

GridPP

UK Computing for Particle Physics

The WLCG Service from a Tier-2 Viewpoint

Duncan Rand
Imperial College London

- Main goal to run LHC VO analysis and production jobs successfully
- Get the data on site as fast as possible
 - CMS requests are approved - UK sites very rapid response time (<1 hour)
 - bandwidth - 'can never have too much bandwidth'
 - quality - need to monitor site transfer quality and flag alarm if it drops
- Run jobs reliably
 - eliminate as many of the single points of failure as possible
 - dual site BDII's, local top-level BDII
 - NetApp reliable storage server for experimental software area, VM images etc
- Understand as well as possible what is going on at the site
 - a good start is to monitor everything pertinent
 - ganglia, Nagios, Cacti (network switches)
 - site and VO status boards

- Different to Tier-0/1
 - OK, cache-like, no custodial responsibilities
 - however, sizes and service levels approaching Tier-1's (some soon to have a petabyte of disk)
- But may have local infrastructure constraints
 - likely to have to negotiate with local staff who are often unfamiliar with demands of HEP computing
 - less dedicated infrastructure: e.g. ad-hoc machine rooms often scattered across campus with poor connectivity; some sites have grown rapidly
 - WAN: big disparity in networking demands compared to a typical university's
- Staff
 - generally low staff numbers - vulnerable to absence, turnover, illness, vacation etc
 - have had sites without staff for extended periods



CE (ceprod00)
 JobSubmit
CAVer
 BrokerInfo
 Csh
 SoftVer
 Rep
 RepCr
 RepGet
 RepRep
 RepDel
 RepFree
 RepISenv

Site Status

Upcoming Downtimes
 No currently scheduled downtimes. :)

dCache Space Usage

| Pool Group | Used Space (GB) | Free Space (GB) | Graph |
|------------|-----------------|-----------------|-------|
| CMS | 292458 (55%) | 240581 | |
| Atlas | 4315 (25%) | 12478 | |
| LHCb | 4540 (15%) | 26654 | |
| Other | 1686 (5%) | 30036 | |

CE (ceprod02)
 JobSubmit
CAVer
 BrokerInfo
 Csh
 SoftVer
 Rep
 RepCr
 RepGet
 RepRep
 RepDel
 RepFree
 RepISenv

CMS Site Status

| Visible Job | Robot | CE | SRM | Prod | Analysis | Running | Pending | Phedex In | Phedex Out | | | |
|-------------|-------|-----|-----|------|----------|---------|---------|-----------|------------|------|-----|-----|
| OK | 100% | 200 | OK | OK | 100% | 941 | 100% | 1922 | 740 | 1228 | n/a | n/a |

CMS Site Readiness

| 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 |
|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|
| W | R | R | R | R | R | R | R | R | R | W | W | R | R | R |

LHCb Site Mask


LCG.UK1-LT2-IC-HEP.uk | cepro02.hep.ph.ic.ac.uk

SRM (gf02)
 GetURLs
 LsDir
 Put
 Ls
 GetURLs
 Get
 Del
 All






CE Estimated Reponse Time

| Experiment | ceprod00 | ceprod02 |
|------------|------------------|------------------|
| CMS | 191069 (853/949) | 191013 (856/934) |
| Atlas | 33212 (15/177) | 33156 (15/177) |
| LHCb | 0 (145/0) | 3 (145/0) |


