

CMS Analysis Operations in STEP-09

Sanjay Padhi

University of California, San Diego (UCSD)

With

James Letts, Thomas Kress, Dave L. Evans, Dashboard
and CRAB developers, Stefano Belforte, Frank Würthwein

WLCG STEP'09 Post-Mortem Workshop, July 9th, 2009

Goals

To assess the readiness of the global Tier-2 infrastructure

From twiki:

<https://twiki.cern.ch/twiki/bin/view/CMS/Step09>

<https://twiki.cern.ch/twiki/bin/view/CMS/Step09T2>

Demonstrate analysis at a scale using all pledged resources at T2

- Close to 16,000 pledged slots
 - 8,000 for analysis (~ 50%)

Explore data placement

- measure how (much) the space granted to physics groups is used
- Replicate “hot” datasets as we expect to do in operations
- Monitor effect on job submission

Job Related Metrics

Weekly Metrics tables since May 1st, 2009.

- ◆ Took “baseline” during May
- ◆ Started additional jobs in June
- ◆ Doing accounting for 6/8 00:00 - 6/21 23:59

Present large tables on twiki.

Present distributions in this talk:

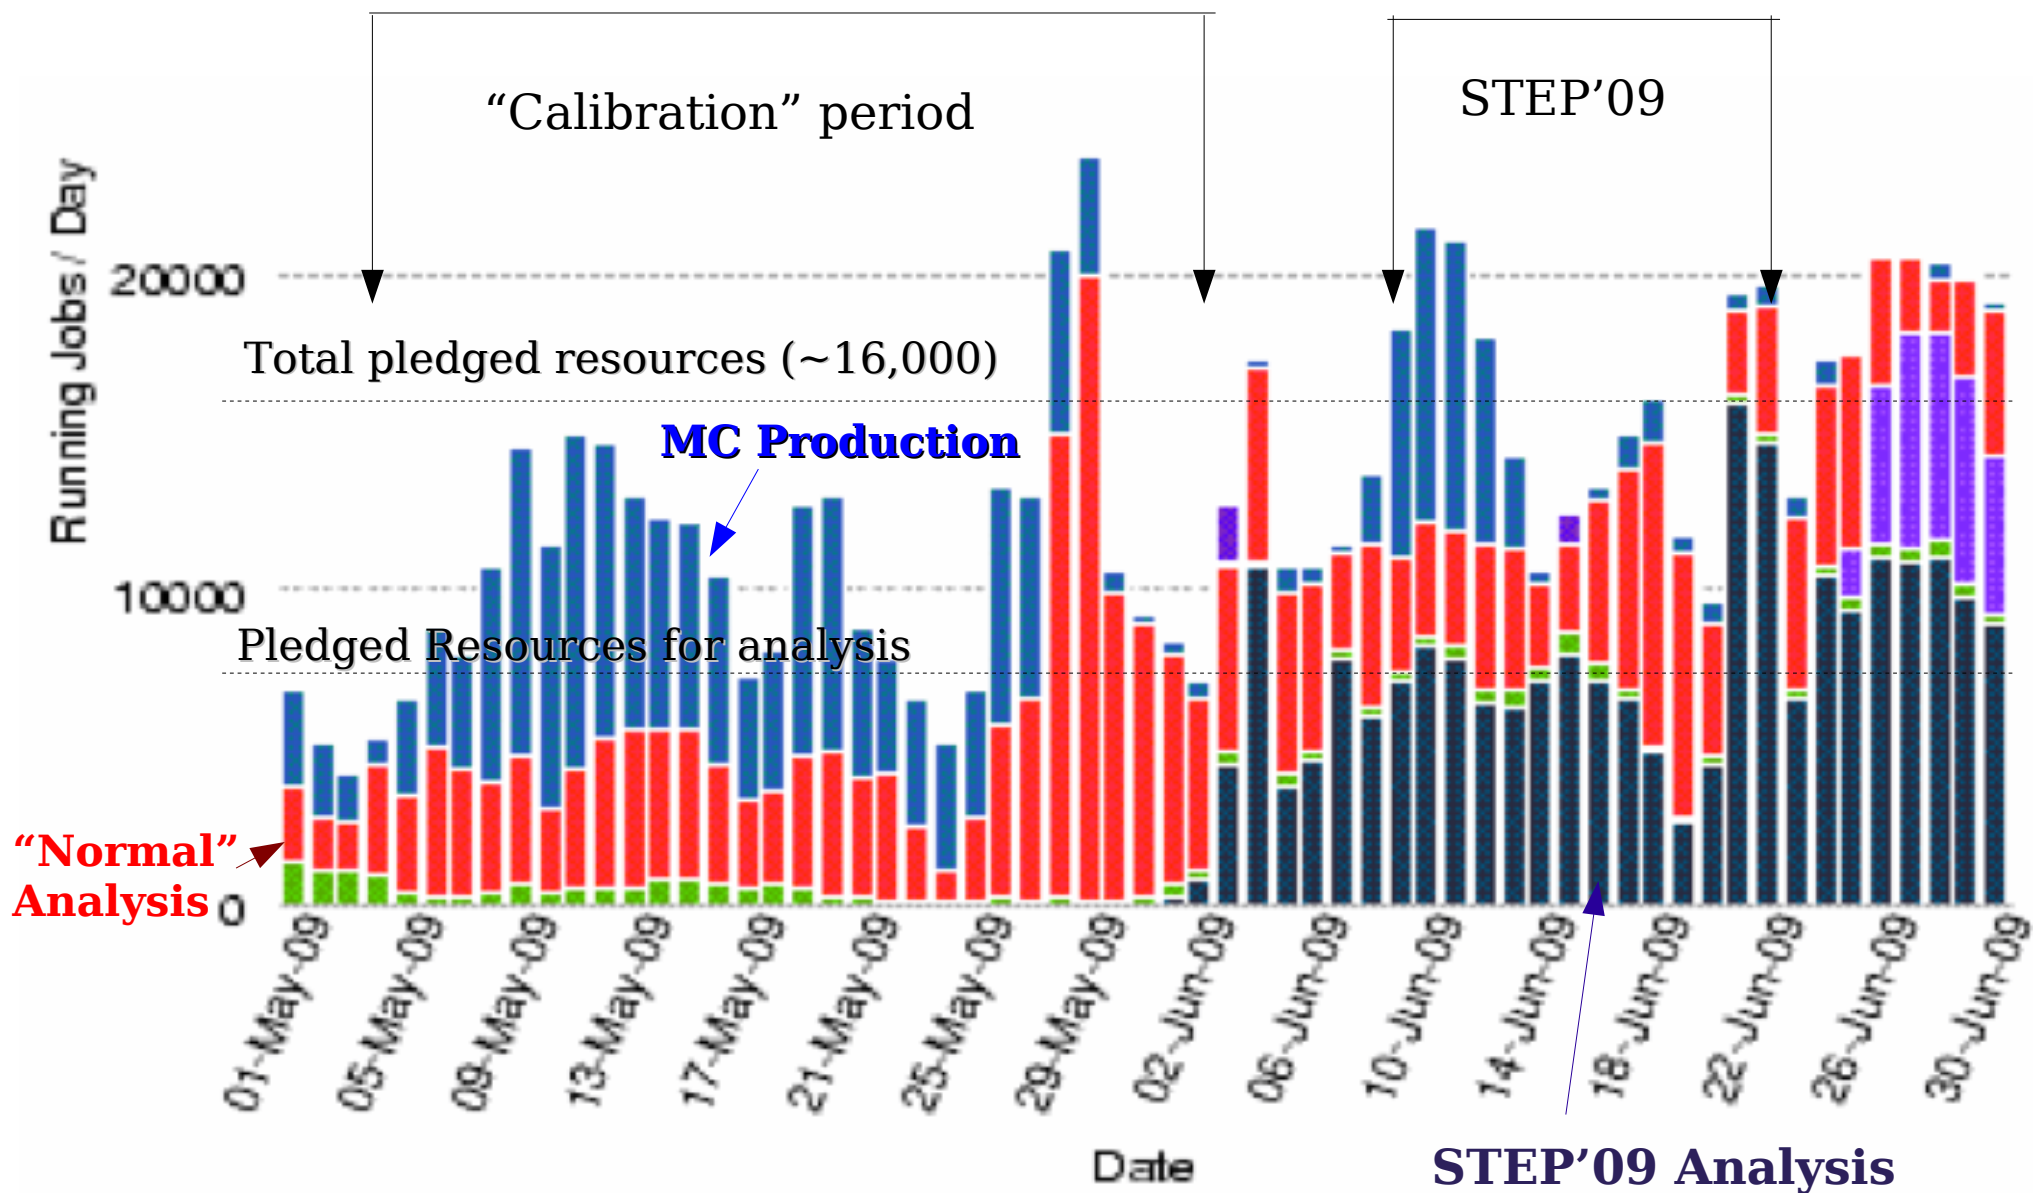
of jobs running vs time

% of analysis pledge used in the two periods.

