Tier 2 site report: CSCS

Gianfranco Sciacca (Bern)

CSCS report (November)

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

	Nr of jobs	Walltime (h)	CPU time (h)
Total	**211136** (264262)	**379540** (353408)	**338439** (297434)
PROD	**58766** (55431)	**307351** (307555)	**288080** (268417)
PILOT	**143424** (210675)	**67561** (40666)	**49495** (27246)
User	**8946** (7156)	**4628** (5186)	**864** (1771)

October figures in brackets

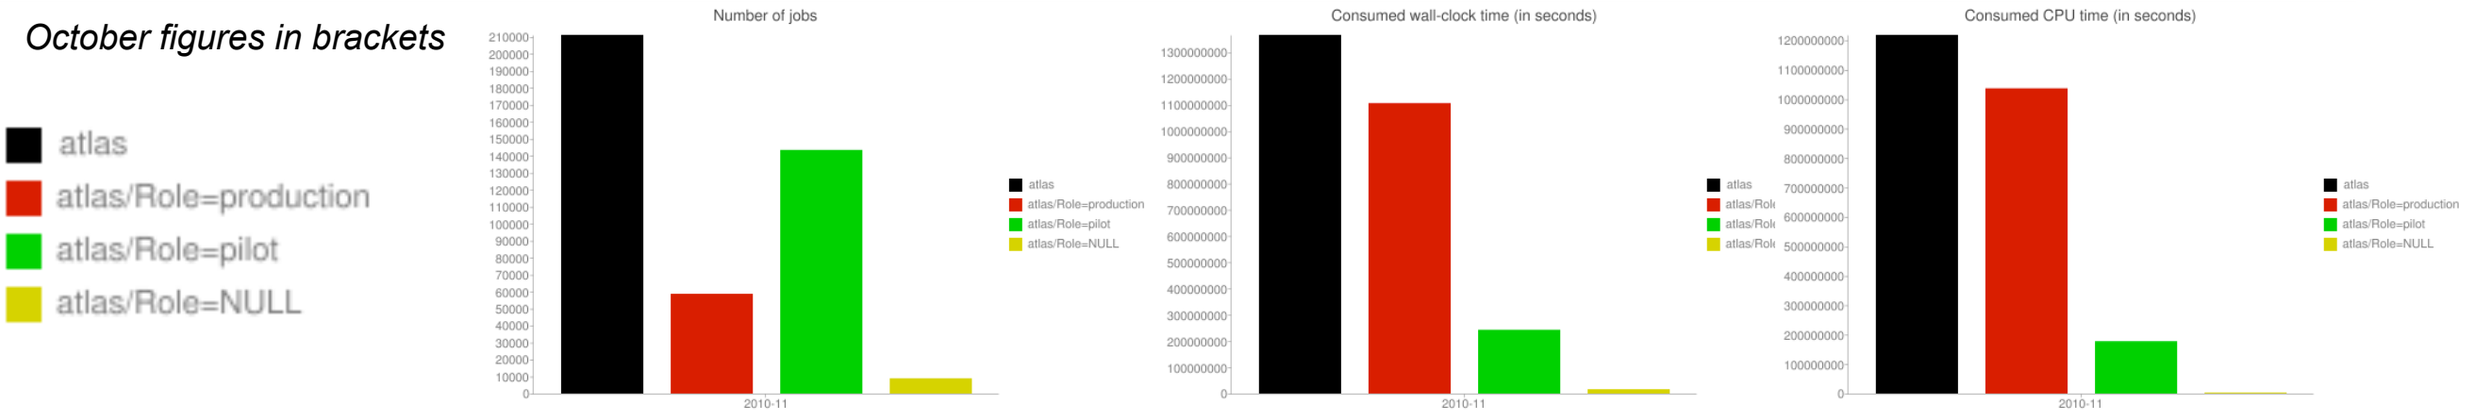

- **Efficiency: 82.0** (83.2 September)
- Still large number of pilots not picking up payloads
- Slightly increased levels (despite some hiccups), analysis still low (but probably in line with other sites)
- ~15% of total jobs through NG, accounting for >50% of total WT (not showing in Panda DE)

CSCS report (November)

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

Report for FZKSITES (UTC 2010-12-02 15:30:03.618009)

Site	Used(G)	Free(G)	Total(G)	%
CSCS-LCG2_DATADISK	117075	22924	139999	83
CSCS-LCG2_GROUPDISK	15798	24202	40000	39
CSCS-LCG2_HOTDISK	812	188	1000	81
CSCS-LCG2_LOCALGROUPDISK	9022	978	10000	90
CSCS-LCG2_MCDISK	37463	57537	95000	39
CSCS-LCG2_PRODDISK	834	9166	10000	8
CSCS-LCG2_SCRATCHDISK	5796	12203	17999	32
Totals (TB)	**186.8**	**127.2**	**314**	
(previous month)	(98.2)	(215.8)		