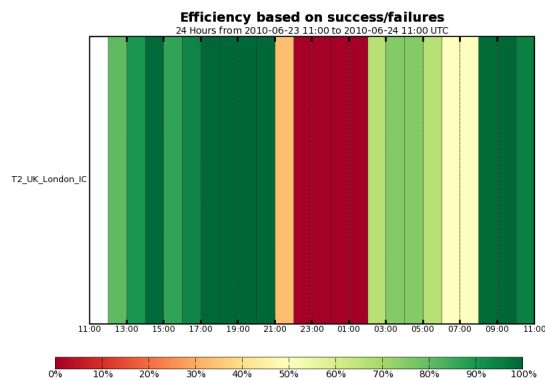


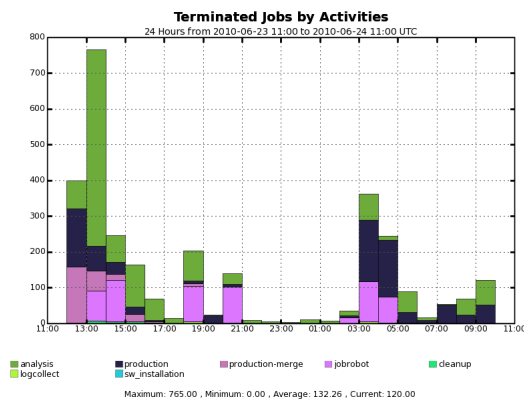

Key: Estimated Response Time Seconds (Running Jobs / Queued Jobs)



Last Updated: Thu Jun 24 12:00:05 BST 2010




- A view of the site from CMS point of view






- UK sites comment starting to find LAN usage growing
- LAN likely to become more of an issue as sites get even more cores per worker node and more storage per disk server
- UK Tier-2 sites generally share a university 1 Gbps WAN connection, some are capped, some are not; several have a dedicated 1 Gbps link
- Some larger universities have 10 Gbps connection but only 1 or 2 Gbps to Tier-2
- With large increase in disk servers RAL Tier-1 now able to push data out to sites at good rates
- Sites are starting to have issues with contention, failing SAM tests, backlogs of transfers etc - many are looking to upgrade WAN
- At Imperial the machine room link has been upgraded from 1 to 10 Gbps
 - got adequate rates (2.5 Gbps) but still less than expected

- Ran a few tests
- Much better rates outside firewall
- Encouraged by other site's positive experiences, we decided to move the grid site outside the firewall









■ T2_UK_London_IC

Maximum: 494.92 MB/s, Minimum: 0.00 MB/s, Average: 108.27 MB/s, Current: 20.10 MB/s

- GridPP has a dedicated storage support team
 - reflects importance and difficulty of running storage successfully
 - lively storage group discusses all aspects of Tier-2 storage: weekly meetings, email list and a recent two day workshop
 - general ongoing request to make the storage software as good as possible
- DPM
 - most common SE in UK Tier-2's
 - simple and easy to use but lacking some desirable features
 - e.g. automatic hot file replication, intelligent pool selection
 - pleased to see DPM support staff numbers increased
- Storm + Lustre/GPFS
 - used at three UK sites (Bristol, QMUL, Edinburgh)
 - QMUL only one using it in production
 - we need to build more UK experience

- dCache

- liked by the two UK sites that run it (Imperial and RALPPD both CMS)
- documentation could be improved however - more help on day-to-day tasks
- dCache workshop a very good idea

| Pool Group | Used Space (GB) | Free Space (GB) | Graph |
|------------|-----------------|-----------------|---|
| CMS | 310033 (58%) | 223155 |  |
| Atlas | 4426 (27%) | 12368 |  |
| LHCb | 4540 (15%) | 26654 |  |
| Other | 2077 (7%) | 29645 |  |

- dCache hot file replication at Imperial

- switched on hot file replication
- files are replicated from hottest fraction of pools
- fill disk up with cached copies of hottest files (green)
- works well - we now see a more even balance of load and network bandwidth access over all 35 disk servers

- Hadoop HDFS model

- some interest in the UK
- nice to be able just to switch off disk servers and maintain a service, but is it worth the 'loss' of disk?

