

Tier 2 site report: CSCS

Gianfranco Sciacca (Bern)


CSCS report (February)

CPU:

<https://mon.lcg.cscs.ch/pbsplots/pbsplots.py?> (inc. WLCG and NorduGrid)


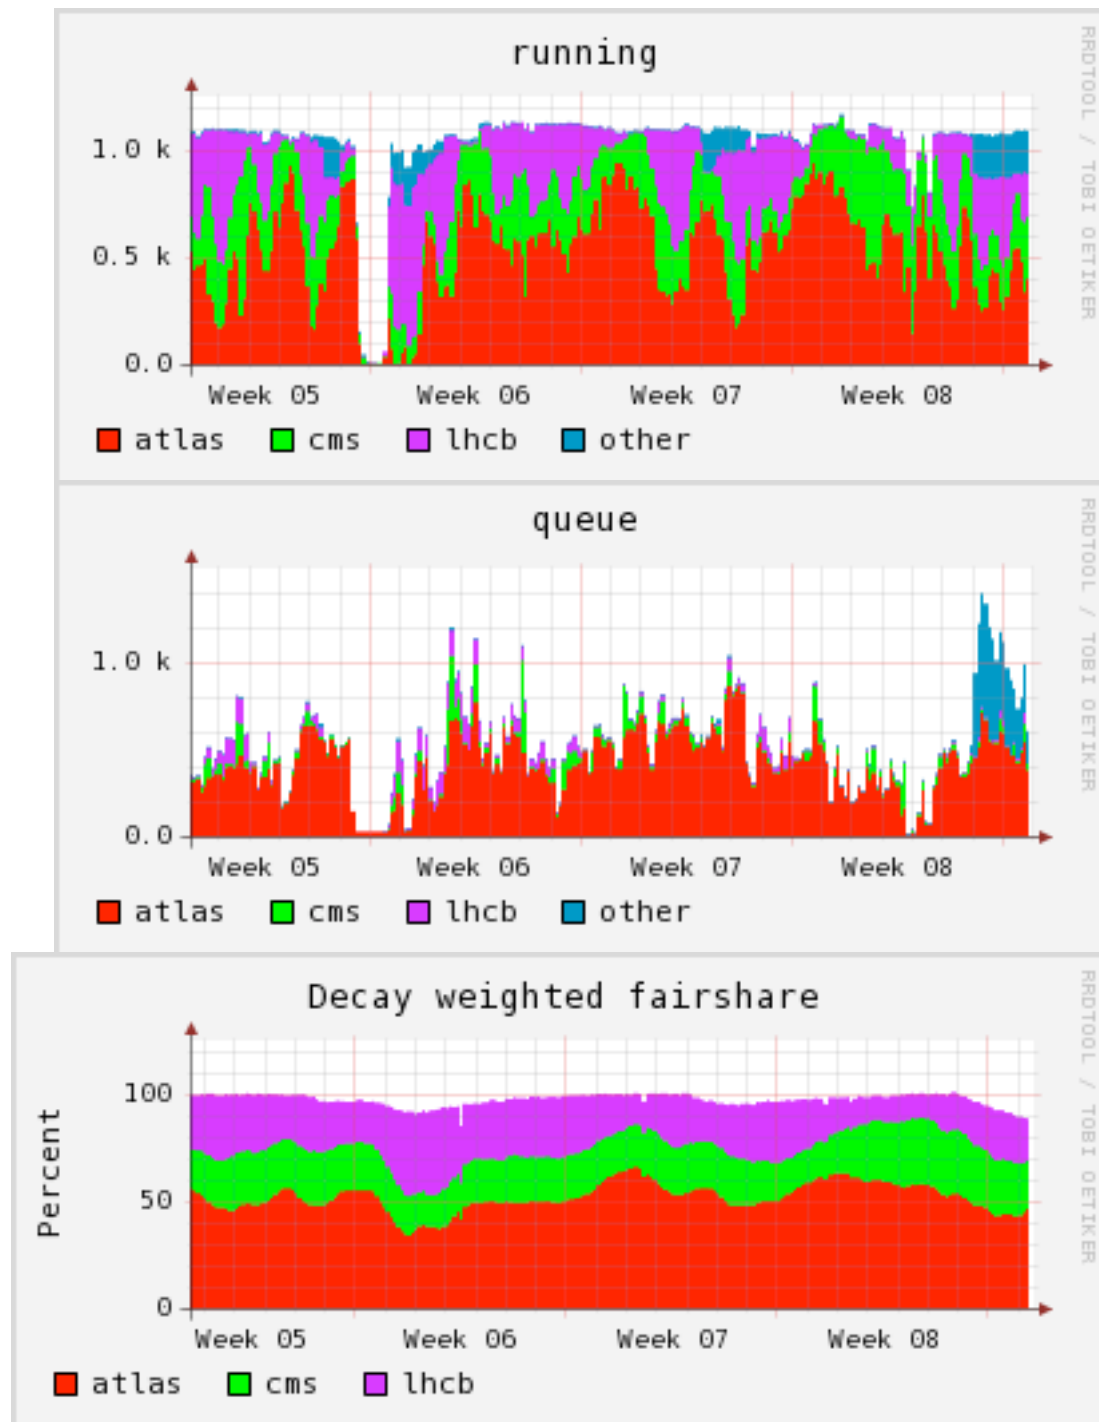
	Nr of jobs	Walltime (h)	CPU time (h)
Total	106313 (135530)	352811 (333435)	297207 (299318)
PROD	52856 (39360)	196696 (270982)	178447 (250608)
PILOT	44782 (87827)	154656 (59705)	118498 (47821)
User	8675 (8343)	1460 (2748)	262 (889)

