

GridPP

UK Computing for Particle Physics

CMS Multicore jobs at RAL


Andrew Lahiff, RAL

WLCG Multicore TF Meeting


1st July 2014

- **Reminder**
 - RAL supports all LHC VOs
 - Batch system is HTCondor
- **Multi-core configuration**
 - Changes made since last RAL report at a multi-core task force meeting
 - Changed the order in which the negotiator considers groups
 - Consider multi-core groups before single-core groups
 - Multi-core slots are 'expensive' to obtain, so don't want to lose them too quickly
 - Changes made since CMS started submitting multi-core jobs
 - None


- CMS multi-core jobs over the past month
 - in addition to 1-3K running CMS single-core jobs




- Wasted CPUs due to draining
 - CPUs are wasted due to draining in order to provide multi-core slots
 - At most around $< 3\%$ of total number of CPUs



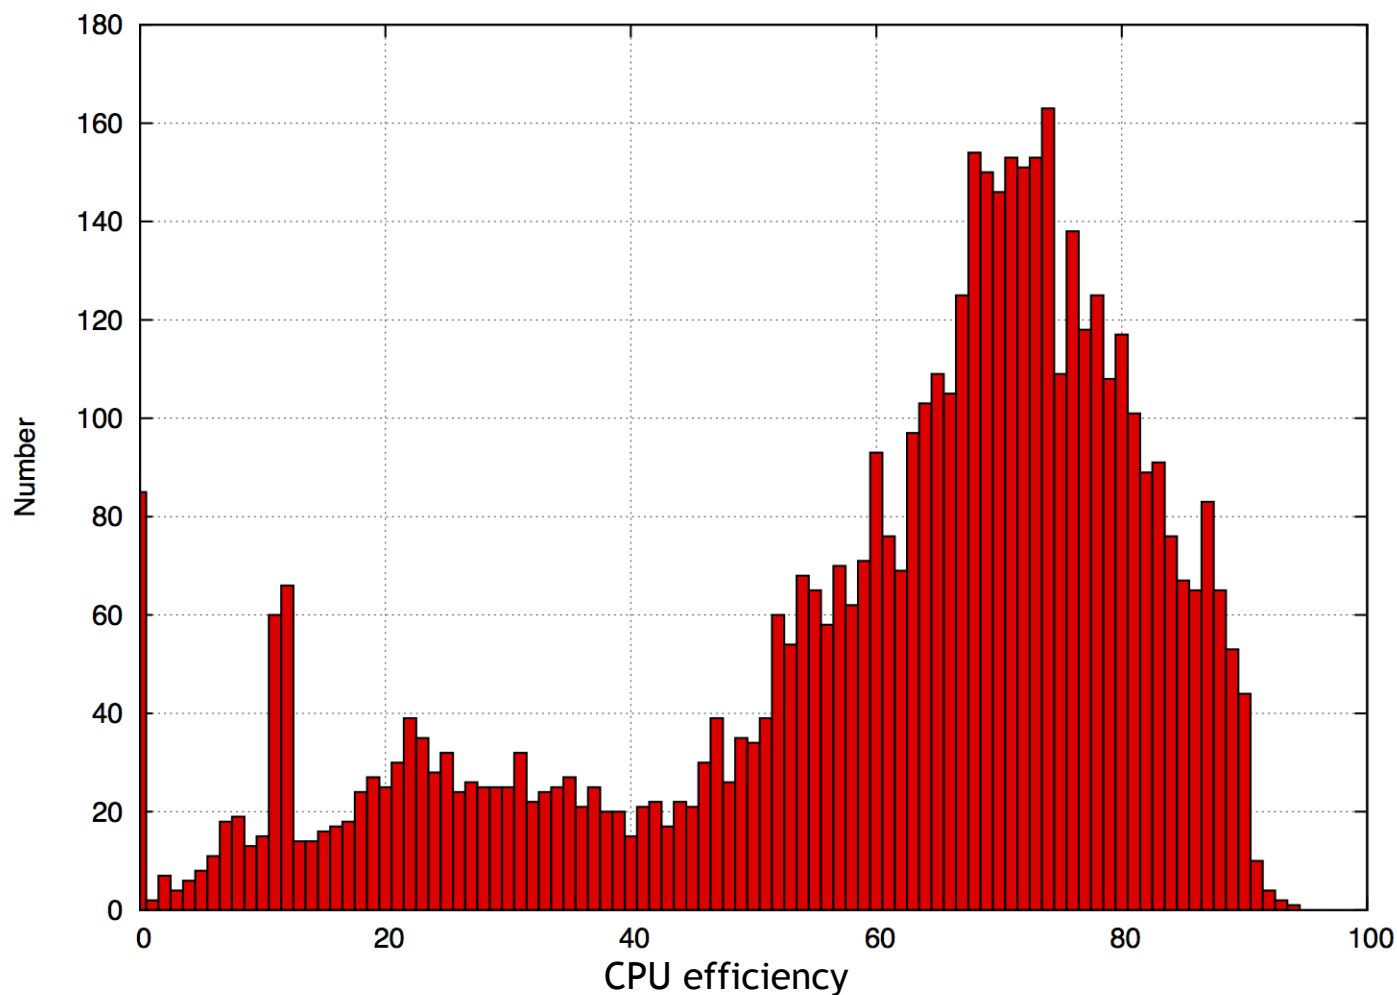
- Multi-core job wall times
 - Since 1st June