- **Communication with WLCG operations**
 - generally successful in UK
 - UK chair of GDB
 - GridPP production manager regularly attends WLCG Ops meeting
 - Tier-2 Coordinators also attend GDB meetings in rotation
- **GridPP internal communication**
 - weekly deployment team meetings
 - recently broadened out to include sys-admins every other week - more of an operations flavour
- **Direct VO communication with Tier-2 sys-admins is also important**
 - in past Tier-2's have often provided 'a Tier-2' with LHC VO's one of many
 - now need to be a part of the VO
 - attendance at relevant VO computing weeks, Jamborees etc

- Imperial, Brunel, RALPPD and Bristol Tier-2's support CMS
- Excellent communications with CMS
 - CMS direct communication with Tier-2s via hyper-news
 - backed up by UK CMS support staff and fortnightly UK CMS meetings
 - most UK Tier-2's have a member of CMS computing at the site
 - Savannah tickets used to communicate with CMS or other sites (e.g. during data transfer problems)
- Issues
 - Remote stage-out is difficult
 - remote lcg-cp often unreliable
 - blocks job slot copying out
 - can be very disruptive to remote site SE
 - recent bug in CRAB made it worse

- CMS site status board and site readiness status pages are excellent - easy for a site to know if it is 'OK' and to drill down to the problem if not

| T2_UK_London_IC | | | | | | | | | | | | | | | | | | | | | | | |
|--|------|------|------|------|------|------|-----|------|------|------|-----|------|------|------|------|-----|-----|------|------|-----|----|----|--|
| Site Readiness Status: W W R R R R R R R R R R W W R R R | | | | | | | | | | | | | | | | | | | | | | | |
| Daily Metric: O O O O O O E E O O O O O O O O E E O O O | | | | | | | | | | | | | | | | | | | | | | | |
| Maintenance: | - | Up | Up | Up | Up | Up | Up | Up | Up | Up | Up | Up | Up | Up | Up | Up | Up | Up | Up | Up | Up | | |
| Job Robot: | 100% | 100% | 100% | 100% | 100% | 100% | 84% | 100% | 100% | 100% | 99% | 100% | 100% | 100% | 100% | 50% | 43% | 100% | 100% | n/a | | | |
| SAM Availability: | 80% | 80% | 80% | 100% | 84% | 80% | 72% | 76% | 80% | 100% | 88% | 100% | 100% | 88% | 100% | 78% | 44% | 100% | 100% | n/a | | | |
| Good T2 links from T1s: | 8/8 | 7/8 | 7/8 | 8/8 | 8/8 | 8/8 | 8/8 | 8/8 | 7/8 | 8/8 | 8/8 | 8/8 | 8/8 | 8/8 | 8/8 | 7/8 | 8/8 | 8/8 | 8/8 | n/a | | | |
| Good T2 links to T1s: | 4/4 | 4/4 | 4/4 | 4/4 | 4/4 | 4/4 | 4/4 | 4/4 | 4/4 | 4/4 | 4/4 | 4/4 | 4/4 | 4/4 | 4/4 | 4/4 | 4/4 | 4/4 | 4/4 | n/a | | | |
| Active T2 links from T1s: | 8 | 8 | 8 | 8 | 8 | 8 | 8 | 8 | 8 | 8 | 8 | 8 | 8 | 8 | 8 | 8 | 8 | 8 | 8 | 8 | 8 | | |
| Active T2 links to T1s: | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | | |
| | 02 | 03 | 04 | 05 | 06 | 07 | 08 | 09 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | |
| | Jun | | | | | | | | | | | | | | | | | | | | | | |

Report made on 2010-06-23 02:30:01 (UTC)

* = Due to operational errors, the metric has been corrected manually (!=SSB).
 [X] = Errors on weekends are ignored on Site Readiness computation for T2s [info]
 "Site Readiness Status" as defined in Site Commissioning Twiki:
 [G] = READY
 [W] = WARNING
 [NR] = NOT-READY
 [SD] = SCHEDULED-DOWNTIME

"Daily Metric" as boolean AND of all individual metrics:
 [G] = OK (All individual metrics above Site Commissioning Thresholds; "n/a" ignored)
 [E] = ERROR (Some individual metrics below Site Commissioning Thresholds)
 [SD] = SCHEDULED-DOWNTIME