% job failures per site

<https://twiki.cern.ch/twiki/bin/view/CMS/Step09AnalysisMetrics>


Average # of cores used @ T2s




Resource usage from WMS via Crabserver

- For STEP'09 CRABServer was used via both glidein and glite WMS

UCSD-CS with glideinWMS

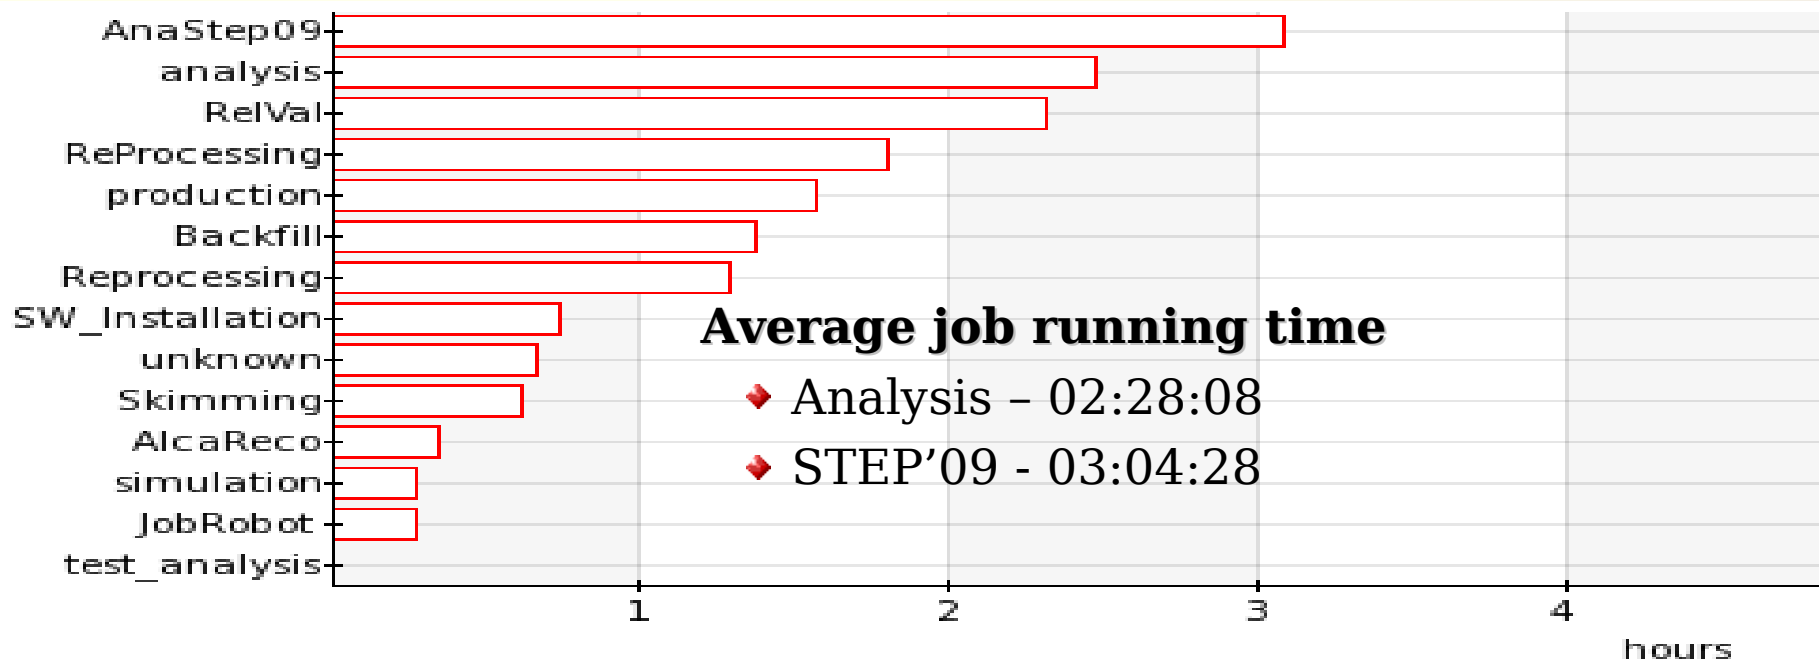


Bari-CS with gliteWMS



- The CRABserver was tested with more than 130,000 submission per day
- Very successful workflow - Running 13,200 cores in parallel
- Submitted to about 50 T2 sites
 - Excellent response from the sites in case of errors (See Next)


CPU/WallClock time usage



- We studied logs from each and every job from glideinWMS as well as statistics from dashboard for gliteWMS
- On average we ran on more than **12,000** cores in parallel per day
- For example:
 - Average running of ~10,000 CPUs in parallel per day for 24 hours
 - Jobs with 3 hour duration
 - This corresponds to $10000 * 24 / 3.0 = 80,000$ CMS jobs per day

[Personally: I do not prefer # number of jobs as the metric]

Resource usage @ CMS T3 sites



- Several T3 sites also participated during the exercise
- Largest and newest among them - T3_US_Omaha
 - Thanks to the site admins for getting ready during very short time

Can we easily use all sites at their pledged levels of analysis CPU ?



- ◆ Analysis CPU = 50% of total pledged CPU

If we increase the load on many sites, does this negatively effect the overall success rate ?

- ◆ We study all logs of jobs we submit from one of our two crabserver.
- ◆ We categorize the errors, and communicate all errors daily to all sites in sufficient detail for sites to fix things.

% of analysis pledge used

Ran on 49 T2s and 8 T3s.



Capable of filling majority of sites at their pledges, or above.

We used in aggregate 1.5 times the analysis pledge.

Caveats:

- Several sites had at least one day downtime during STEP09.
- We did not queue jobs at all sites all the time.
- See backup slides for queued and running for all sites.