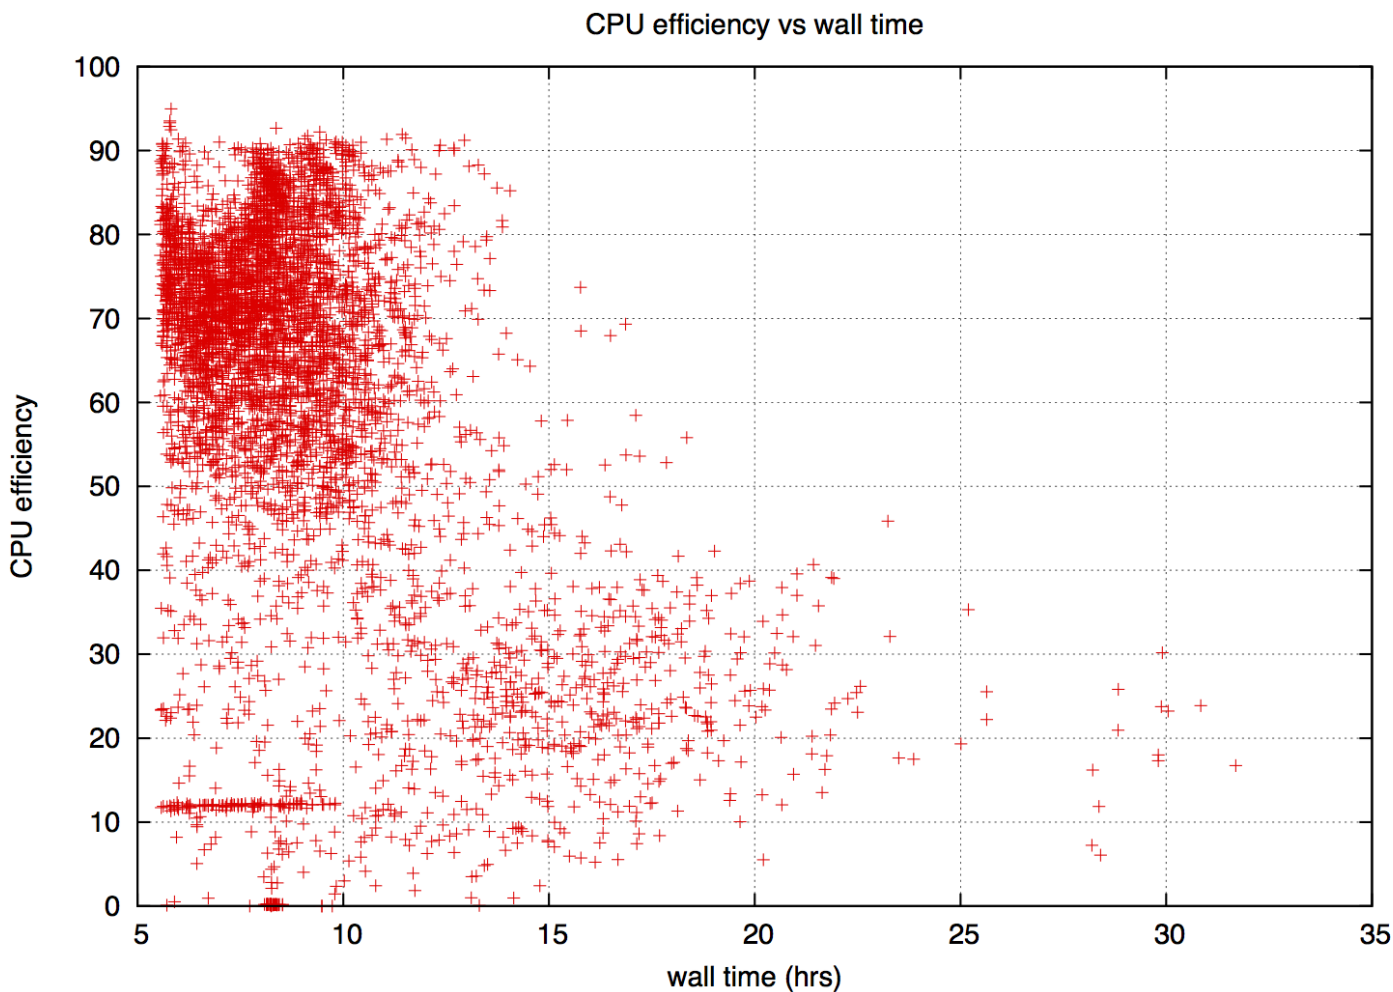
- CPU efficiency of single-core jobs
 - Since 1st June, jobs with wall time > 5 hours



- CPU efficiency of multi-core jobs
 - Since 1st June, jobs with wall time > 5 hours



- CPU efficiency vs wall time for multi-core jobs
 - The longest jobs generally have poor CPU efficiency




- Overall CPU efficiency for CMS – June 2014
 - Single-core jobs: 80.1%
 - Multi-core jobs: 58.2%

- CMS multi-core pilots are internally running single-core jobs
 - How many single-core jobs are running as a function of time?
- Simple analysis
 - Script running as cron on the CE every 20 mins which does:


```
condor_ssh_to_job <job id> "ls glide_*/execute | wc -l"
```


for all running CMS multi-core jobs
 - Gives number of single-core jobs running within the multi-core pilot
 - Use this data to reconstruct number of single-core jobs running as a function of time within each multi-core pilot

- Example job
 - Pilot ran for ~9.5 hours but only used all 8 CPUs for < 3 hours



- Another example job
 - Ran for over 4 hours using only 1 of the 8 requested cores



- CMS multi-core pilots running single-core jobs
 - Resources wasted in providing the multi-core slots
 - Resources wasted due to inefficient use of CPUs within the multi-core pilots
- Until multi-core jobs at sites are more common than single-core jobs, perhaps it's best for multi-core pilots to only run multi-core jobs?
 - Longer-lived pilots, however, would help with the above issues