CMS: How scrutiny reports are prepared

WLCG Accounting TF – CERN – 23/06/2016

J. Flix (PIC/CIEMAT) [remote]

Document structure

- A single document with the Resource Utilization (typically covering the last year) and with the Resource Requests (typically, corrections for next year and indications for current+2 year)
- The Resource Utilization covers:
  - HLT farm
  - Tier-0 and CAF
  - Tier-1 and Tier-2 resources
  - Production and analysis activities
  - Data popularity and disk access
  - Dynamic Data Management
  - Data Transfers

Document structure

- A single document with the Resource Utilization (typically covering the last year) and with the Resource Requests (typically, corrections for next year and indications for current+2 year)
- The Resource Utilization covers:
  - HLT farm **[dashboard]**
  - Tier-0 and CAF **[dashboard]**
  - Tier-1 and Tier-2 resources **[WLCG repository]** **[EGI Accounting Portal]** **[dashboard]**
  - Production and analysis activities **[dashboard]**
  - Data popularity and disk access **[CMS internal monitoring]**
  - Dynamic Data Management **[CMS internal monitoring]**
  - Data Transfers **[PhEDEx]**

Document structure

- A single document with the Resource Utilization (typically covering the last year) and with the Resource Requests (typically, corrections for next year and indications for current+2 year)
- The Resource Utilization covers:
  - HLT farm **[dashboard]**
  - Tier-0 and CAF **[dashboard]**
  - Tier-1 and Tier-2 resources **[WLCG repository]** **[EGI Accounting Portal]** **[dashboard]**
  - Production and analysis activities **[dashboard]**
  - Data popularity and disk access **[CMS internal monitoring]**
  - Dynamic Data Management **[CMS internal monitoring]**
  - Data Transfers **[PhEDEx]**