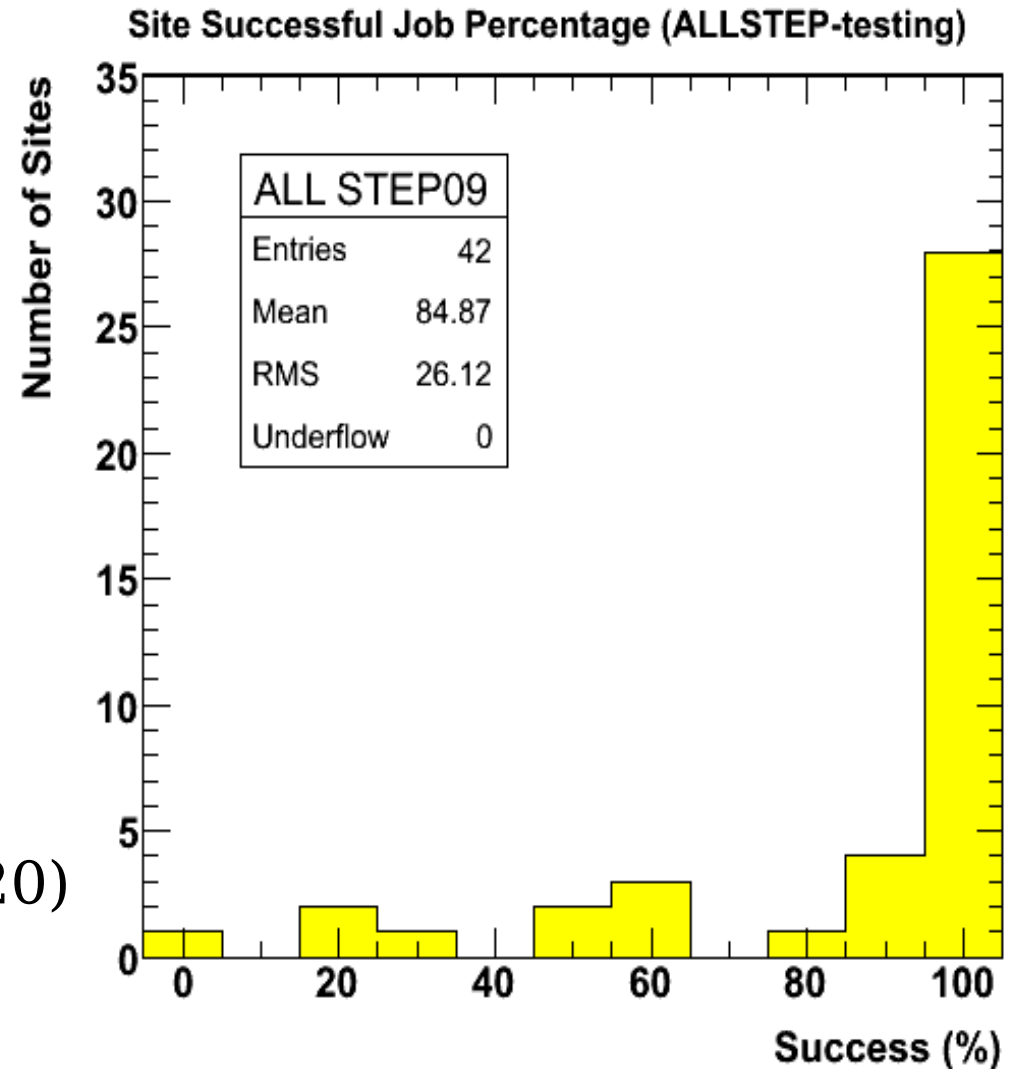
Error statistics for STEP'09

Standard analysis jobs were used:

- CMSSW reading data from SE
- No stageOut

Roughly 80% success rate