- INDIVIDUAL METRICS -


"Scheduled Downtimes": site maintenances
 [G] = Site is not declaring Scheduled-downtime
 [SD] = SD=full-site; SE-SD: All CMS SE(s) in SD; CE-SD: All CMS CE(s) in SD
 [~] = Some SE or CE services (not all) Downtime
 "SAM Availability":
 [G] = SAM availability is ≥ 80%
 [E] = SAM availability is < 80%
 "Active T2 links from T1s":
 [G] = Site has ≥ 4 DDT-commissioned links from T1 sites
 [E] = Otherwise

"Job Robot":
 [G] = Job success rate is ≥ 80%
 [E] = Job success rate is < 80%
 [Y] = Jobs submitted but not finished
 [n/a] = Job success rate is n/a
 "Active T2 links to T1s":
 [G] = Site has ≥ 2 DDT-commissioned links to T1 sites
 [E] = Otherwise

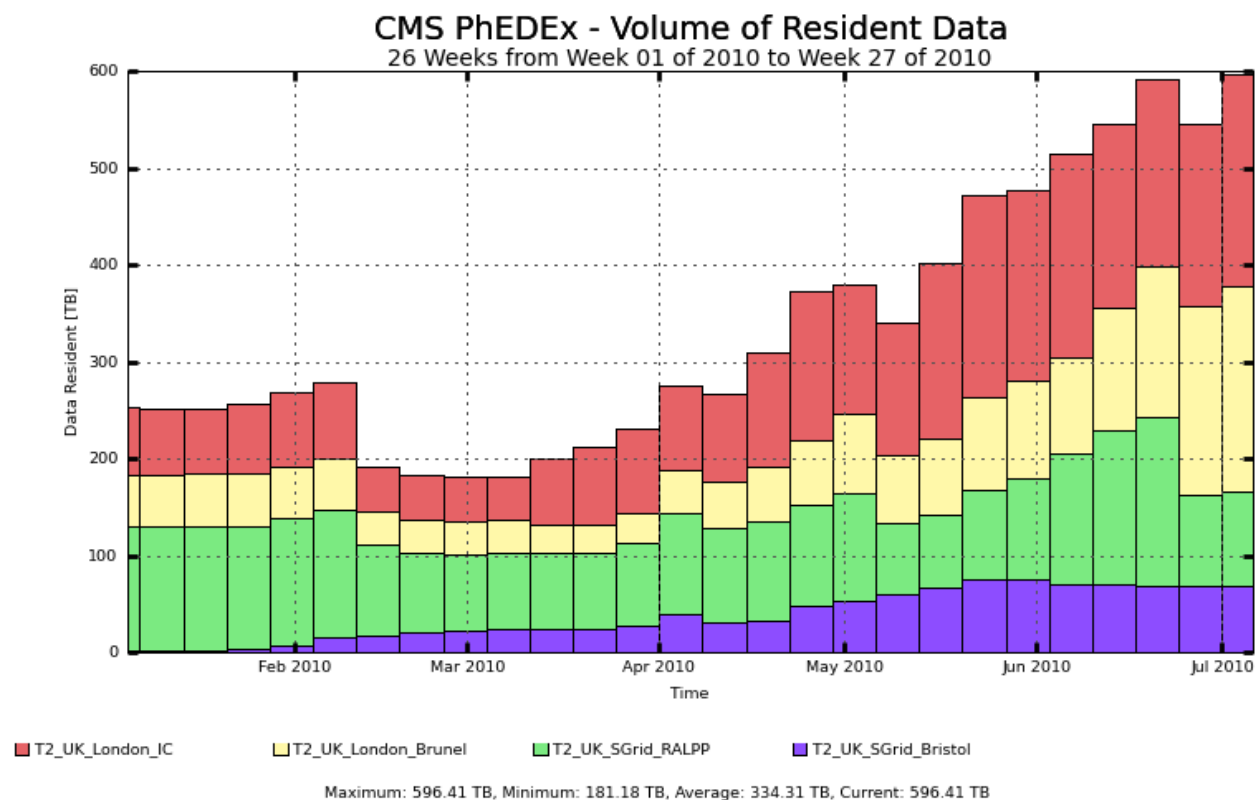
"Good Links":
 [G] = at least half of links have 'good' transfers (i.e. with transfer quality > 50%)
 [E] = Otherwise

