

# On monitoring and accounting

Antonio Pérez-Calero Yzquierdo

2016-07-14

# Monitoring and accounting

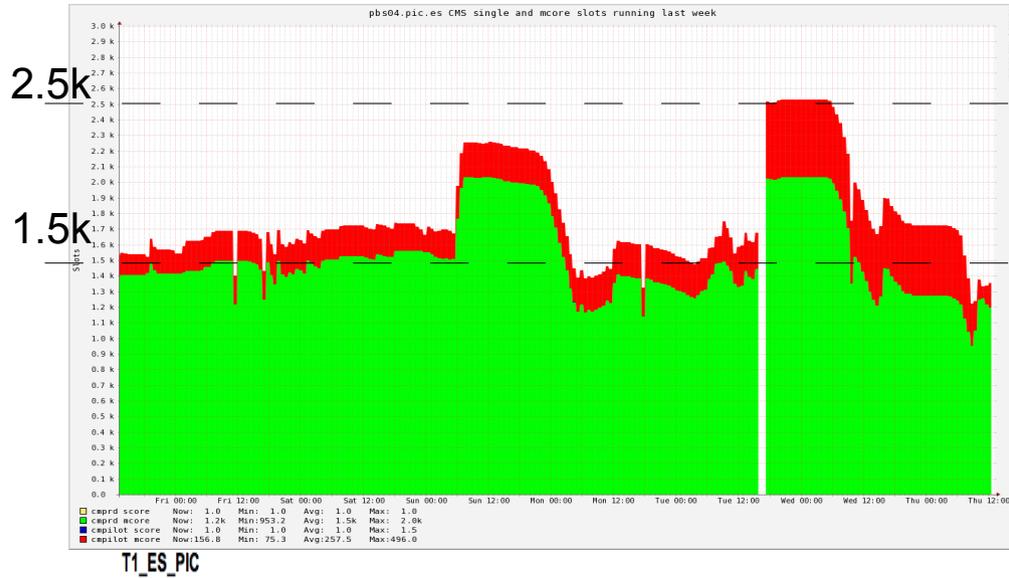
- Not a sexy topic but needs to be addressed!
- I am taking part on the WLCG accounting task force, as site/CMS contact
  - Its mandate is to review how LHC VOs and sites report and account for resource utilization
  - Specially CPU utilization!
- First point of disagreement: VOs tend to look at **payloads**, while sites see the whole **pilot**
- Not a huge difference if the pilot model is trivial or extremely efficient
  - ~OK for T2s up to Spring (single core pilots)
  - Not so good for T1s, running mcore pilots already in 2015
- With the move of CMS SI towards the "multicore global pool", this difference now needs to be understood and accounted for
  - Ultimately, we should be able to combine our own payload and pilot accounting information and reproduce site view (to a reasonable level of accuracy)
- As SI team, **the responsibility for the pilot part is on us**

