Tier 2 site report: CSCS

Gianfranco Sciacca (Bern)

CSCS report (October)

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

	Nr of jobs	Walltime (h)	CPU time (h)
Total	**264262** (130246)	**353408** (410101)	**297434** (321121)
PROD	**55431** (50646)	**307555** (267709)	**268417** (209158)
PILOT	**201675** (73830)	**40666** (137851)	**27246** (110463)
User	**7156** (5770)	**5186** (4541)	**1771** (1500)

September figures in brackets

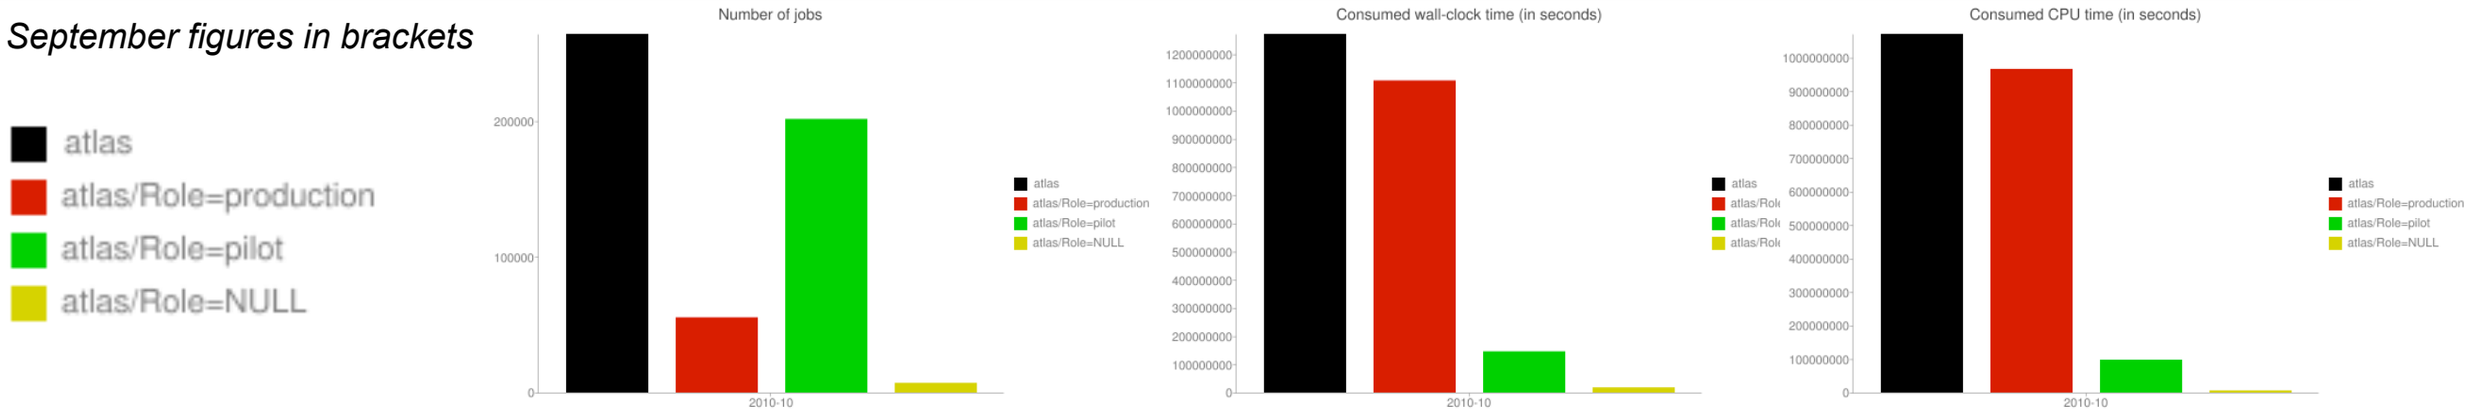

- **Efficiency: 83.2** (81.4 September)
- Twice as many jobs, lower CPU time
- Considerably lower analysis volume (concern), more production

CSCS report (October)