| Site Name | Visible | JobRobot | SAM TESTS | | Production | Analysis | Site usage | |
|-----------------|---------|-----------|-----------|-----|------------|------------|------------|---------|
| | | | CE | SRM | | | Running | Pending |
| T2_UK_London_IC | On | 100%(200) | On | On | n/a | 100%(1652) | 5/4 | 1 |


- Data has been copied to UK Tier-2's at good rates



- Space used has increased significantly in last 4 months: 0.6 PB of data is now hosted





- CMS jobs are running...



- Majority of UK Tier-2 sites support ATLAS
 - a petabyte of ATLAS data now hosted at UK Tier-2's
- Tier-2 communications with ATLAS work well
 - support devolved to UK 'cloud'
 - UK weekly operations meeting very useful
 - 'Jamborees' are a good way for sys-admins to understand what ATLAS are trying to achieve
- ATLAS Monte Carlo production works really smoothly
- Issues
 - ATLAS software has sometimes been difficult to get installed in the past
 - a number of UK sites recently ran out of space in their space tokens - probably more of a VO issue but still a problem for the site
 - Tier-2 hardware purchases affected by current method of staging to local disk; should sites buy multiple local disks and should they be higher performance drives?

- Excellent: really easy to see site status, error rate and to drill down to logs. But it is often slow!
- Panglia - easy to monitor what is going on

| UK Sites | Job Nodes | Jobs | Latest | Pilot Nodes | defined | assigned | waiting | activated | sent | running | holding | transferring | finished | failed tot | trf | other |
|---------------------------------|-----------|------|-------------|-------------|---------|----------|---------|-----------|------|---------|---------|--------------|----------|------------|-----|--------|
| Site Name | 995 | 1199 | 06-30 13:31 | 5590 | 1128 | 0 | 0 | 23 | 0 | 310 | 34 | 4796 | 1199 | 20% | 18% | 2% |
| ANALY_BHAM ✓ | 24 | 22 | 06-30 13:30 | 142 | 2 | 0 | 0 | 1 | 0 | 13 | 1 | 0/0 | 15 | 22 | 59% | 3% 57% |
| ANALY_CAM ✓ | 35 | 2 | 06-30 13:25 | 166 | 2 | 0 | 0 | 4 | 0 | 16 | 0 | 0/0 | 243 | 2 | 1% | 1% 0% |
| ANALY_ECDF ✓ | 23 | 0 | 06-30 13:31 | 499 | 1 | 0 | 0 | 0 | 0 | 12 | 1 | 0/0 | 12 | 0 | 0% | 0% 0% |
| ANALY_GLASGOW ✓ | 248 | 28 | 06-30 13:31 | 1018 | 578 | 0 | 0 | 0 | 0 | 19 | 2 | 0/0 | 571 | 28 | 5% | 2% 3% |
| ANALY_LANCS ✓ | 133 | 15 | 06-30 13:29 | 562 | 0 | 0 | 0 | 0 | 0 | 27 | 0 | 0/0 | 862 | 15 | 2% | 2% 0% |
| ANALY_LIV ✓ | 42 | 71 | 06-30 13:31 | 288 | 1 | 0 | 0 | 0 | 0 | 20 | 3 | 0/0 | 268 | 71 | 21% | 5% 16% |
| ANALY_MANC ✓ | 77 | 6 | 06-30 13:31 | 142 | 0 | 0 | 0 | 0 | 0 | 0 | 6 | 0/0 | 111 | 6 | 5% | 5% 0% |
| ANALY_OX ✓ | 50 | 900 | 06-30 13:31 | 286 | 3 | 0 | 0 | 0 | 0 | 18 | 2 | 0/0 | 107 | 900 | 89% | 89% 0% |
| ANALY_QMUL ✓ | 160 | 59 | 06-30 13:30 | 197 | 147 | 0 | 0 | 2 | 0 | 47 | 6 | 0/0 | 1345 | 59 | 4% | 4% 1% |
| ANALY_RAL ✓ | 0 | 0 | | 750 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0/0 | 0 | 0 | | |
| ANALY_RALPP ✓ | 92 | 6 | 06-30 13:31 | 761 | 0 | 0 | 0 | 0 | 0 | 16 | 1 | 0/0 | 145 | 6 | 4% | 3% 1% |
| ANALY_RHUL ✓ | 45 | 72 | 06-30 13:31 | 283 | 10 | 0 | 0 | 16 | 0 | 92 | 10 | 0/0 | 149 | 72 | 33% | 32% 1% |
| ANALY_SHEF ✓ | 52 | 18 | 06-30 13:31 | 398 | 384 | 0 | 0 | 0 | 0 | 17 | 0 | 0/0 | 882 | 18 | 2% | 1% 1% |
| ANALY_UCL ✓ | 14 | 0 | 06-30 13:30 | 98 | 0 | 0 | 0 | 0 | 0 | 13 | 2 | 0/0 | 86 | 0 | 0% | 0% 0% |