# The situation right now

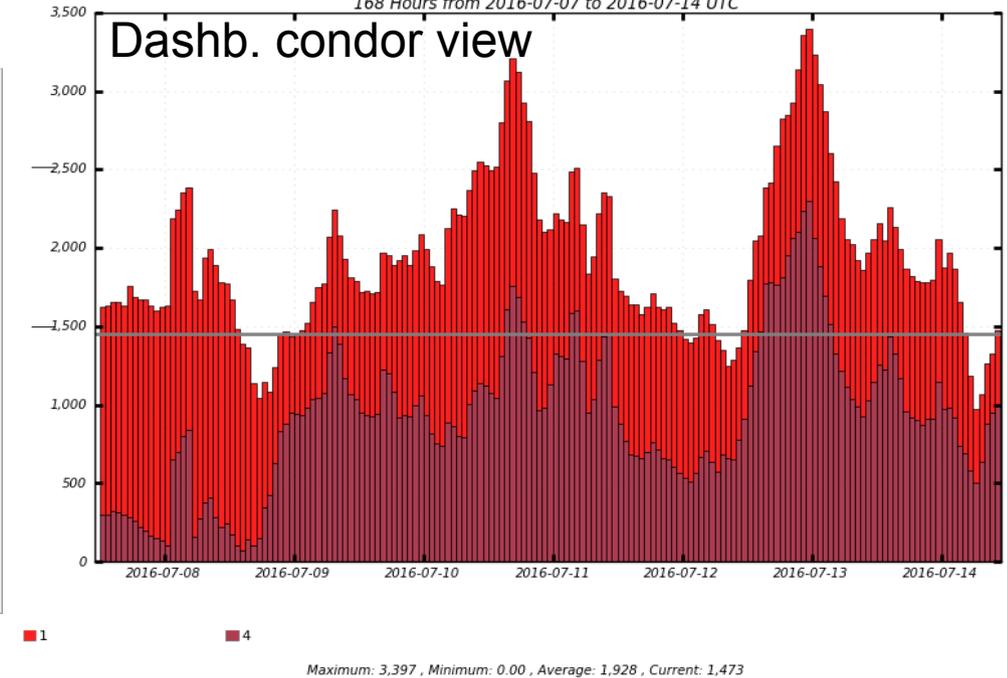
- Payload:
  - Dashboard historical view
  - Dashboard condor view
  - gwmsmon
- Pilots:
  - GlideinWMS FE and factory monitors
  - ad-hoc built monitors
  - gmwsmon
- Sites: their own accounting and monitoring
- Things to avoid:
  - Fragmented information
  - Not reliable information (!!)