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

Report for FZKSITES (UTC 2010-11-02 19:09:21.360748)

Site	Used(G)	Free(G)	Total(G)	%
CSCS-LCG2_DATADISK	47847	92152	139999	34
CSCS-LCG2_GROUPDISK	13282	26718	40000	33
CSCS-LCG2_HOTDISK	781	219	1000	78
CSCS-LCG2_LOCALGROUPDISK	8644	1356	10000	86
CSCS-LCG2_MCDISK	24754	70246	95000	26
CSCS-LCG2_PRODDISK	757	9243	10000	7
CSCS-LCG2_SCRATCHDISK	2129	15870	17999	11
Totals (TB)	**98.2**	**215.8**	**314**	

- 10TB moved from GROUPDISK to LOCALGROUPDISK
- Slight impact on deployed/pledged
- New disk commissioning under way: 110TB (January 2011)
- Plan to get aligned with pledges by April 2011
- Storage still considerably underused (effect on analysis volume)
- Hard to motivate funding agencies if disk not used
- Need to understand data distribution (effect of new dynamic placement?)
- Will work on GROUPDISK, but what about DATADISK and MCDISK?

CSCS report (October)

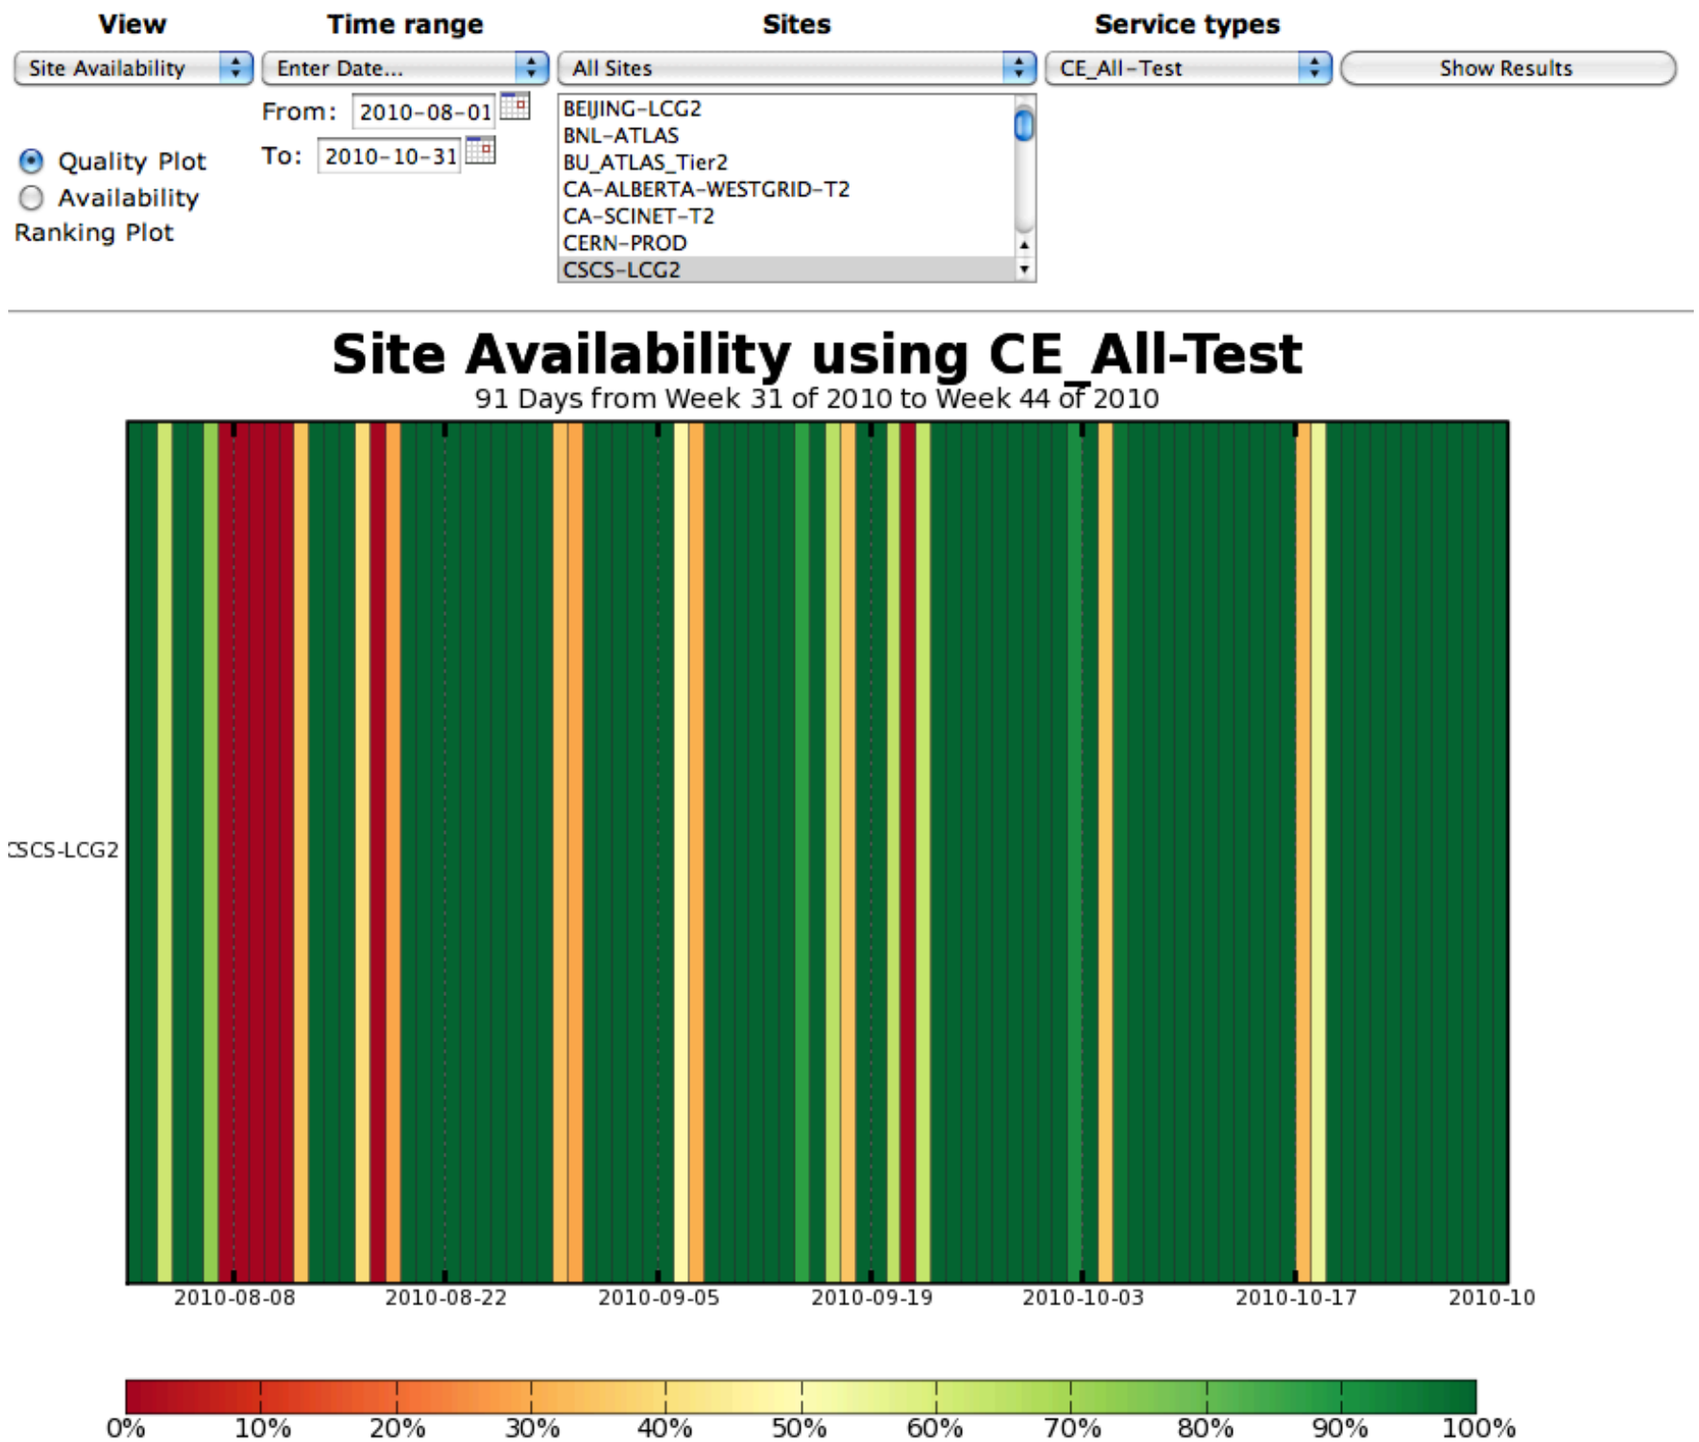

Site availability (CE_All-Test): **smooth operation last month**

CSCS report (October)