Approximately 90% of the failures
are due to read failures (error 8020)



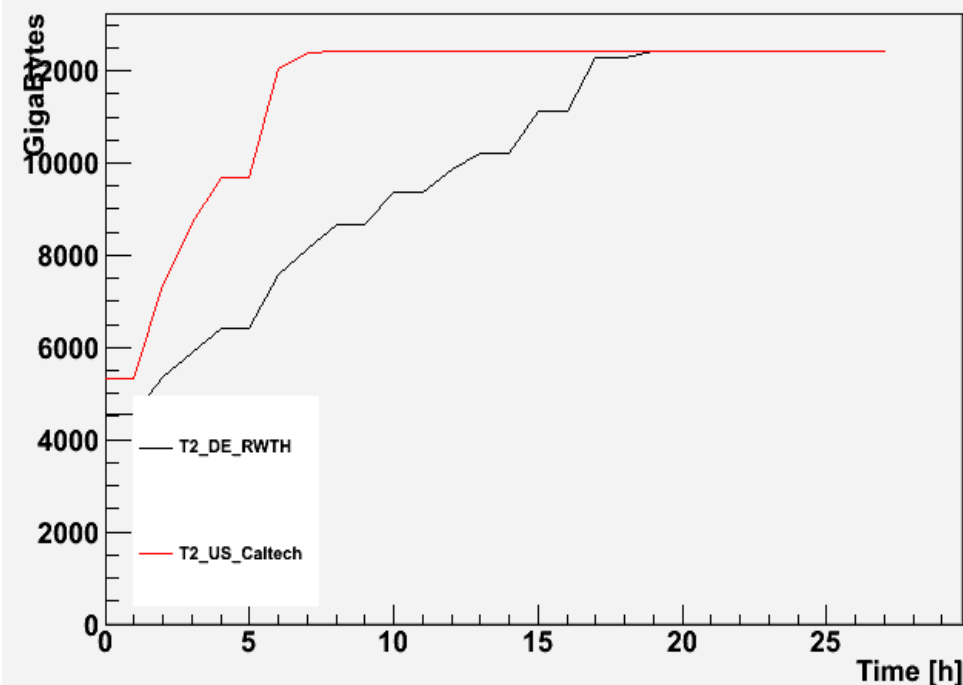
Most sites were very successful !!!