# Compare PIC last week from 4 different views

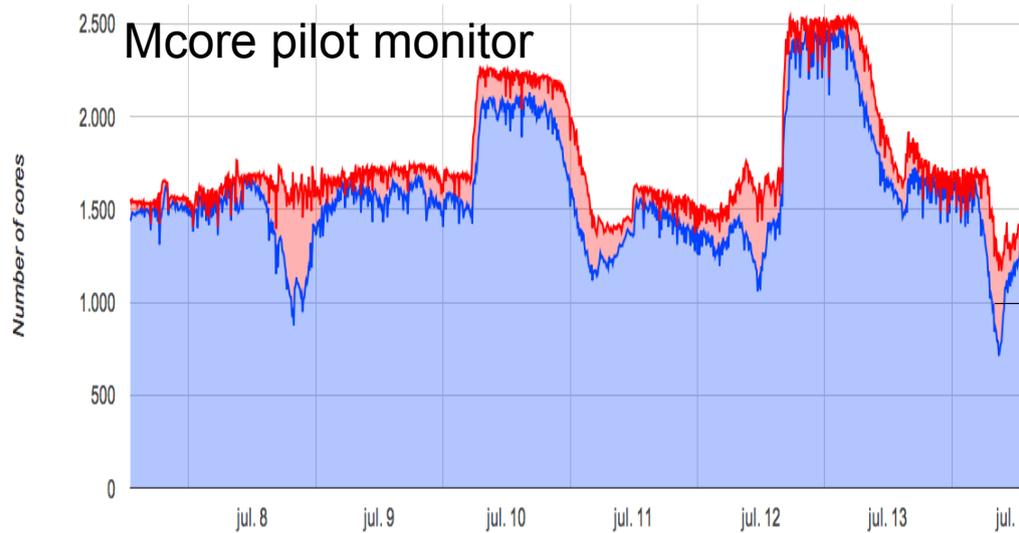
## PIC internal monitor



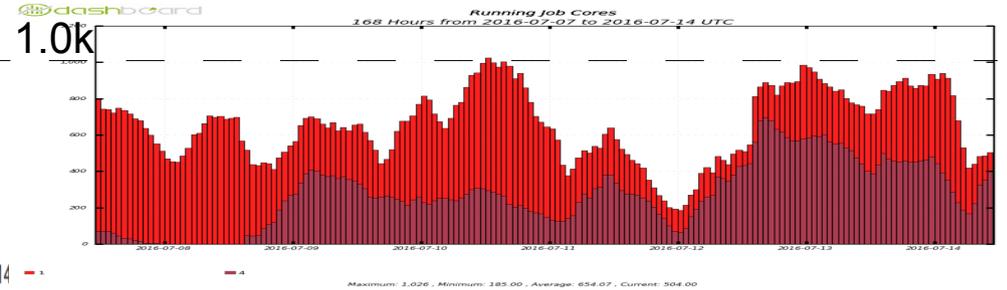
## Dashb. condor view



## Mcore pilot monitor



## Dashb. historical view



# Observations

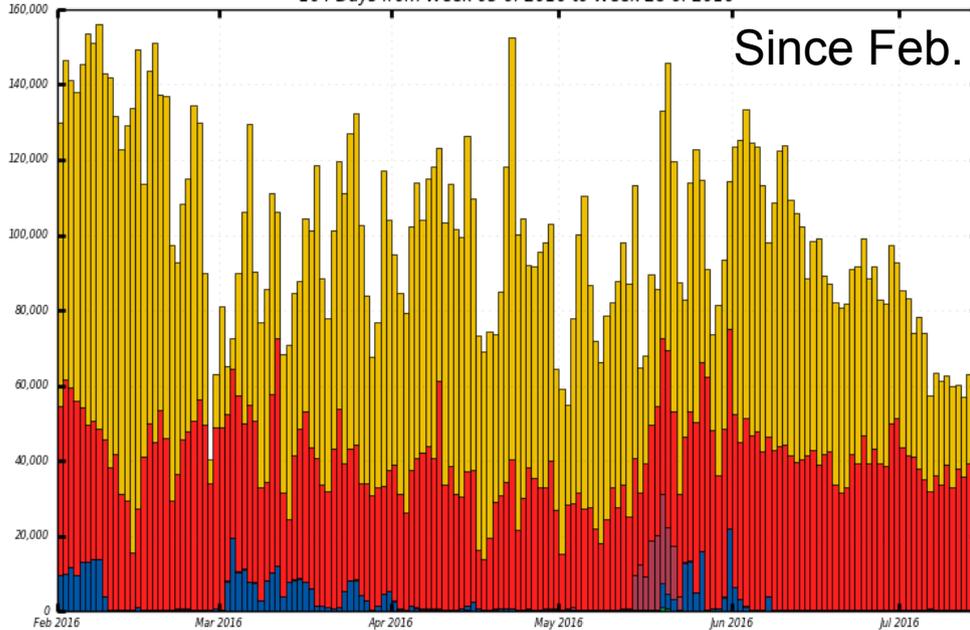
- At the level of resource allocation, in the example for PIC, its internal monitoring view and CMS multicore pilot monitoring match
- Being used quite efficiently as it is a T1 site which runs basically production jobs
  - Cores wastage in pilot ~5%
- We can use pilot core utilization to estimate number of running payloads:
  - It does not match the dashboard historical view nor the condor view
  - Both need(ed) to be verified

# Dashboard views problems



Running Job Cores  
164 Days from Week 05 of 2016 to Week 28 of 2016

Since Feb.