January figures in brackets



- **CPU Efficiency: 84.3** (89.8 December)
- Compared to previous month, similar total usage, but more than doubled analysis volume
- Fraction of total jobs: cream 23.2%, ARC 24.5%, lcg-CE 51.3%

CSCS report (February)



Nominal fair-shares:
atlas=40:cms=40:lhcb=20

CSCS report (February)

Disk:

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

Report for FZKSITES (UTC 2011-03-01 11:55:03.247962)

Site	Used(G)	Free(G)	Total(G)	%
CSCS-LCG2_DATADISK	154091	180908	334999	45
CSCS-LCG2_GROUPDISK	17095	22905	40000	42
CSCS-LCG2_HOTDISK	824	176	1000	82
CSCS-LCG2_LOCALGROUPDISK	9533	467	10000	95
CSCS-LCG2_PRODDISK	1013	8987	10000	10
CSCS-LCG2_SCRATCHDISK	12589	5410	17999	69

Totals (TB)

(previous month)

195.1 218.8 414

(186.2)

(127.8)

(314)

47.1

- New disk came online on 2 Feb 2011 (100TiB)
- ~50TB should additionally come in the next few weeks

CSCS report (February)

General news

- Smooth operation during February
- Upgrade of cream01 to 3.2.10-0 just completed
- cream02 upgrade on 7 March (downtime announced in GOCDB for the node)
- Moab scheduler updated to v 6.1 (decouples client queries from scheduling process)
- New hardware to be put online (will likely have this online before next meeting):
 - 2000 HEP-SPEC06 (~20% increase in capacity)
 - Total capacity will be 13488 HEP-SPEC06 with fair shares atlas=40:cms=40:lhcb=20 (AMD Opteron 6172 @ 2.1GHz with 3GB of memory per job slot)
 - New 10TB GPFS shared file system (in addition to the existing 10TB Lustre area)
 - ~50 TB more storage
- Establishing timeline for full decommissioning lcg-CEs, will need input from ATLAS on this

CSCS report (February)

General news

- Observed sizeable number of (analy) jobs killed for exceeding 3GB vmem limit (limit is $1.5 * \text{required_mem}$ with `required_mem` defaulting to 2GB)
 - => In February estimated ~5% of ATLAS WallTime wasted
 - => Not a huge problem per se, but undesirable nevertheless (specially from user point of view)
 - => Some WN HW peculiarity involved
 - current WNs with 2GB RAM/job slot; no local disk; use shared FS for scratch; almost no local space for swap
 - new WNs will have 3GB RAM and 3GB swap per job slot
 - => Got valuable feedback from Rod about real life memory usage of jobs (also feedback from sites who apply limits has been appreciated)
 - => We are discussing a possible solution to prevent this from occurring again in the future
 - e.g. creating a shared swap space for the “old” WNs as a safety net, could allow raising the vmem limit to $\geq 4\text{GB}$ (New WNs come with their own swap per node on local disks)
 - => Question: is it likely (or technically possible) to expect in the longer term that jobs will be submitted with a memory request that is passed through the MW to the batch system?