Tier-1 Resources

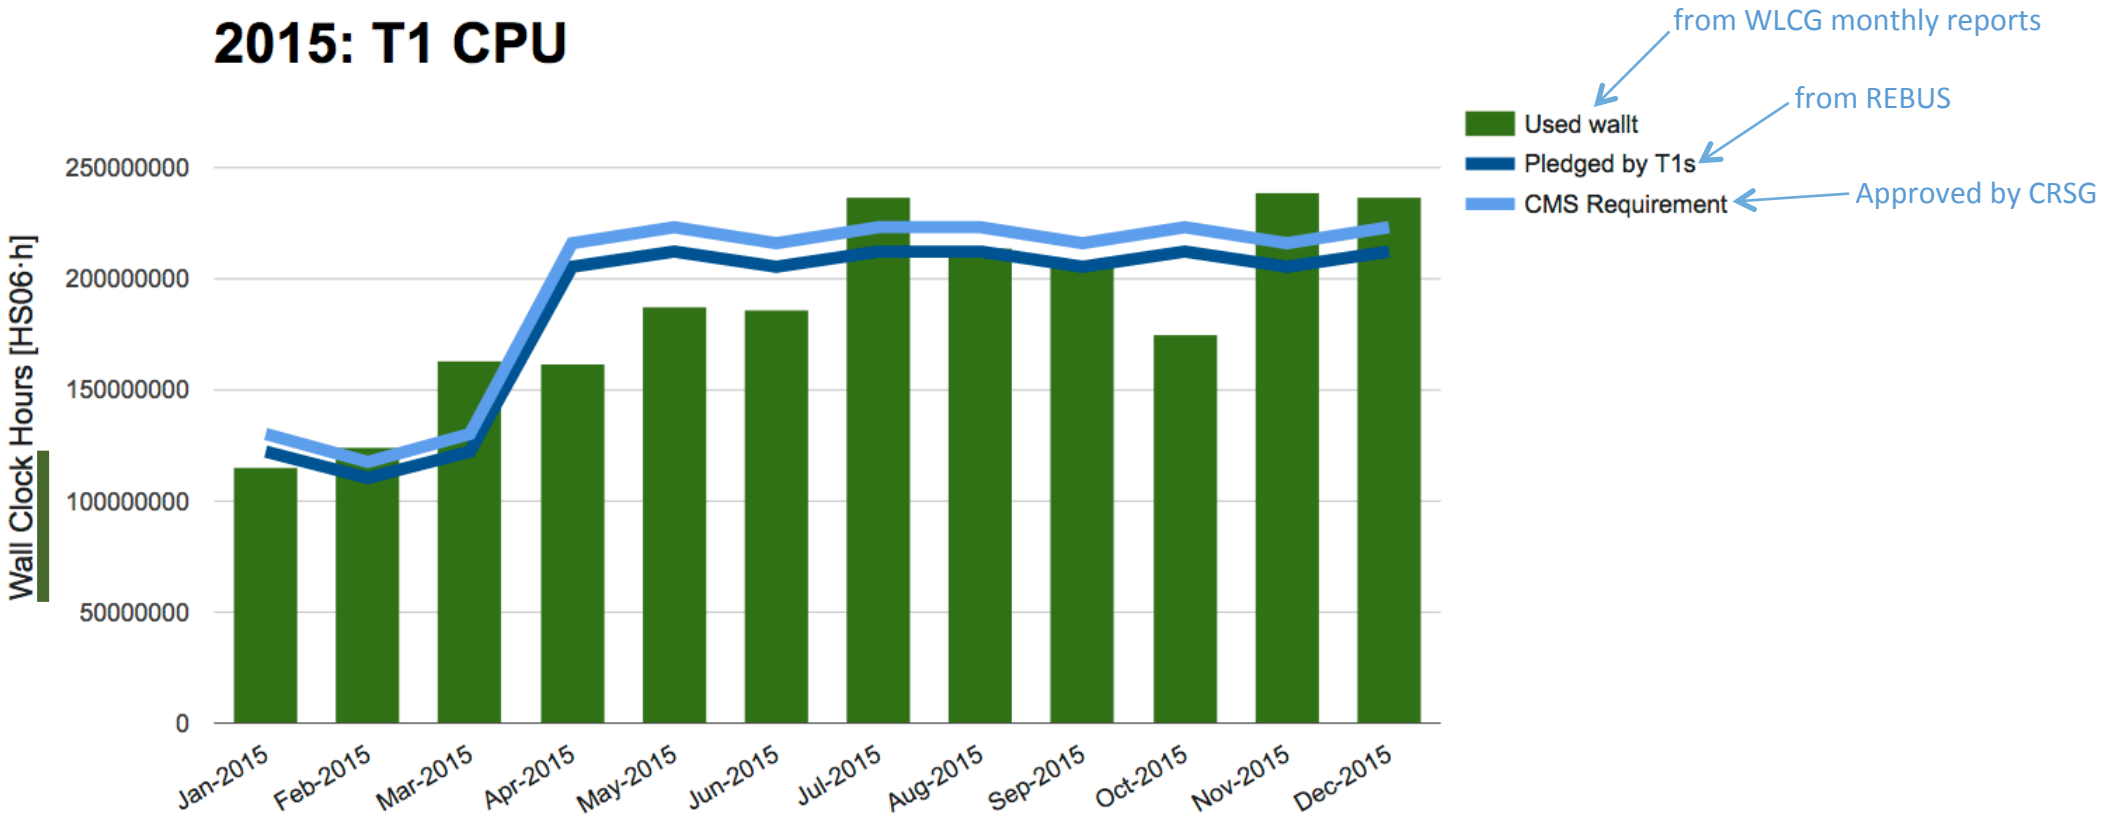

- The accounting information from the WLCG monthly reports (pre-filled with EGI Accounting Portal numbers)
  - Small corrections applied for some sites and some months, since MultiCore accounting was not correct (it was much better this year!)
  - Since reported CPUTime was correct, corrections for Walltime applied taking into account the yearly averaged CPU_{eff} of the site (for example)
  - We don't know the installed CPU for CMS at the Tier-1s (the report contained the total Tier-1 CPU installed capacity)