- Almost doubled amount of data compared to last month (good!)
- Still 60TB short of the pledge
- New disk commissioning under way: 110TB (January 2011)
- Plan to get fully aligned with pledges by April 2011
- GROUPDISK became full on 26 Nov, but ~25TB disappeared over the w/e
- 10TB missing from GROUPDISK (prefer 5+5TB from PROD+SCRATCH if needed)

CSCS report (November)

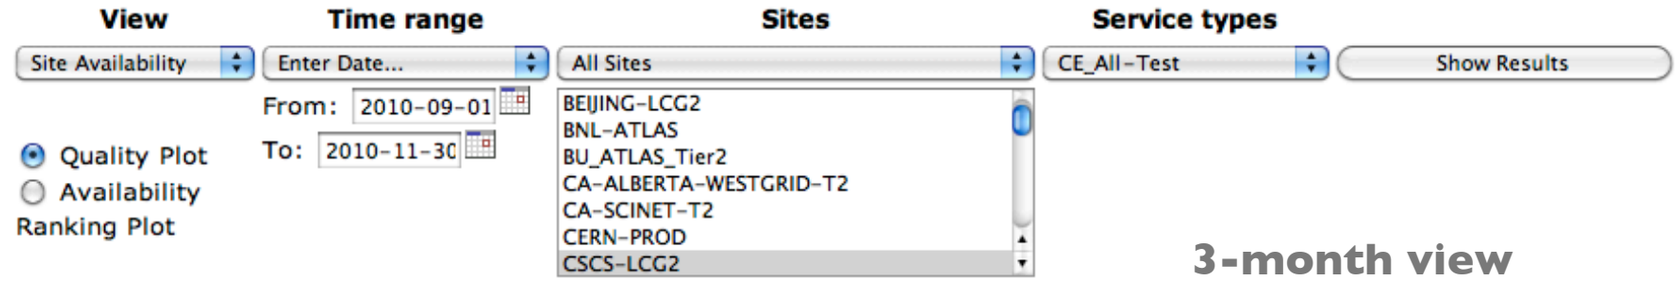

3-month view

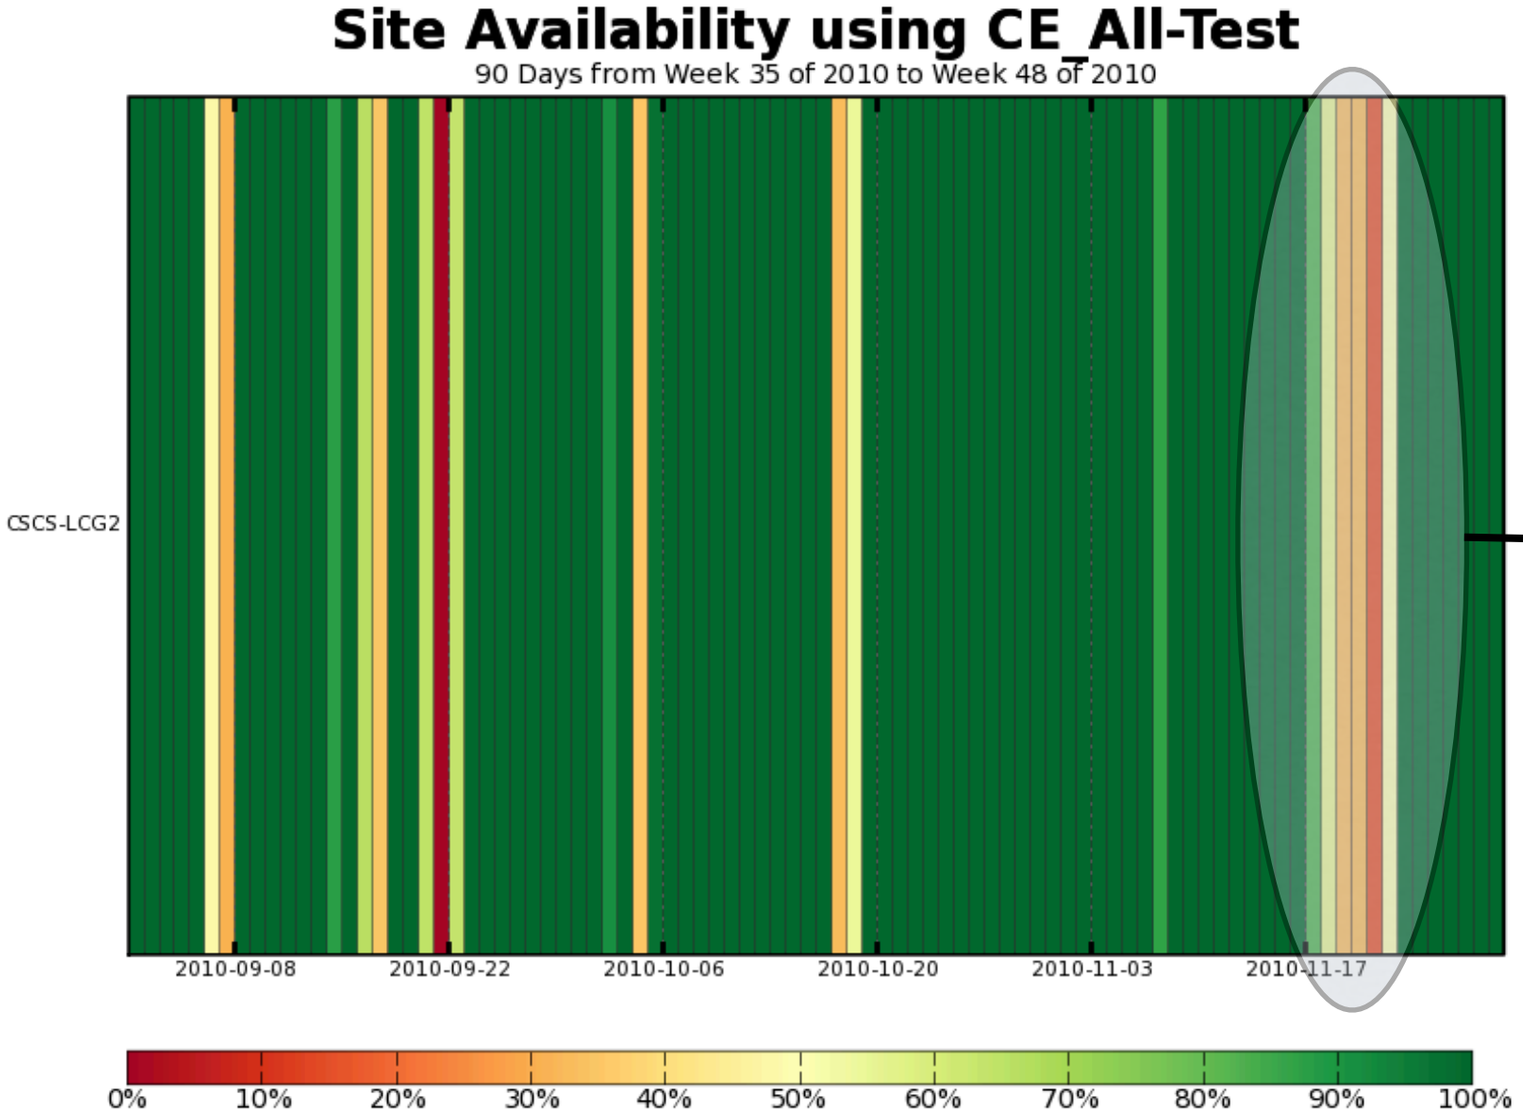

17 Nov: dCache dies overnight, fixed next morning early but delay in bringing the site back in Production

18 Nov: Redundant NFS dies overnight

(combination of disk-controller firmware and driver)

19 Nov: more dCache woes

all gridftp doors that run in the Solaris pools die after some time, with a very generic error message:

```
Got an IO Exception ( closing server ) :
java.net.SocketException: Invalid argument
```

22 Nov: problem identified, site back online, started moving all Solaris pools to Linux (not affecting uptime)

✳ NG jobs kept running all along

Site availability (CE_All-Test): **dCache/NFS problems caused downtime** (over w/e)

CSCS report (November)

General news

- More HC stress tests in early November: Panda direct I/O dcap (left) vs File-Stager (right)
- We're inclined to switch from direct I/O dcap to File-Stager, but no decision taken yet