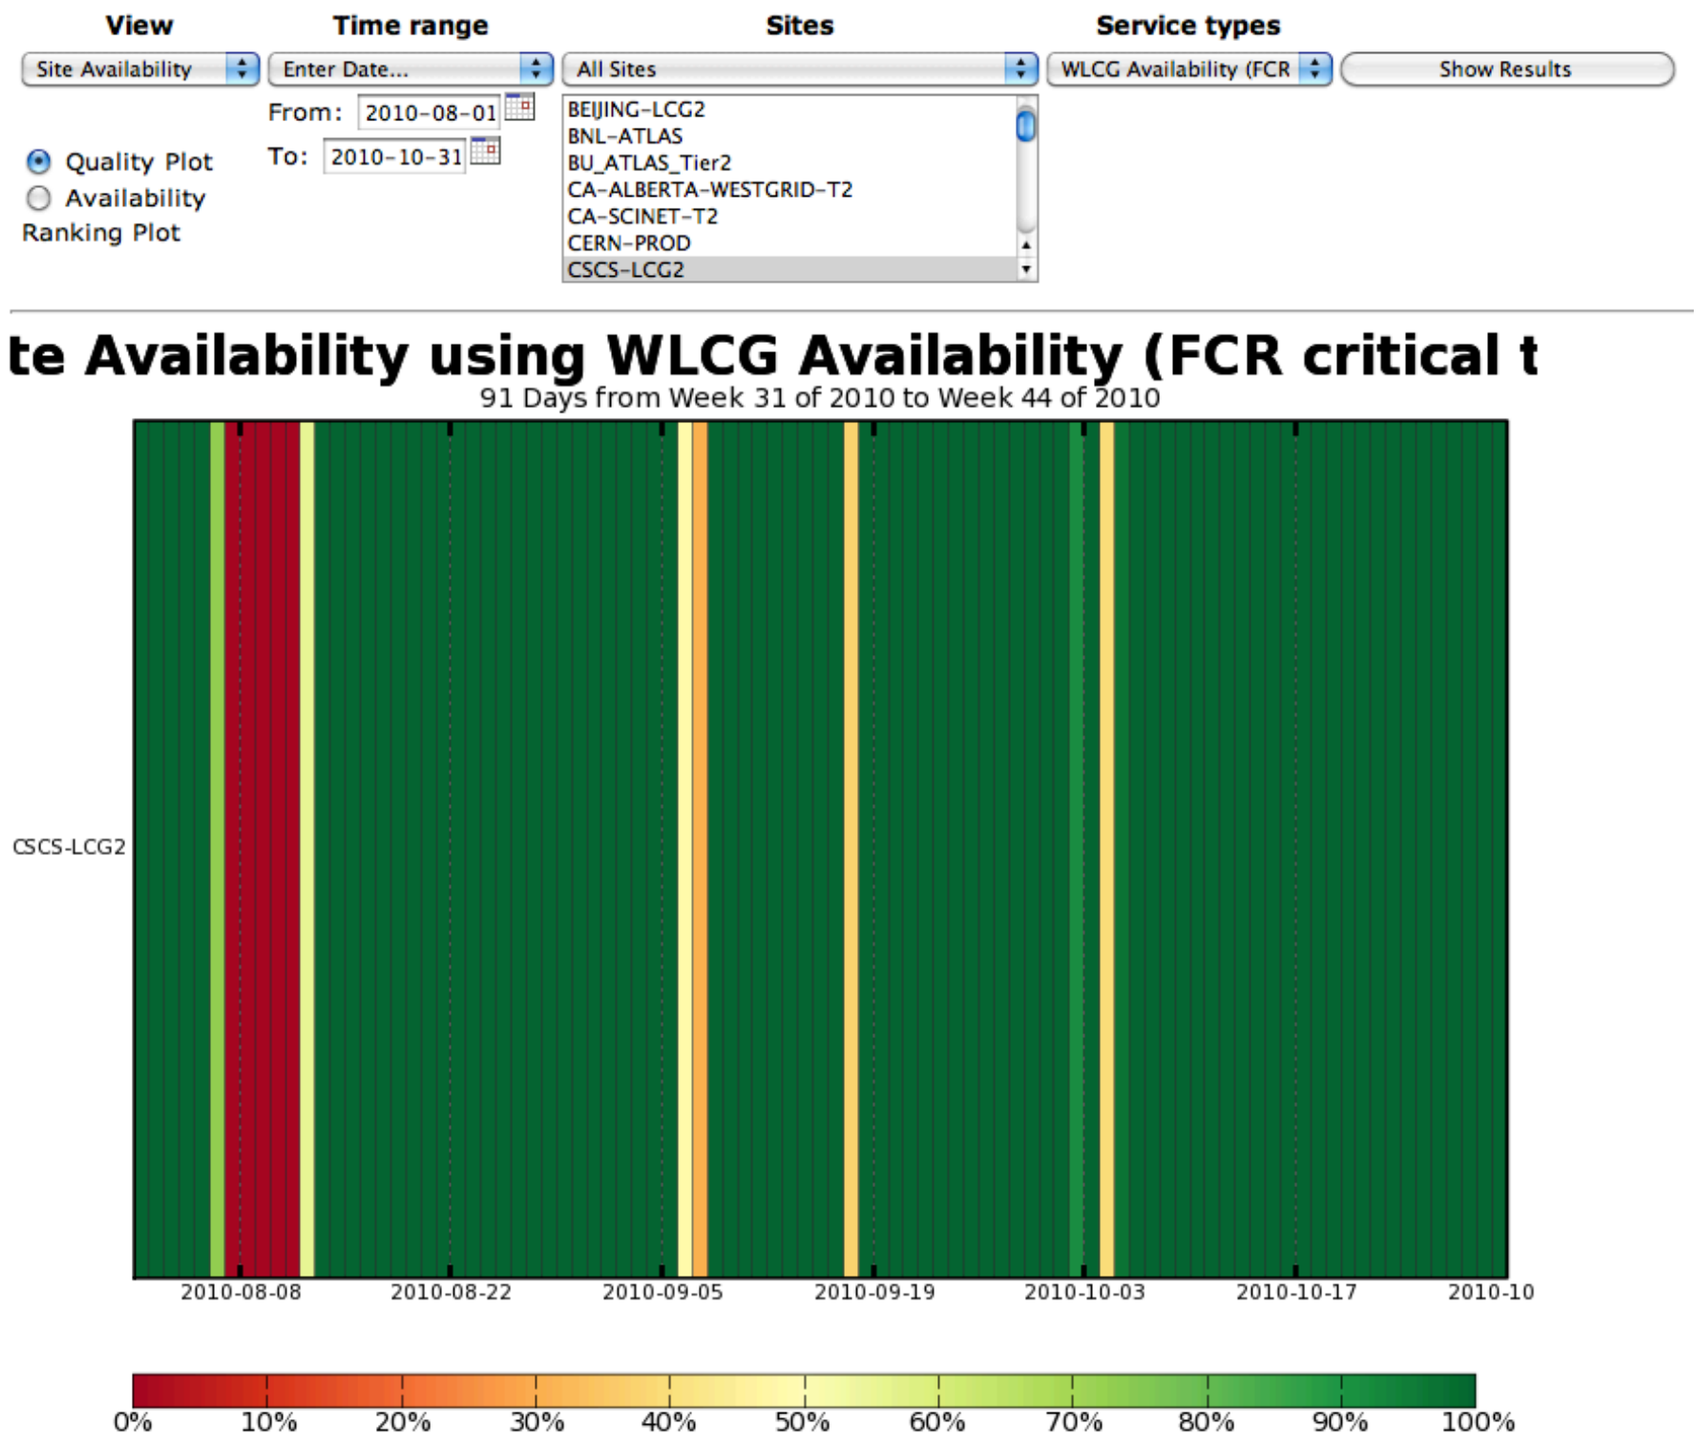

Site availability (WLCG - FCR critical tests): **smooth operation last month**

CSCS report (October)

HammerCloud tests

Test	CPU/WT	Event Rate	Software Setup	Download / Prepare Input	Athena time	Output Storage	Events/ Athena (s)	URL
Panda default (Direct I/O dcap)	55.9	15.1		38.8	4400	22.3	15.8	http://hammercloud.cern.ch/atlas/10001384/test/