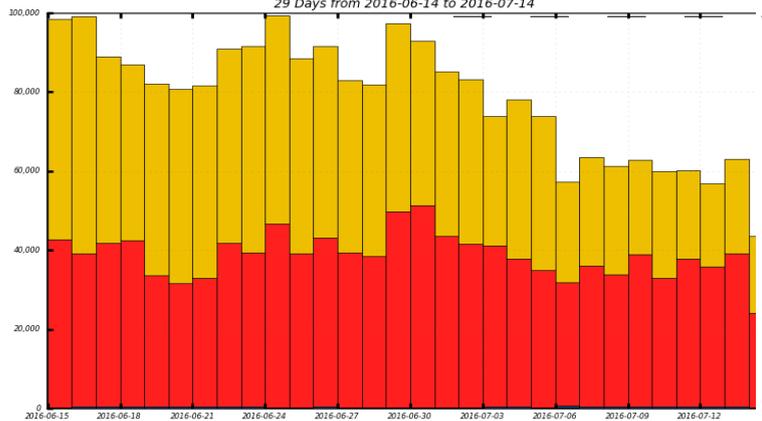


production analysis integration test unknown

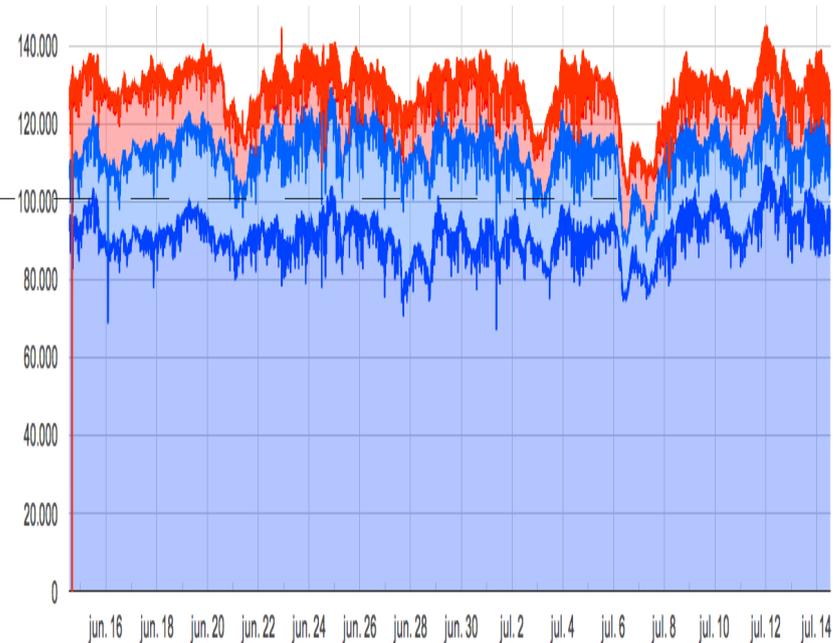
Maximum: 155,949 , Minimum: 40,398 , Average: 99,575 , Current: 43,500

Last month

29 Days from 2016-06-14 to 2016-07-14



- Historical view: bug in WMAgent monitoring reporting module, analysis jobs are ok.
- From Alan: bug introduced on June 10<sup>th</sup> cmsweb upgrade: wfs injected from that moment not reporting from WN to dashb.
- Corrected in the cmsweb upgrade yesterday, new wfs should be reporting again correctly
- **Unnoticed for quite a while!!**



# Dashboard views problems

- **Condor view:**
  - Traditionally used as a cross check to historical view:
    - stopped to be used as numbers did not make sense since the introduction of multicore jobs
  - Not including analysis!
    - code written before the adoption of the global pool
  - Not maintained! Orphan code
- So (I'd say) we need to
  - Correct numbers: make it reliable again
  - Add analysis: should be easy, just make it check and report info from crab schedds
  - Take **responsibility of the reporting code into SI**
    - move code to our gitlab area
- See discussions at <https://github.com/amaltaro/scripts/issues/>