Tier-1 Resources

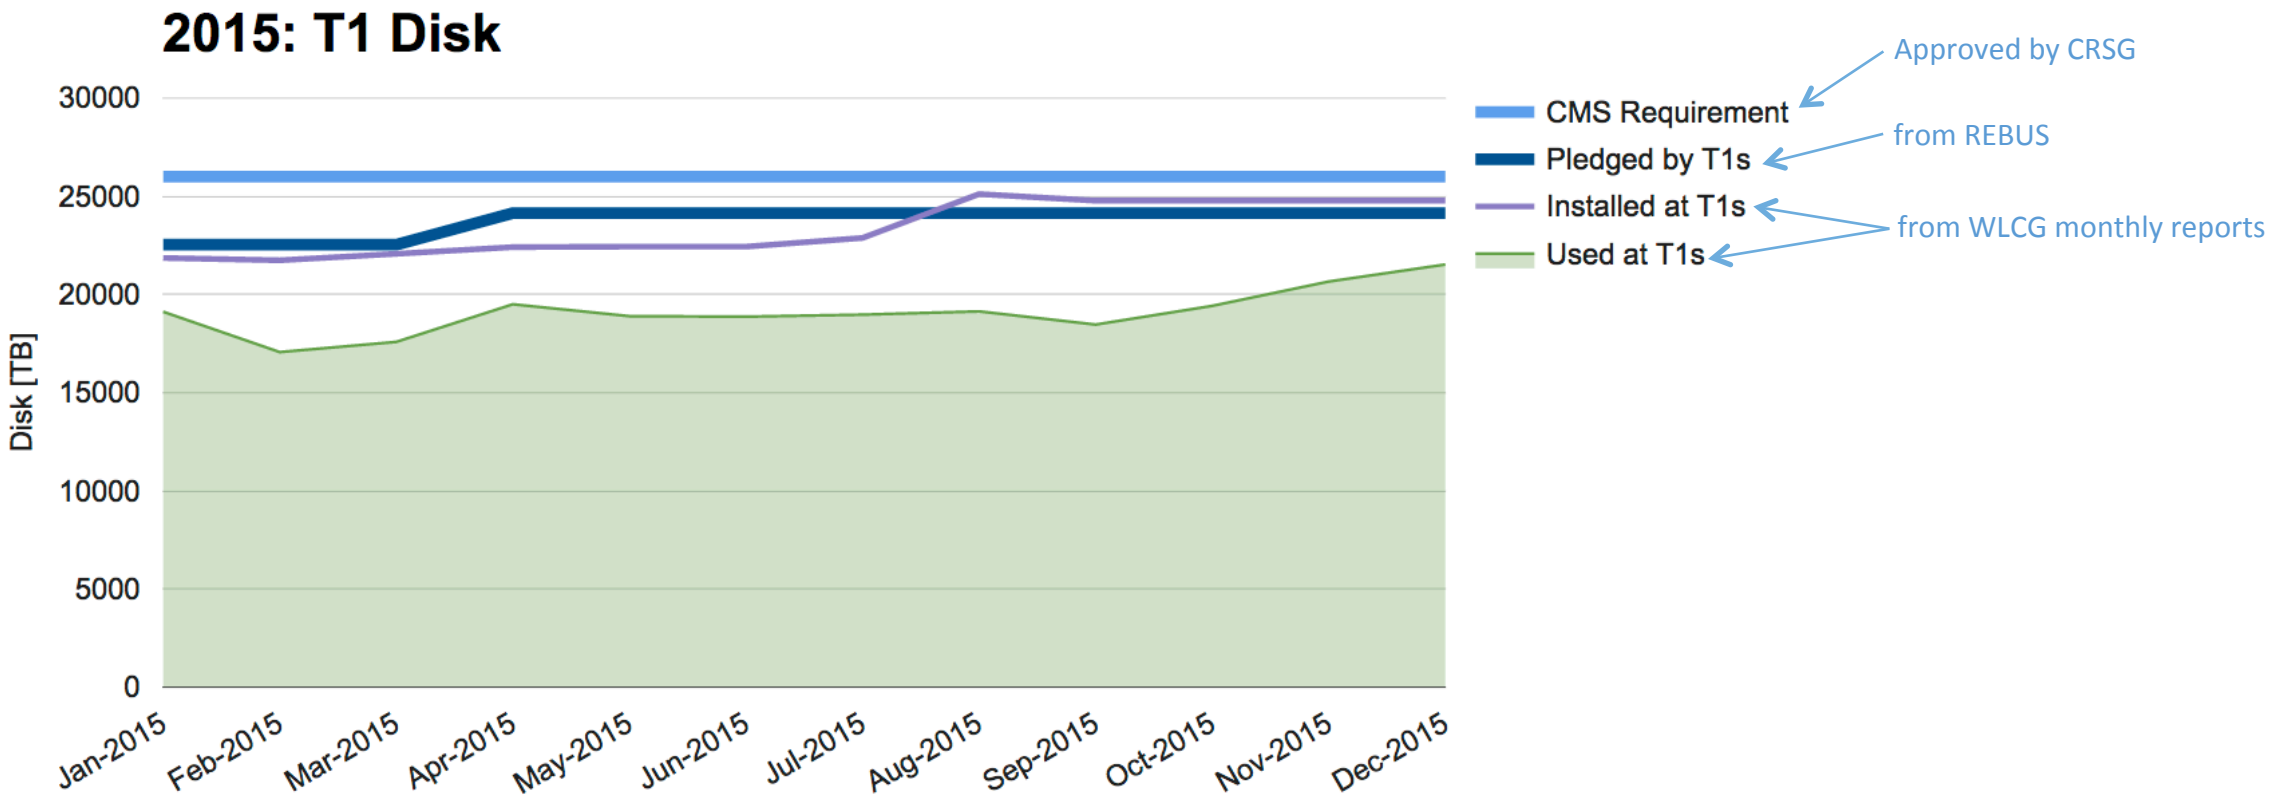

- The accounting information from the WLCG monthly reports
  - The most reliable way to know on installed and used disk at the Tier-1s
  - Sometimes one month is not ok: we report to the site, ask for a correction. If not, extrapolation from previous/next month is applied (not so often)

