CSCS
Swiss National Supercomputing Centre

Tier 2 site report: CSCS

Gianfranco Sciacca (Bern)

CSCS report (June)

CPU:

(inc. WLCG and NorduGrid)

	Nr of jobs	Walltime (h)	CPU time (h)
Total	**208392** (187430)	**444872** (414442)	**405024** (381599)
PROD	**49199** (54028)	**302710** (238275)	**289255** (238275)
PILOT	**158359** (132114)	**142100** (165384)	**115767** (138392)
User	**7** (1)	**0** (0)	**0** (0)

May 2011 figures in brackets

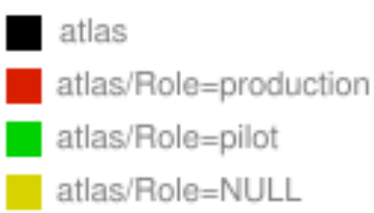

- Very similar to previous month, slight increase for ATLAS figures
- Breakdown per CE: cream 87.7%, ARC 12.1% (jobs); cream 50% ,ARC 49.1% (WT)
- Analy/Prod rate: 47.6%
- **CPU Efficiency: 91.1% [85.5% cream; 96.8% ARC]** (92.1 in May; 89.6% Efficiency for all VOs)

CSCS report (June)

Disk: http://bourricot.cern.ch/dq2/accounting/t2_reports/FZKSITES/

Report for FZKSITES (UTC 2011-07-01 09:24:36.899487)

Site	Used(G)	Free(G)	Total(G)	%
CSCS-LCG2_DATADISK	227806	204193	431999	52
CSCS-LCG2_GROUPDISK	17095	32905	50000	34
CSCS-LCG2_HOTDISK	877	2123	3000	29
CSCS-LCG2_LOCALGROUPDISK	9533	467	10000	95
CSCS-LCG2_PRODDISK	821	9179	10000	8
CSCS-LCG2_SCRATCHDISK	20590	7409	27999	73
Totals (TB) (April figures)	**276.7** (210.6)	**256.3** (322.4)	**533** (533)	51.9

- No news of relevance
- Increase of disk usage (good!)
- Online: 533TB / Pledged: 468.5TB
- dCache 1.9.5-26

CSCS report (June)