- LHCb are pretty unobtrusive
 - in general few issues arise, jobs come, jobs go
 - UK VO support works well
- However, some problems copying data back from some UK Tier-2's
 - remote lcg-cp from WN
 - seems to be related to use of Network Address Translation
 - UK site-admins have helped LHCb debug this
 - recently resolved at Glasgow and Liverpool by tuning WN TCP parameters

- GGUS/savannah ticketing systems work well in the background
- Change from SAM testing to Nagios has been disruptive just when Tier-2's are concentrating on first data
- 'Gstat2' is taking some getting used to - not clear how to easily get info that 'gstat1' used to provide (e.g. state of queues, VO tag info)
- Tier-2's always appreciate clear and timely instructions from the VO's
- Good middleware is important
- WLCG could benefit more from the expertise at Tier-2's
 - not very easy to feedback ideas for middleware improvements - sys-admins often focused on their site - could benefit WLCG project as a whole by being better engaged
 - what about having middleware developer/sys-admin workshops?

- How to run chaotic analysis well at a Tier-2?
 - Users themselves make mistakes e.g. user submitted thousands of jobs with wall-time of about 10-20 minutes, result: low job efficiency
 - Need to educate users about characteristics of a good job
 - Challenging to separate 'site' from 'user' error
- Just because 'no complaints' doesn't mean a site is doing well
 - Inevitably a large gap between the best and the least effective sites
 - Users gravitate to better sites
- Sites need to work proactively with VO's and users to get the most out of their site
- VO site black lists offer a lot of valuable customer feedback
 - Can the fact that it has been blacklisted by a user be fed back to a site?

- UK Tier-2's are providing a good service to the WLCG
 - Data arriving on site, jobs run, work gets done
 - Communications with WLCG and with VO's are effective
 - GridPP storage support is excellent
 - Some issues related to pressure on site infrastructure: firewalls, machine rooms, WAN rates etc
 - Always room for improvement of course
- Good VO monitoring pages are very useful for a site
- Remote stage-out can be problematic

- Now we need to tune sites for optimal running
 - Requires sites to be proactive
 - Understand their site well and have comprehensive monitoring in place
 - Work closely with the LHC VO's computing teams
 - React rapidly to problems and communicate well with the users
 - Remember that Tier-2's are providing a service for the LHC VO's