Attempt to more effectively use - the global T2 system


- **Pick 5 “hot” datasets.**
- **Distribute it to 10 sites**
 - Can we successfully move the data ?
 - Does this increase the use of the sites ?

Experience moving data

T-1 -> T2 analysis transfers



T-1 -> T2 analysis transfers



Ability to move data is somewhat mixed

Roughly 3 categories of sites:

Easy: 0.4-1 TB/hour with ease

Fine: ~100 GB/hour with some stops and starts

Difficult: week(s) to transfer a few TB (Not shown)

Note: Not all sites had to transfer same amount of the data !

Data utilization, once moved



Moved /InclusivePPmuX/Summer08_IDEAL_V11_redigi_v1/GEN-SIM-RECO to JINR and CSCS
Find that use of that dataset picked up after it was moved.

Note:

- Beijing, JINR, CSCS are roughly same size.
- Beijing had the dataset already before STEP09.
- This indicates that not all users use white lists!

Summary and Conclusions

STEP'09 exercise was needed to evaluate the readiness for data taking.

Exercised the T2/T3 system at the level expected for early data taking.

- actually 1.5 times the analysis pledge

The job submission infrastructure is adequate for the initial collisions.

The storage infrastructure will need some improvement.

- 90% of all job failures are due to read failure.

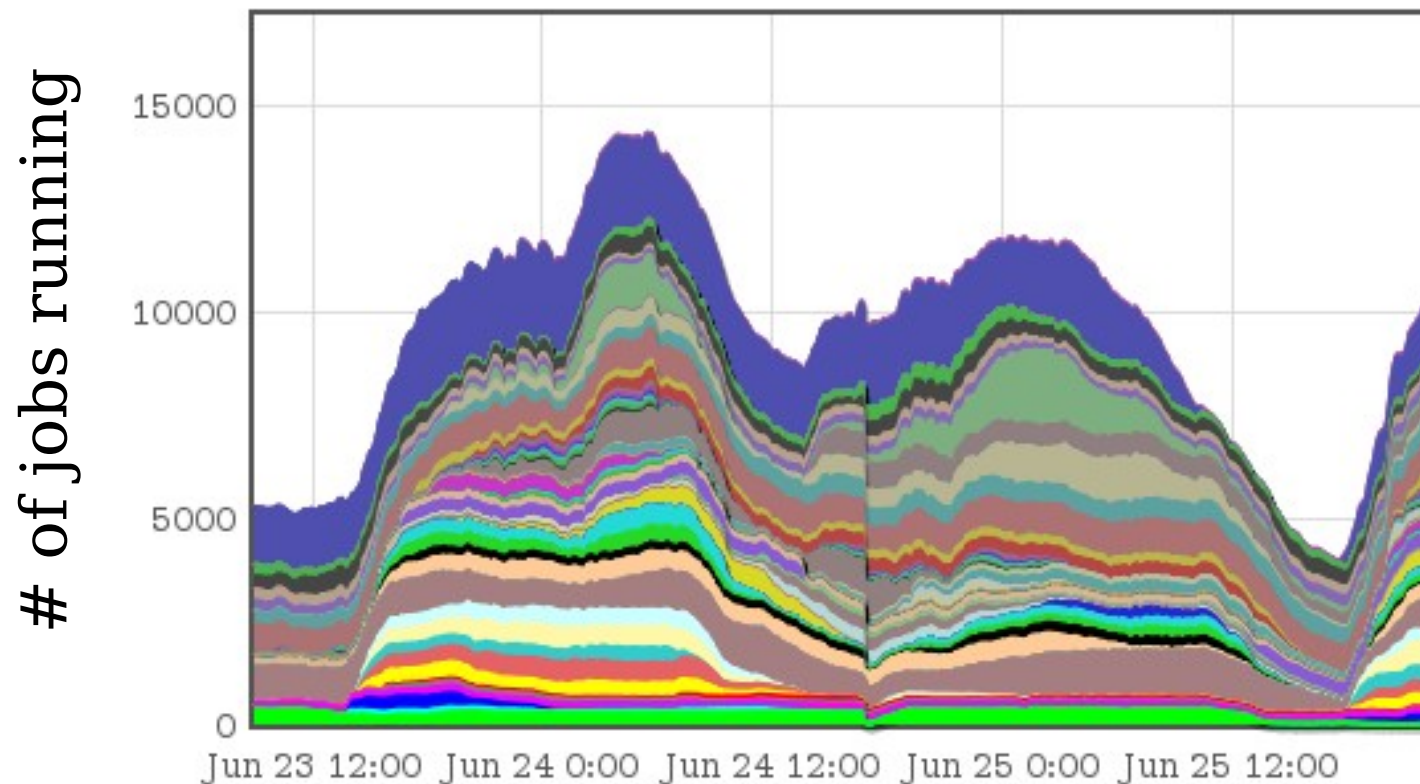
If the read failures could be reduced by a factor 5