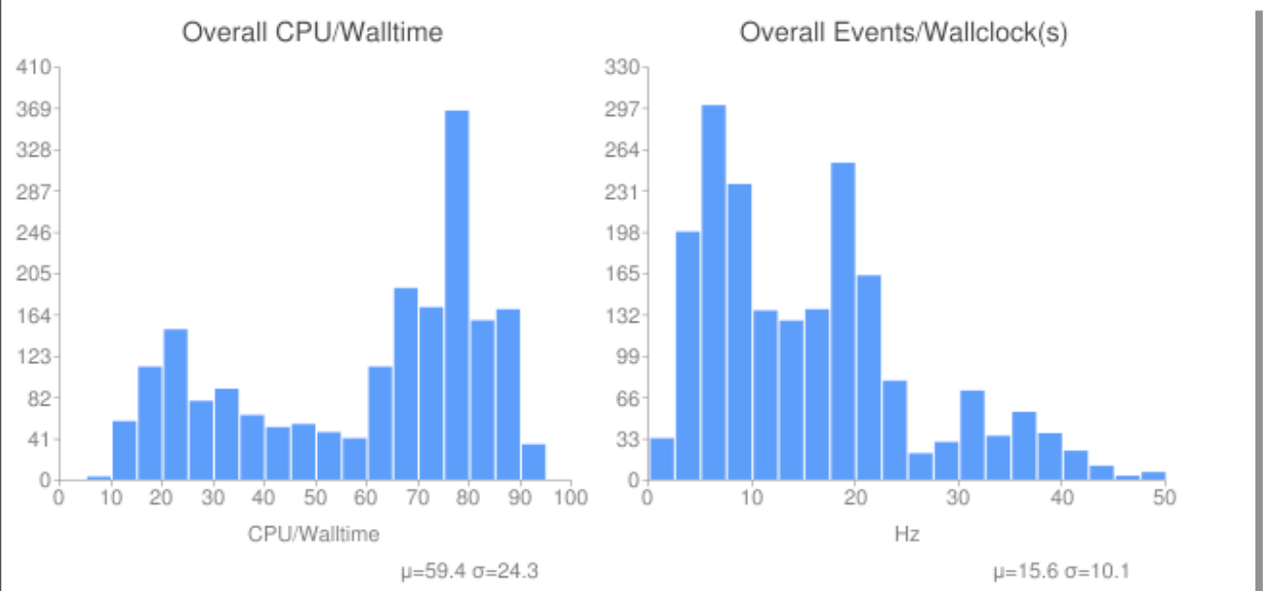

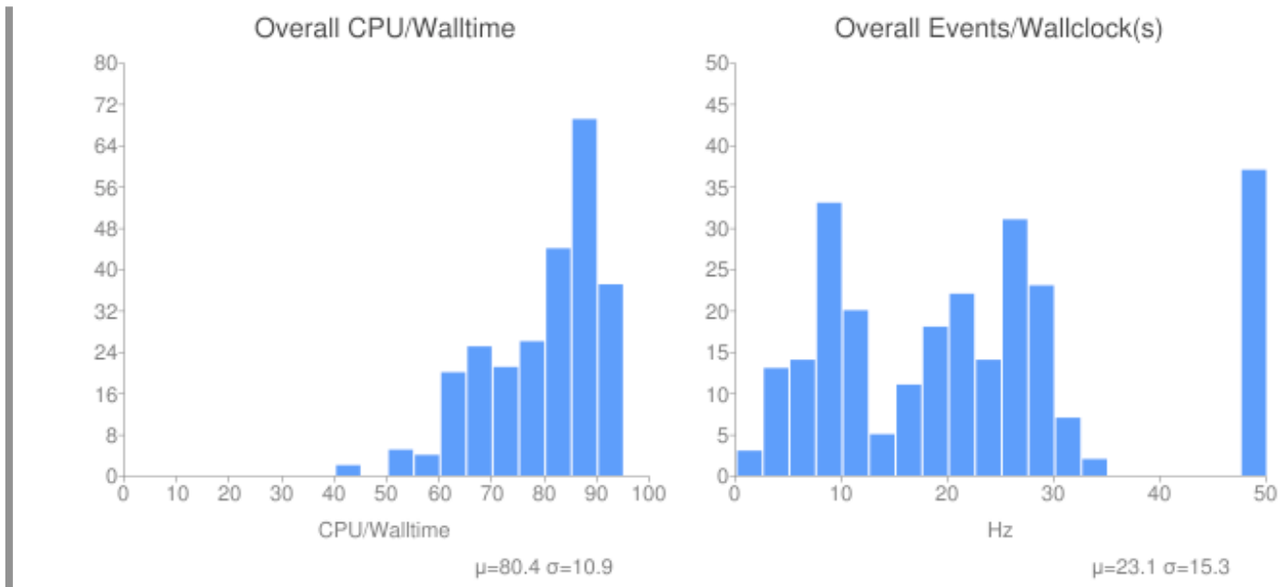

Additional Issues

- Jobs using up to 4GB vmem caused disruption to several WNs (through ANALY)
- Limits now introduced: 1.2*requested or 2GB default
- Feedback welcome: are there any good reasons not to introduce limits (or different levels)?
- Do jobs actually come in with a memory request ?