Tier-1 Resources

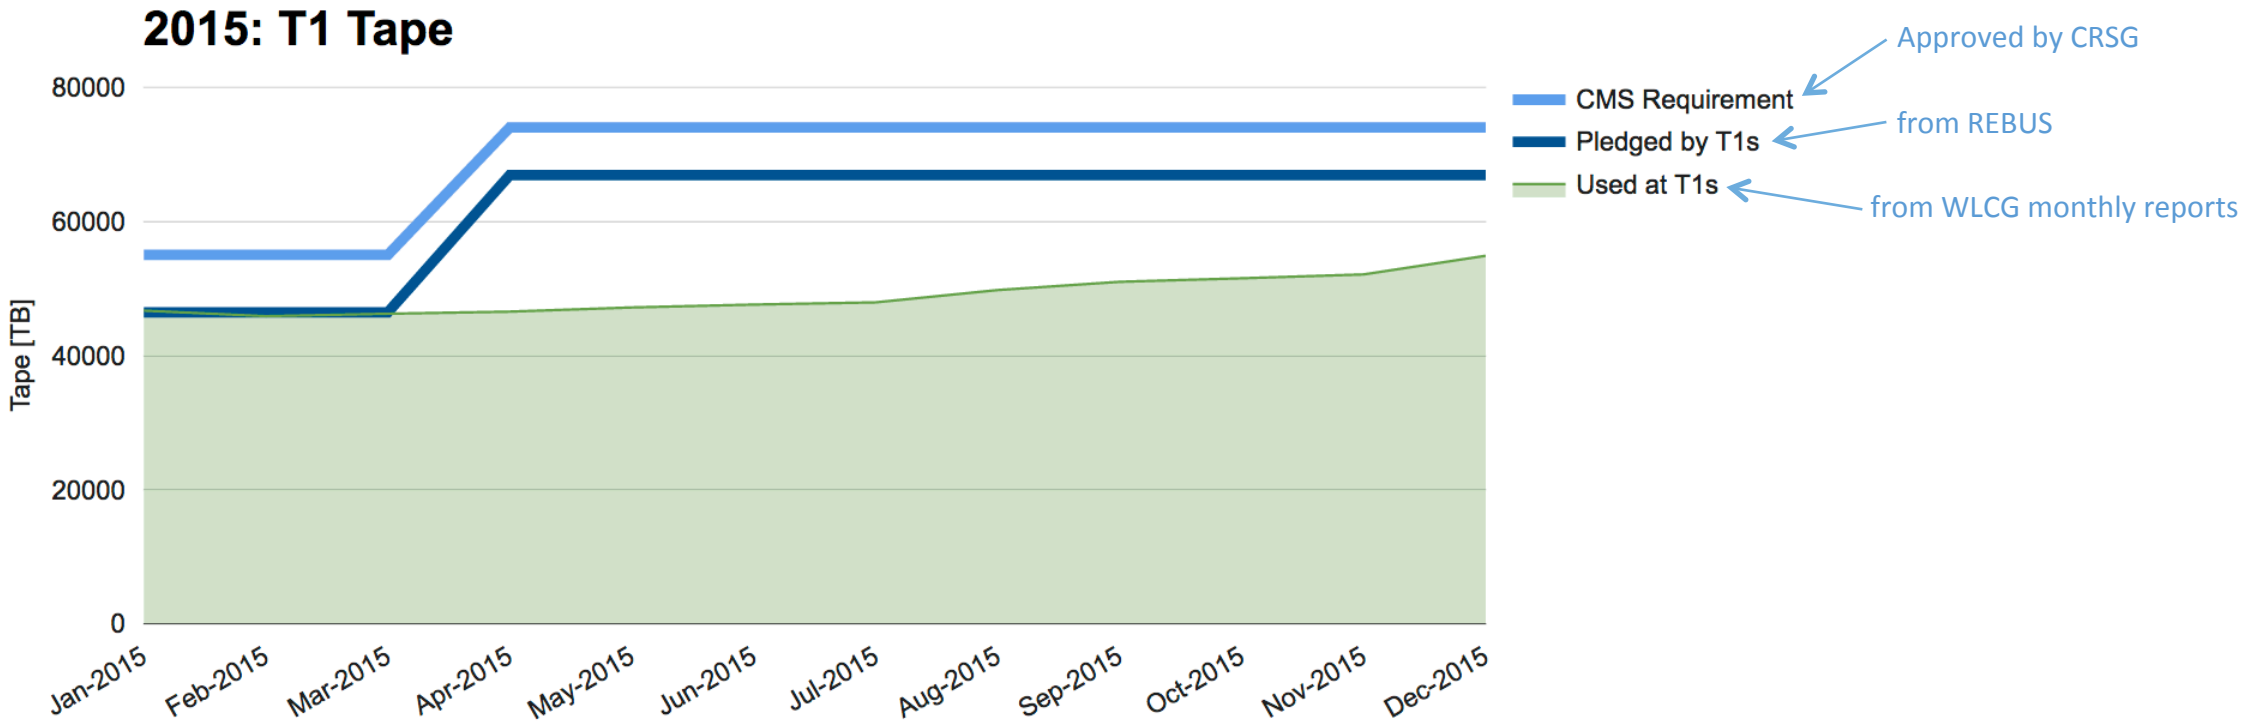

- The accounting information from the WLCG monthly reports
  - The most reliable way to know on used tape at the Tier-1s
  - The reports contain total Tape installed for WLCG, not no breakdown per experiment
  - Sometimes one month is not ok: we report to the site, ask for a correction. If not, extrapolation from previous/next month is applied (not so often)