- we should be able to operate at >95% success rates.

Thanks to the sites-admins who made this exercise a real success !!!

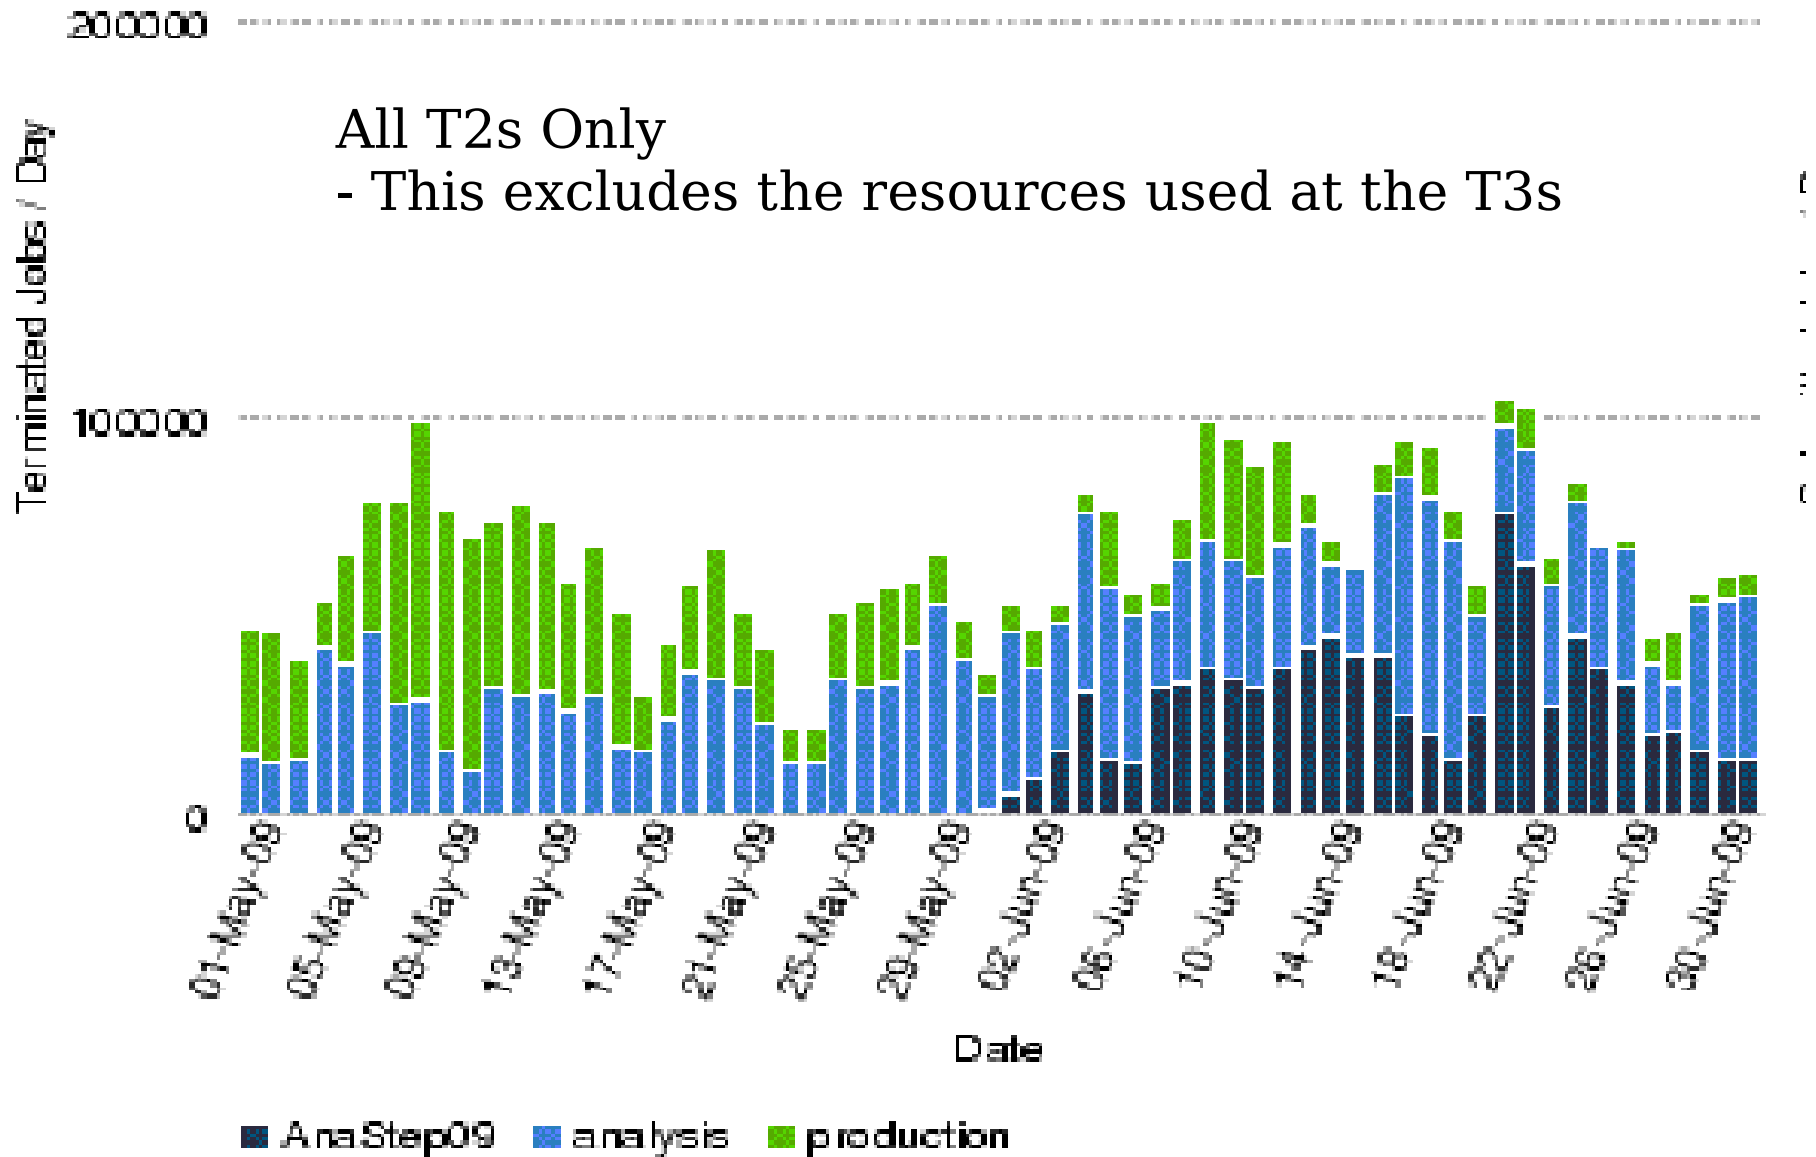
Backup slides

Total running for glidein based crabserver




This shows that we can fill the entire analysis pledge Worldwide from a single submission point.

of jobs finished - as per Dashboard



WallClock time used - as per Dashboard

Distribution of WallClock Time utilised by analysis and anastep09 jobs at the T2



Maximum: 1,338,974,952 , Minimum: 0.00 , Average: 695,236,611 , Current: 402,765,282