

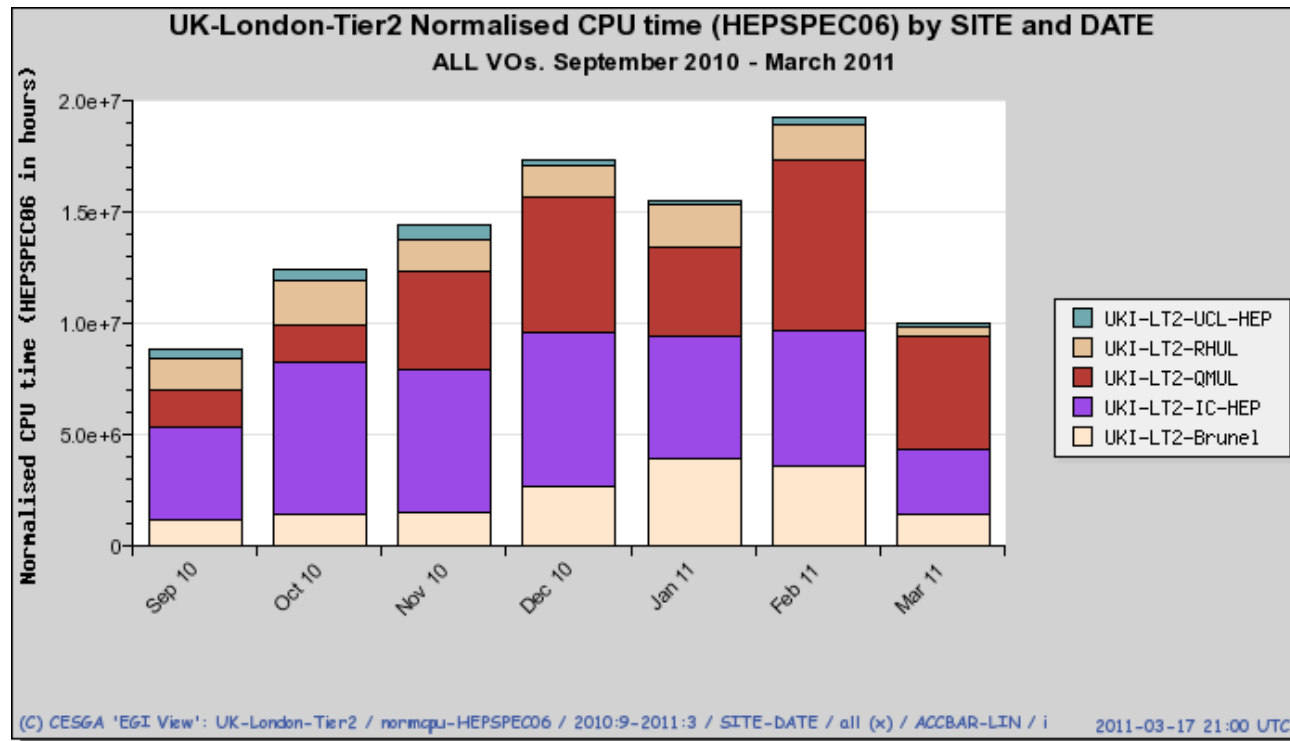
LondonGrid Status

Duncan Rand

on behalf of the London site admins

GridPP26

LondonGrid usage by site



Tier-2 CPU efficiencies

CPU Efficiency (%) by SITE and VO				
SITE	atlas	cms	lhcb	Total
UKI-LT2-Brunel	95.0	74.0	95.5	81.6
UKI-LT2-IC-HEP	96.8	71.0	66.7	76.0
UKI-LT2-QMUL	89.0	89.0	71.1	86.4
UKI-LT2-RHUL	83.6	92.1	96.1	85.9
UKI-LT2-UCL-HEP	84.9	4.4	92.3	86.1
Total	89.6	75.8	73.6	82.6


[Click here for a csv dump of this table](#)
[Click here for a EXTENDED csv dump](#)

Key: 0% <= eff < 50%; 50% <= eff < 60%; 60% <= eff < 75%; 75% <= eff < 90%; 90% <= eff < 100%; eff >= 100% (parallel jobs)

CPU Efficiency (%) by SITE and DATE								
SITE	Sep 10	Oct 10	Nov 10	Dec 10	Jan 11	Feb 11	Mar 11	Total
UKI-LT2-Brunel	70.6	80.4	73.1	78.8	90.4	89.1	75.0	81.6
UKI-LT2-IC-HEP	67.3	83.6	78.5	84.1	65.6	78.5	68.2	76.0
UKI-LT2-QMUL	72.9	63.7	91.5	91.6	80.3	86.4	94.4	86.4
UKI-LT2-RHUL	91.0	94.9	70.2	80.1	91.9	85.3	93.2	85.9
UKI-LT2-UCL-HEP	81.9	85.0	88.2	79.3	84.6	92.3	94.2	86.1
Total	72.7	79.1	83.4	87.1	78.0	84.9	87.8	82.6

[Click here for a csv dump of this table](#)
[Click here for a EXTENDED csv dump](#)


Key: 0% <= eff < 50%; 50% <= eff < 60%; 60% <= eff < 75%; 75% <= eff < 90%; 90% <= eff < 100%; eff >= 100% (parallel jobs)



Usage
by VO

Lots of production being done in Tier-3 role

Should this be 'rewarded/recognised' in some way?