Panda File Stager	66.5	19.6		52.2	3124	13.1	20.8	http://hammercloud.cern.ch/atlas/10001450/test/
WMS DQ2_Local (Direct I/O dcap ??)	79.5	19.9	45.1	123.1	5131	111.1	19.8	http://hammercloud.cern.ch/atlas/10001421/test/
WMS File Stager	77.5	20.6	75.6	160.8	4843	121.6	20.5	http://hammercloud.cern.ch/atlas/10001429/test/

- Default method: lowest performance
- Other methods ~ equivalent
- But wrong access (gsiftp) for WMS File Stager due to misconfiguration at site
- Same for Panda File Stager?
- WMS DQ2_Local: is this Direct I/O over dcap?

CSCS report (October)

HammerCloud tests

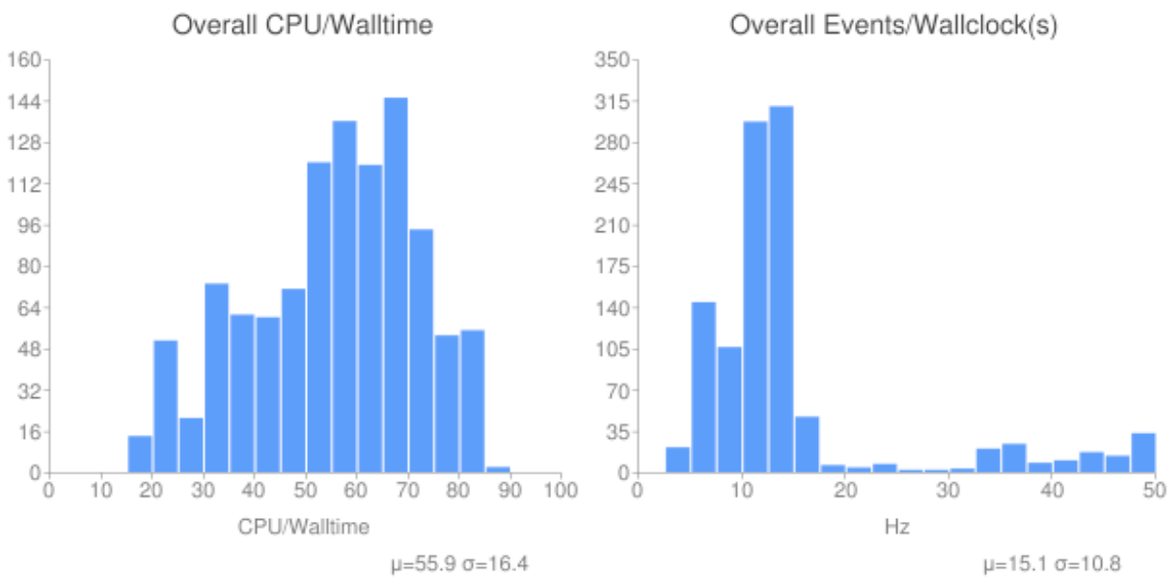

Panda default (Direct I/o dcap)