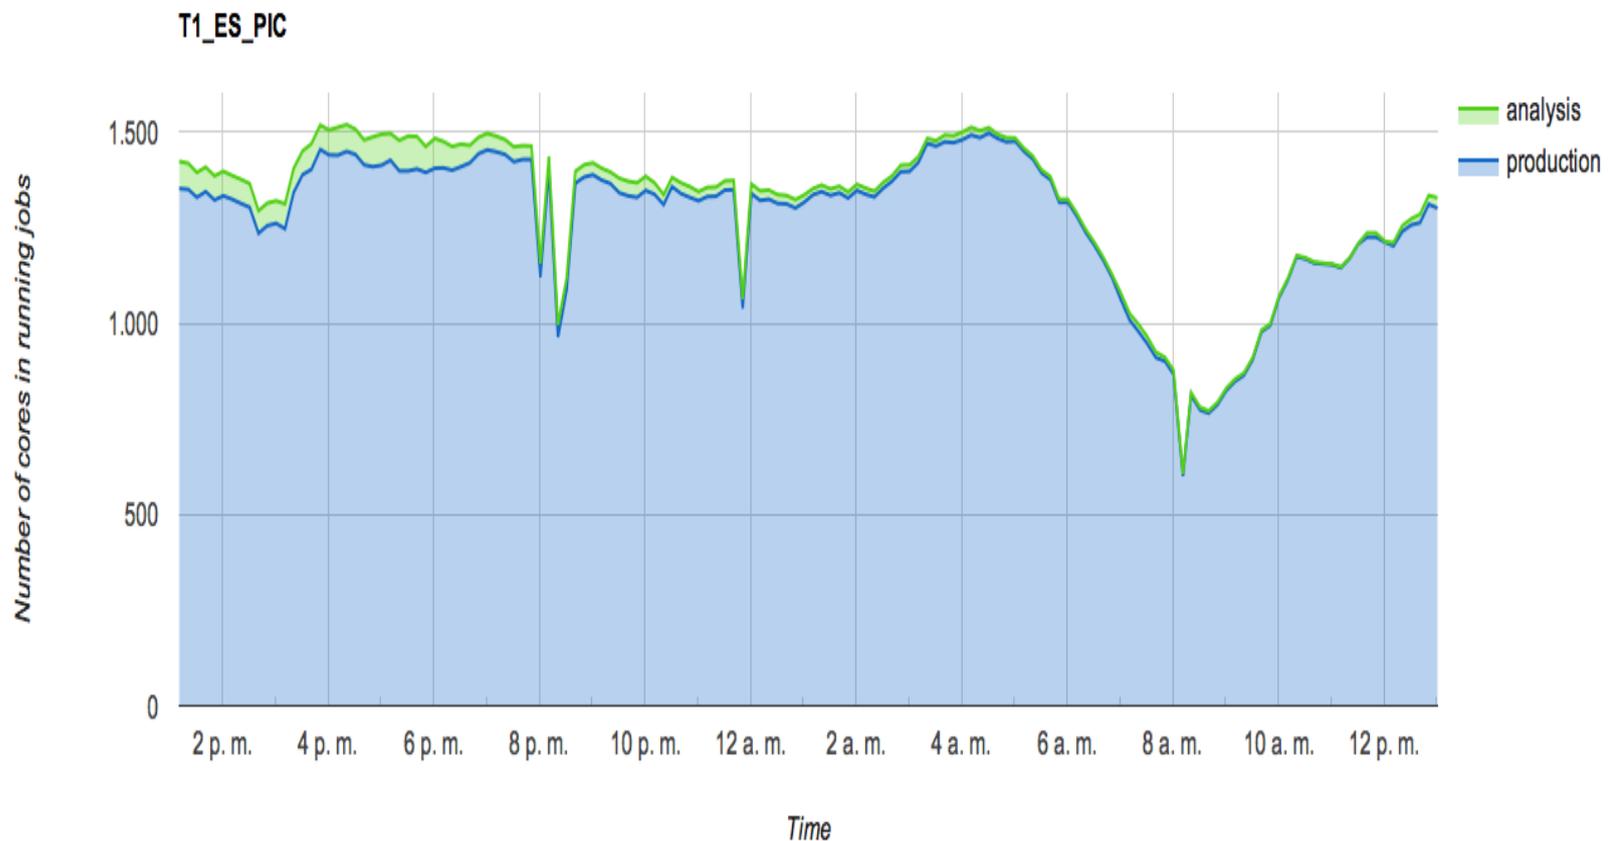
# Some new per site job plots

- Jobs for each site from queries to the global pool CM:

[http://submit-3.t2.ucsd.edu/CSstoragePath/aperez/HTML/JobInfo/jobstatus\\_T1\\_24h.html](http://submit-3.t2.ucsd.edu/CSstoragePath/aperez/HTML/JobInfo/jobstatus_T1_24h.html)