Tier-1 Resources

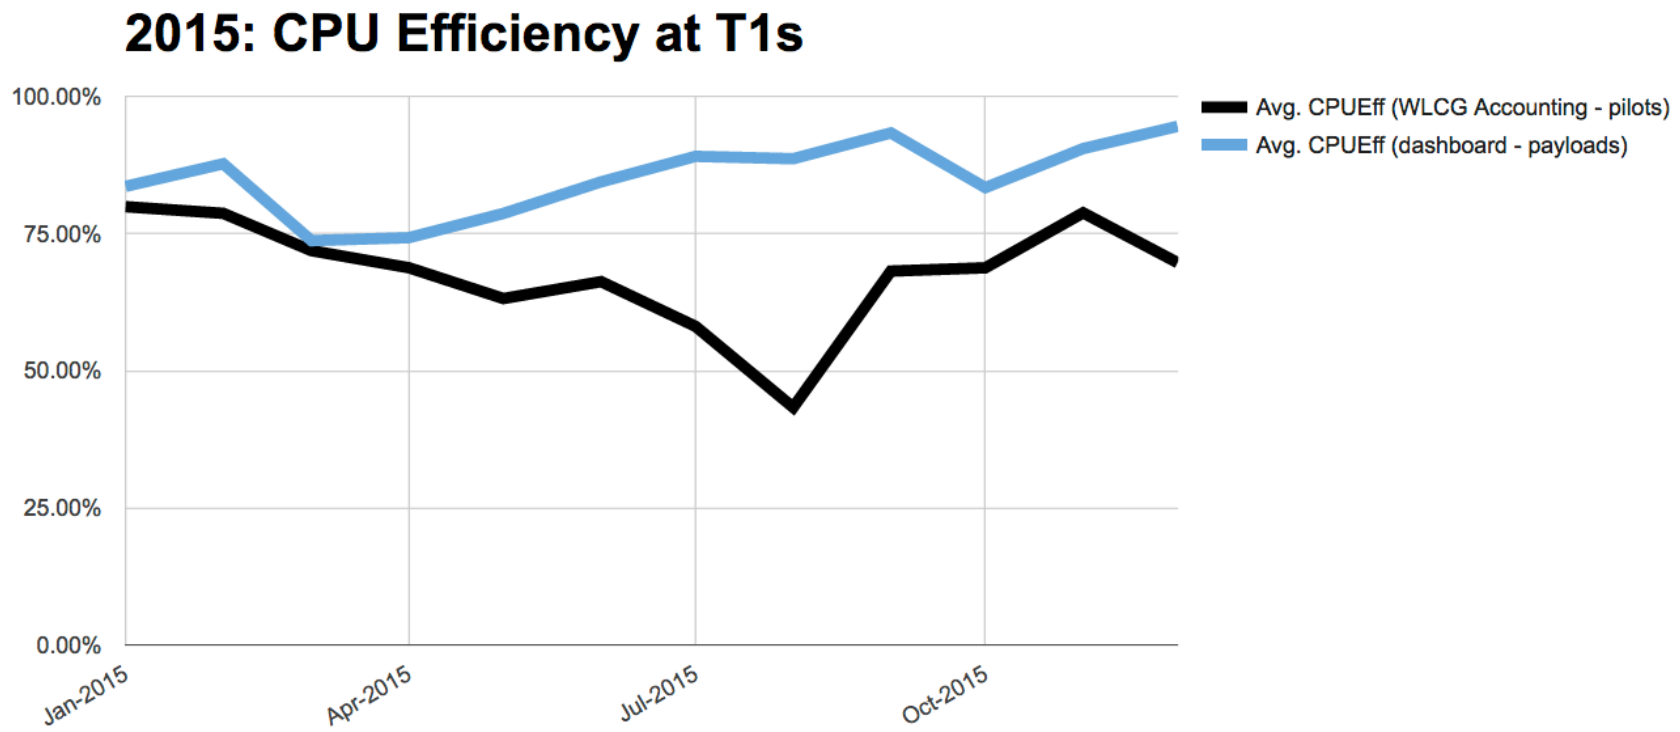

- EGI Accounting Portal hold pilots information – dashboard shows the payload information
- CMS pilots run multiple payloads → Check A. Pérez-Calero talk today for more details
  - The measured pilots average CPU efficiency at the Tier-1 sites in this period was about 70%, dominated by some (known) effects on the multi-core pilots, the transition to this new model, and the running of high-memory activities.
  - The payloads CPU efficiency was about 86%. A number of improvements have been deployed to improve the situation.

