
UKI-SCOTGRID-GLASGOW	146471	17746	89.2%
UKI-NORTHGRID-MAN-HEP	47421	66097	41.8%
UKI-LT2-QMUL	68093	9047	88.3%
UKI-NORTHGRID-LANCS-HEP	37956	14201	72.8%
UKI-LT2-RHUL	44674	3066	93.6%
UKI-SCOTGRID-ECDF	36183	3604	90.9%
UKI-NORTHGRID-LIV-HEP	33028	4918	87%
UKI-NORTHGRID-SHEF-HEP	30665	3162	90.7%
UKI-LT2-IC-HEP	22937	3812	85.7%
UKI-SOUTHGRID-CAM-HEP	18676	1969	90.5%
UKI-SOUTHGRID-RALPP	17156	1551	91.7%
UKI-SOUTHGRID-OX-HEP	10707	3890	73.4%
UKI-SCOTGRID-DURHAM	3909	274	93.4%
UKI-LT2-Brunel	3093	383	89%
UKI-SOUTHGRID-BHAM-HEP	940	152	86.1%
UKI-LT2-UCL-HEP	0	205	0%
<i>total</i>	<i>983882</i>	<i>339465</i>	<i>74.3%</i>

Some comments on Resilience

- RHUL experience demonstrates importance of good machine room infrastructure and good cluster build
- Build resilience into operation e.g. ATLAS pilot jobs
- Chaotic analysis will really test SE quality - users (not unreasonably) expect 100% success rates
- Automate checking of the consistency of files at Tier-2's and automate replacement of missing/damaged files
- At Imperial we're considering Hadoop distributed file system (in use by CMS at Nebraska Tier-2) which has built in resilience through multiple copies of files - at expense of losing disk volume. How much disk space loss is worth improved reliability?

Risks associated with staff recruitment

- We've had our first crisis and it lasted a year!
- London experience highlights risks associated with shortage of man-power
- Are we too distributed?
 - If only one person at a site then if the post is vacant then have no-one at the site
- How can we speed up recruitment?
- The default policy of relying on Tier-2 coordinators to cover for absent staff has allowed us to muddle along last year but I doubt it is sustainable during data taking

Summary

- Manpower situation improving
 - Brunel, QMUL and RHUL recently appointed new system administrators
 - but LeSC is losing an experienced admin
 - there is a steep learning curve for new staff
- We have large sites with new kit and so we are doing well with the numbers
- Improved balance across sites
 - CMS: Imperial and Brunel
 - ATLAS: QMUL and RHUL

Outlook

- Get new staff up to speed and well-integrated with ATLAS and CMS operations
- Issues:
 - Improve Imperial SE performance
 - Successful transition to new data-centres at Brunel and RHUL
 - Get Storm working in production at QMUL
 - Bring UCL Central fully into production and analysis for ATLAS
- Looking to demonstrate sites' ability to download and analyse ATLAS and CMS data reliably and rapidly