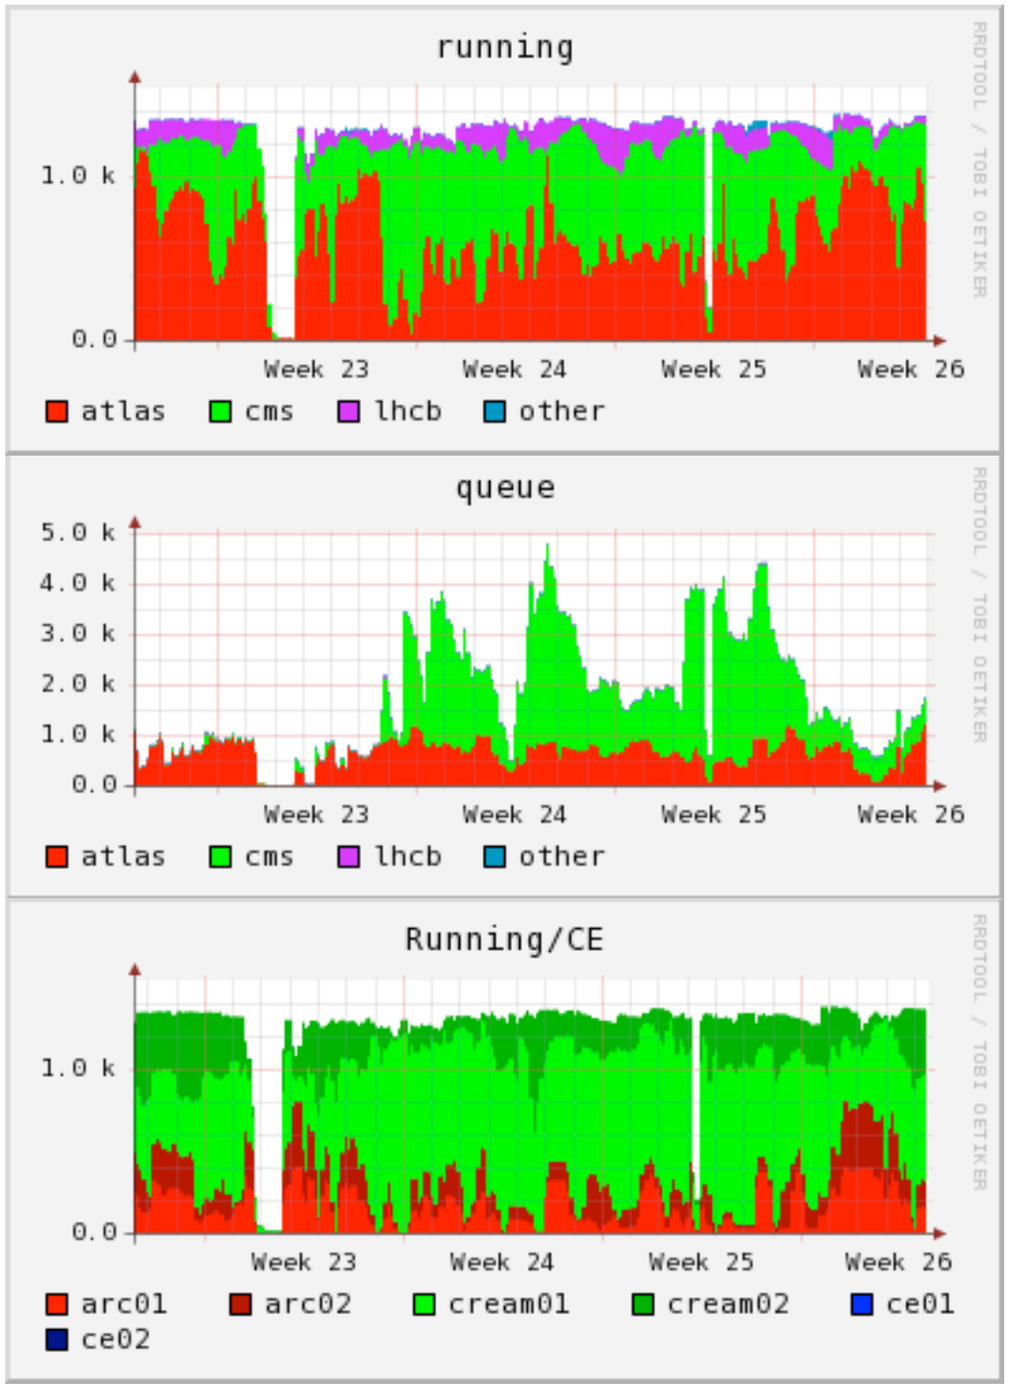

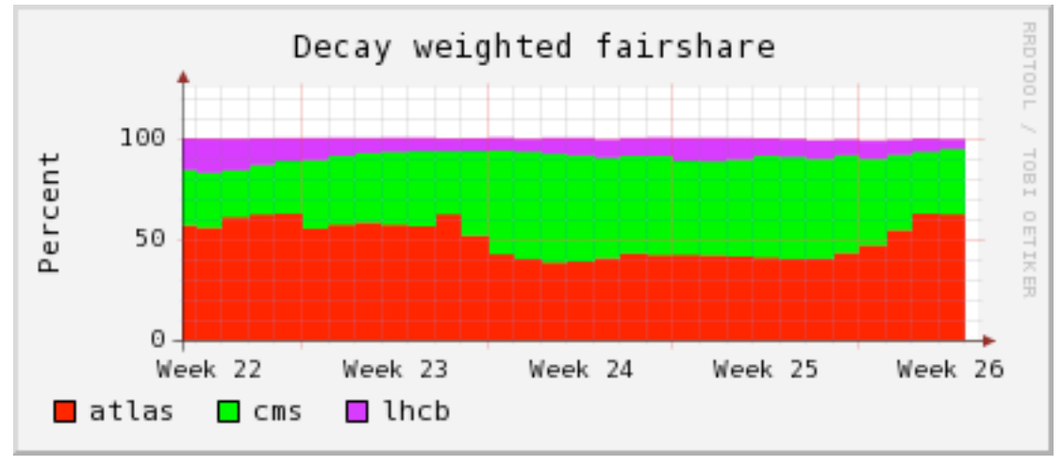

Target fair-shares:

atlas=40:cms=40:lhcb=20

Measured fair-shares:

atlas=49.9:cms=41.6:lhcb=8.5

CSCS report (June)

General news

- Torque upgrade to 2.4.13 (might solve problems with HA instabilities of 2.4.11)
  Current not running in High Availability mode
- Deployed ARGUS authorization service
- Installation of EMI1 CREAM-CE on cream02 (early adopters)
- Enabled glexec capability on CreamCEs (now advertised via the BDII)
- Installation of glite-GLEXEC_wn on the Worker Nodes

CSCS report (June)

HammerCloud tests

http://hammercloud.cern.ch/atlas/10004095/test/

(Standard muon trigger analysis job using recent mc10*AOD* datasets)

highest event rate (25.7), average CPU/WT (67.5% vs 66.4%) - 100% success

(no appreciable rise in load on shared file system)

- Need to better understand discrepancy between event rate and CPU/WT performance
- Provided this is mainly due to the job preparation time (pre and post), then some improvement is possible (input files, download from where? copy output, to SCRATCHDISK? Fetch Panda job peak at 100?)

http://hammercloud.cern.ch/atlas/10004292/test/

(Root-SUSYD3PD analysis which runs over some 50 datasets - test suffered from many small datasets with few files)

below average event rate (46.4), below average efficiency (49.2% vs 58.7%) - 100% success

http://hammercloud.cern.ch/atlas/10004978/test/

(Improved version of previous test) - Missed it because PANDA queue still offline following maintenance downtime

http://hammercloud.cern.ch/atlas/10005454/test/

Better than average efficiency (27.1% vs 23.9%), average event rate (but much higher than 10004978 : ~300 vs ~170) - 100% success

CSCS report (June)

Tier-3 activities

No Tier-3 activities at CSCS

Swiss ATLAS Tier-3's at Bern (456 cores) and Geneva (342 cores), but affiliated with ND cloud (ARC CE's)

Bern: production (Panda: ARC-T2), analysis (Panda: ANALY_ARC), hi-prio tasks (Panda: ARC), local (no direct batch submission, via grid only)
- 180TB in LOCALGROUPDISK (in ND cloud) (rfio access from interactive machines), 40TB NFS
- Likely to add 768 cores in spring 2012

Geneva: production (Panda: ARC-T2), local analysis
- 191TB in LOCALGROUPDISK, SCRATCHDISK, TRIG-DAQ tokens (in ND cloud)

- Essentially: Tier-2 like activities. Both sites in GOCDB and NGI_CH (same as CSCS)

- We've never ruled out the possibility to aggregate some of the T3 resources with those of CSCS (move to a federated T2 model within the German cloud)

- But we haven't found a strong enough motivation yet

- Open to future re-thinks