Site availability

	ops	CMS	ATLAS	LHCb
Brunel	100%	100%	93%	99%
Imperial	97%	95%	97%	97%
QMUL	90%	73%	85%	89%
RHUL	89%	85%	86%	87%
UCL	98%	-	97%	94%

Steve Lloyd tests, last 6 months

ATLAS disk pledges

Federated Tier-2 totals (2011 pledge)

Federation	Pledge (TB)	Actual (TB)
northgrid	1539	1453
scotgrid	1237	1043
londont2	1342	529
southgrid	582	549
Total	4700	3575

Actual > WLCG Pledge Actual < WLCG Pledge
(Excludes LOCALGROUPDISK)

Tier-2 site view (2011 pledge)

Site (click for detail)	WLCG 2011		GridPP (includes LOCALGROUPDISK)				
	Pledge(TB)	Actual (TB)	Pledge (TB)	Actual (TB)	Not-free (TB)	Free (TB)	Usage
UKI-LT2-Brunel	0	4	0	4	2	2	
UKI-LT2-IC-HEP	0	17	0	17	1	16	
UKI-LT2-QMUL	733	263	916	304	199	105	
UKI-LT2-RHUL	155	225	194	266	139	127	
UKI-LT2-UCL-HEP	10	20	12	43	5	38	

WLCG pledges 2011

	CPU (HEPSPEC06)		Disk (TB)	
	Pledge	Expected	Pledge	Expected
ATLAS	7794	27683	1342	529 -> ~2000 ⁺
CMS	18407	17613 -> ~22000 [*]	1097	1397
LHCb	1985	2132	1	34

⁺ including extra disk at QMUL and RHUL

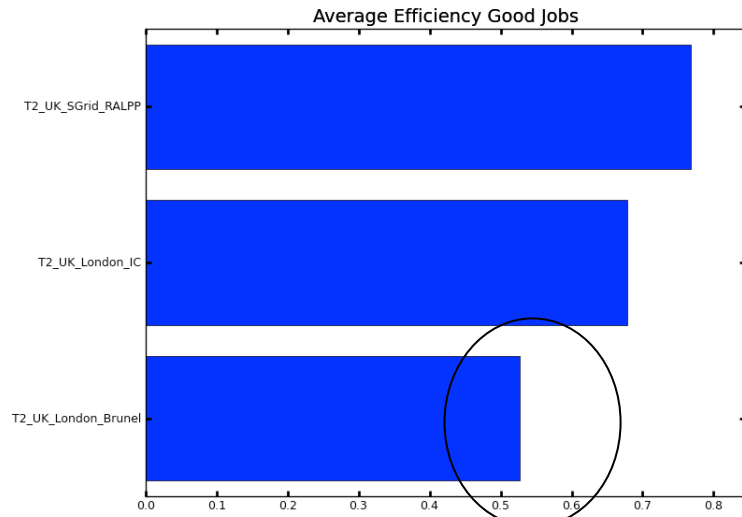
^{*} including extra WN at Imperial and Brunel

http://gstat-wlcg.cern.ch/apps/capacities/pledge_comparison/


Brunel

- Ivan Reid joining in “Grid experiment support” role
- 7201 HEPSPEC06 and 450 TB online
 - 2800 HEPSPEC06 recently added
 - another ~1900 HEPSPEC06 WN ordered
 - 120 TB more disk expected online soon
- JANET agreed to fund 10 Gbps WAN
 - but planned upgrade of local equipment due 2012-2013 so needs to be brought forward
 - Brunel HEP to get 4 Gbps of that
- data centre back-up generator ready to go; so large needed council planning permission - now painted khaki!
- new UPS batteries installed - all servers now have backup

Sept- Dec 2010: eff=0.52



Jan- Apr 2011: eff=0.79



CMS analysis
efficiency
improvement
at Brunel

Imperial


- Tim Watts moved on at the start of March
- 12400 HEPSPEC06 and 1.1 PB now online
 - extra 4660 HEPSPEC06 and 650 TB recently added
 - Another 5290 HEPSPEC06 new WN racked; moving it to another machine room because of floor loading problems
- CMS /store/user permanently replicated
- issues:
 - increase of local batch farm size resulted in dcache head-node running out of memory - now 48GB RAM in machine
 - room thermostat issue