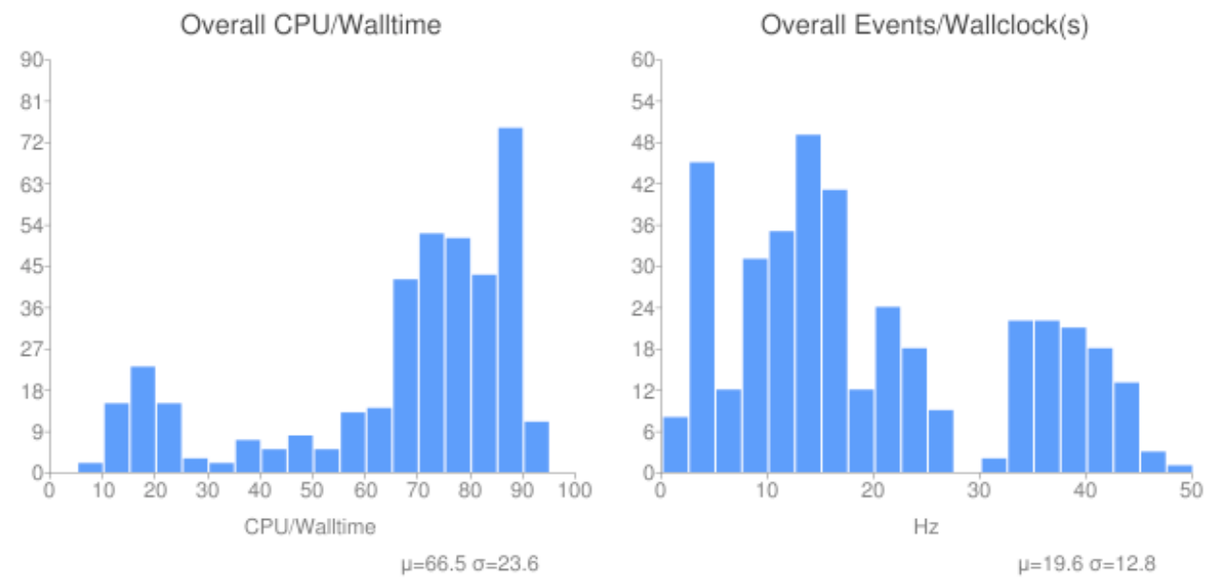

Panda File Stager

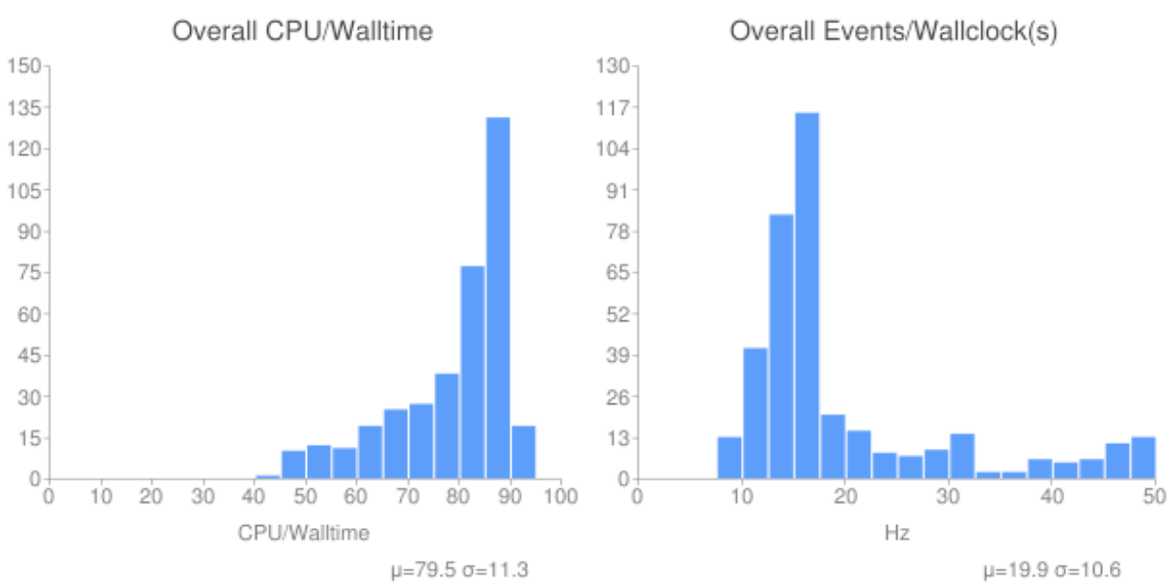

WMS DQ2_Local

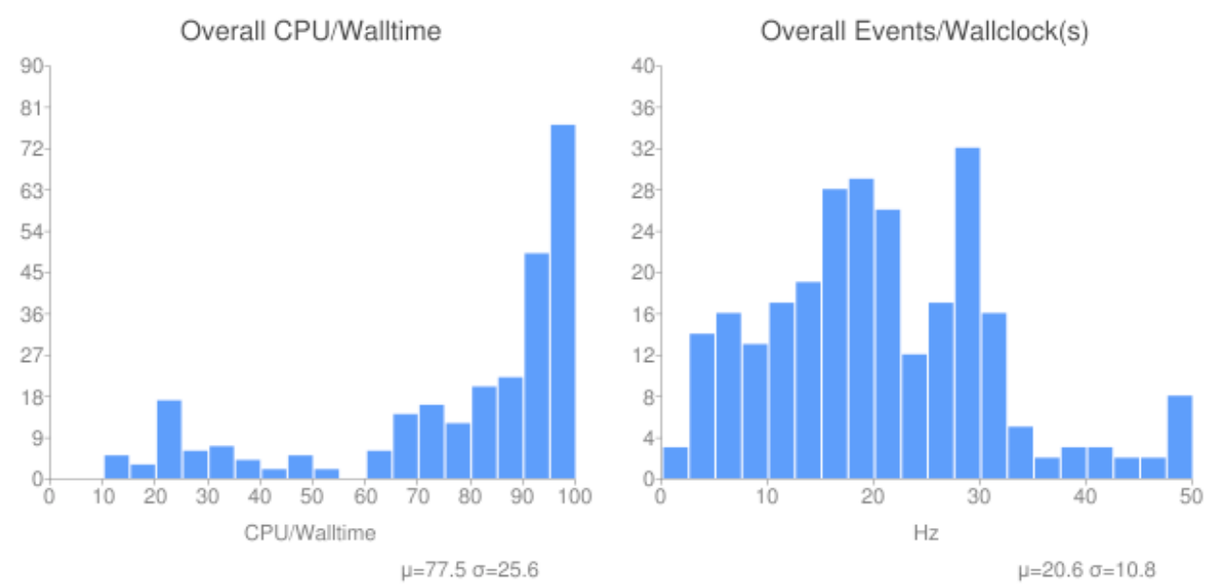

WMS File Stager

CSCS report (October)

General news

- Assorted security patching went smoothly (no downtimes)
- At CSCS, monthly maintenance window (1st Monday, usually)
- Reduced ATLAS queue closing time (before scheduled downtime) to 20h (down from 60h)
  WT in VO card: 41.6h
  Over last 5 months 95% of jobs finished with WT < 19h
- New High Availability setup for Torque/Moab

Open Issues (recap)

- Low levels of analysis are of concern
- Low occupancy of storage (DATADISK 34%, MCDISK 26%)
- HC: default access (direct I/O dcap) lowest performer. To what extent should we worry?