[http://submit-3.t2.ucsd.edu/CSstoragePath/aperez/HTML/JobInfo/jobstatus\\_T1\\_24h.html](http://submit-3.t2.ucsd.edu/CSstoragePath/aperez/HTML/JobInfo/jobstatus_T1_24h.html)

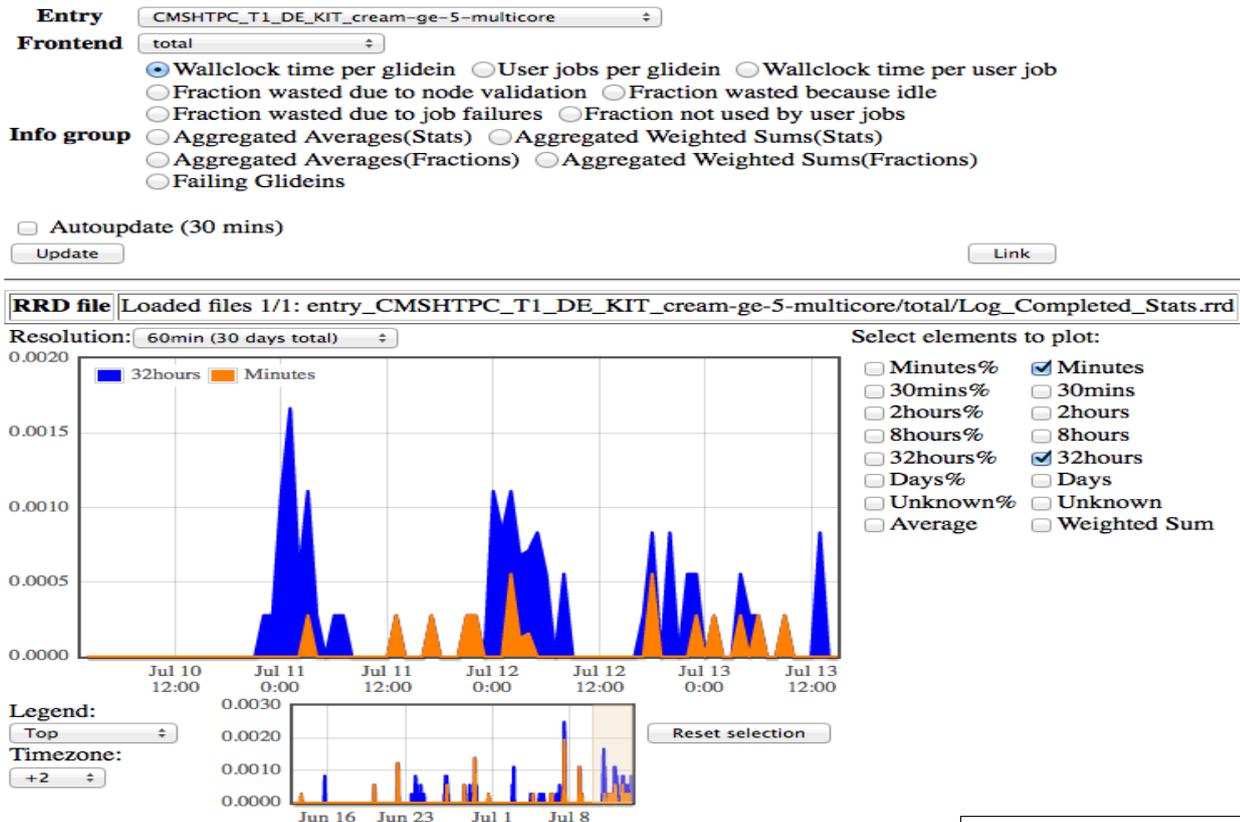
- A light-weight additional cross check, see PIC in the last 24h



# Pilot accounting

- Still missing systematic monitor & accounting of per-pilot information
- Factory monitoring from pilot jobs: fragmented to need to access and aggregate pilot logs
  - To dashboard?

## Glidein factory status from logs - v3\_2@CMS-CERN



**401 Authorization Required**

- Logs collected by Brian at Kibana/ElasticSearch node at hcc-metrics.unl.edu (?)