Imperial: fabric management


- load-balancer (Linux virtual server) in-front of site & top bdiis
 - spreads the load and can stop requests going to one if it's broken
- all ldap servers and NFS (netapp) are redundant
- a lot of our service machines (CEs, BDIs) etc... float atop a (minimalist) VM infrastructure so they can be moved round if a VM host fails
- Using the spare space on dCache to cache files increases our CPU efficiency (mentioned before)

Imperial: dynamic capping of CMS analysis jobs

- Analysis essentially limited by storage IOPS
- Implemented a static cap on number of running CMS analysis jobs
- However, some jobs less I/O bound than others – can increase the cap
- Also load on dcache increases when data being copied in over WAN – then need to reduce cap
- So we implemented a dynamic cap
$$\text{cap} = 1100 - \text{sum}(\text{load}_{15} \text{ of dcache servers})$$




1. queue of CMS analysis jobs empties




2. number of running jobs falls

4. cap (shaded area) on number of running analysis jobs rises




3. dcache load and outgoing network traffic fall



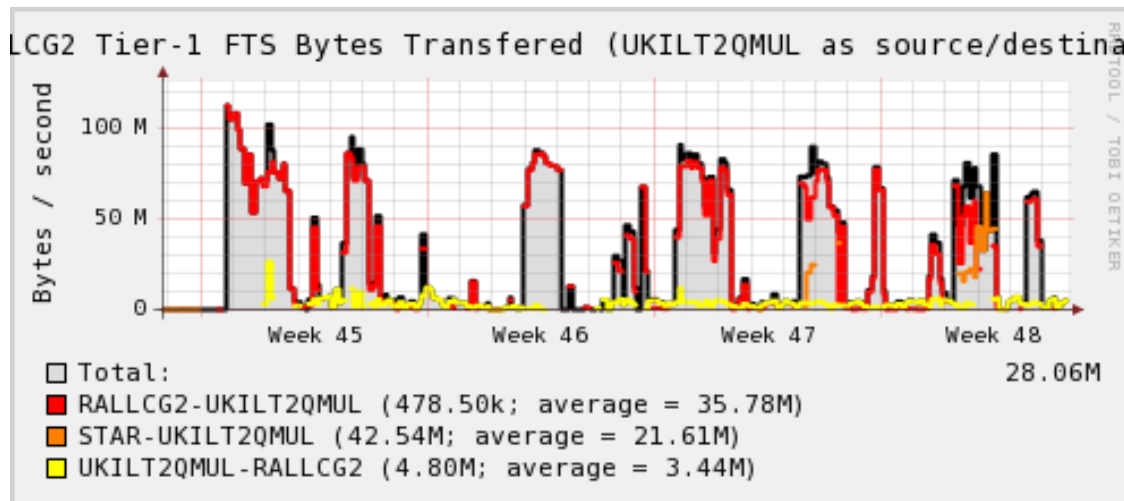
Queen Mary

- Sergey Dolgobrodov left in February
- 25343 HEPSPEC06 and 320 TB currently online
 - 100 new Intel X5650 dual CPU machines with 24 GB RAM
 - Received 1PB of storage (Dell R510) - to be deployed shortly
- Deployed CVMFS for ATLAS software area
 - significantly reduced load on Lustre metadata server
 - dramatically increased job throughput
 - QMUL was the first ATLAS Tier-2 site to switch to CVMFS
- Site did well in ATLAS sonar tests – Tier-2d?
- Some machine room power supply/ air conditioning issues

ATLAS production jobs at QMUL



2900 jobs



Dramatic increase in WAN traffic

- now regularly saturates external 1Gig link
- implemented traffic shaping (ingress policing) on incoming traffic in order to give some room for job submission, catalogue lookups etc.

Royal Holloway

- New data centre working well
- Just taken delivery of new hardware
 - Extra 420 TB (Dell R510) and 6400 HEPSPEC06 (DELL C6100)
 - Alcys system integration
- Have been limited to 300-450 Mbps WAN cap
 - not ideal but hasn't demonstrably affected capacity to do analysis
- Issues
 - failure of SE head node resulted in corrupted database and took a few days to restore
- Plans:
 - shortly to upgrade WAN to 1 Gbps (dedicated)
 - install Argus/glexec



- Two of the racks of new Dell hardware in RHUL data centre

- Hot aisle at back, cooled air comes out of 'loud-speakers'

UCL

- Adam Davison joined in October
 - helping with integration of Legion cluster with UCL-HEP site
- ~1400 HEPSPEC06 new CPU ordered
- Disk still <50 TB disk so unable to receive ATLAS data - however 120TB almost ready to go online
- CREAM CE working on HEP cluster
- Legion cluster ready for use next step is to install a CREAM CE in front of it

End