Tier-1 Resources

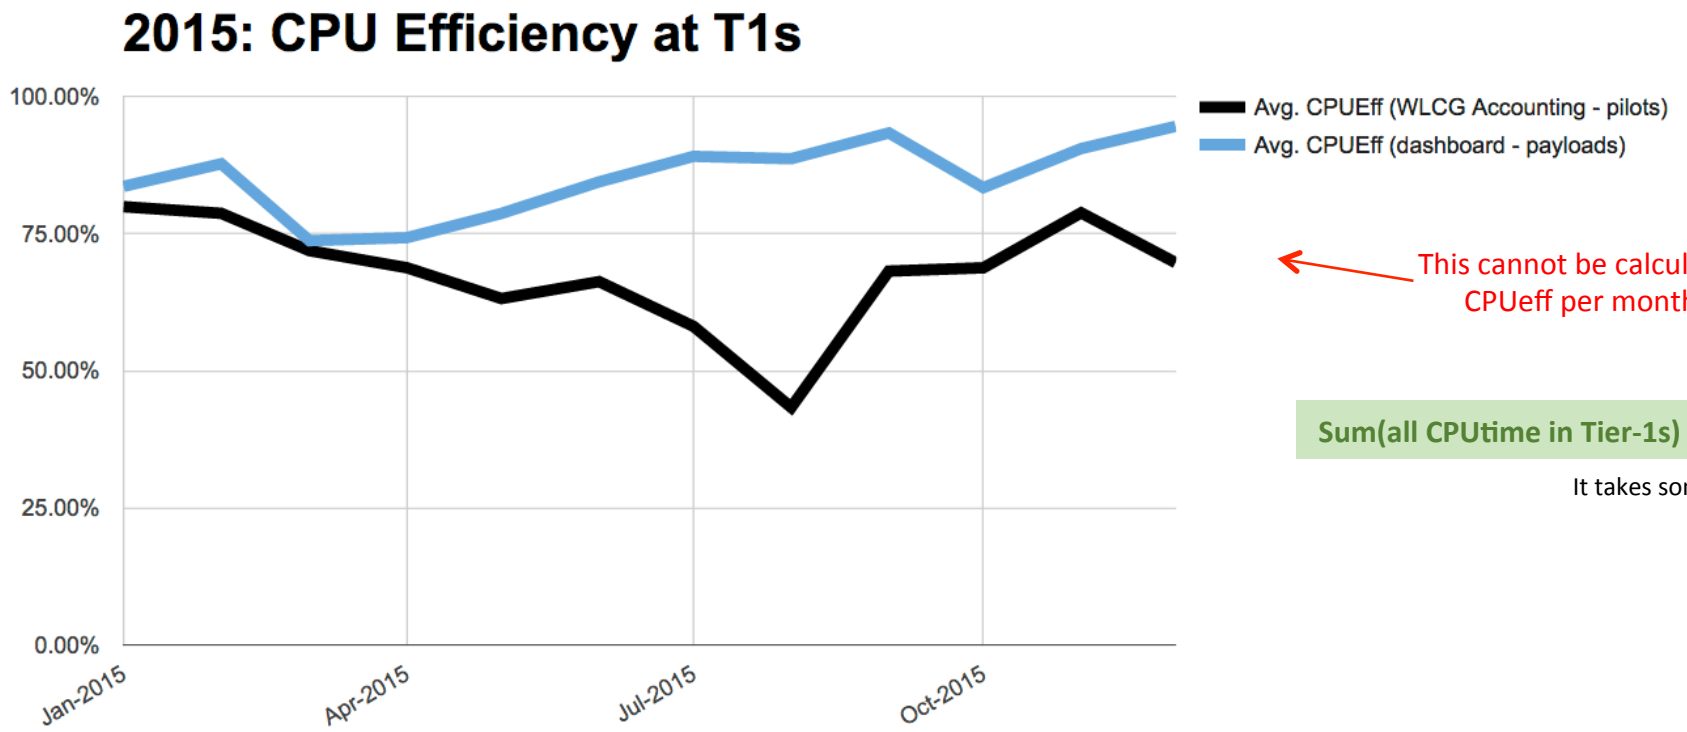

This cannot be calculated as the average of CPUeff per month at the Tier-1 sites

Sum(all CPUtime in Tier-1s) / Sum(all Walltime in Tier-1s)

It takes some effort…

- EGI Accounting Portal hold pilots information – dashboard shows the payload information
- CMS pilots run multiple payloads → Check A. Pérez-Calero talk today for more details
  - The measured pilots average CPU efficiency at the Tier-1 sites in this period was about 70%, dominated by some (known) effects on the multi-core pilots, the transition to this new model, and the running of high-memory activities.
  - The payloads CPU efficiency was about 86%. A number of improvements have been deployed to improve the situation.

Tier-2 resources

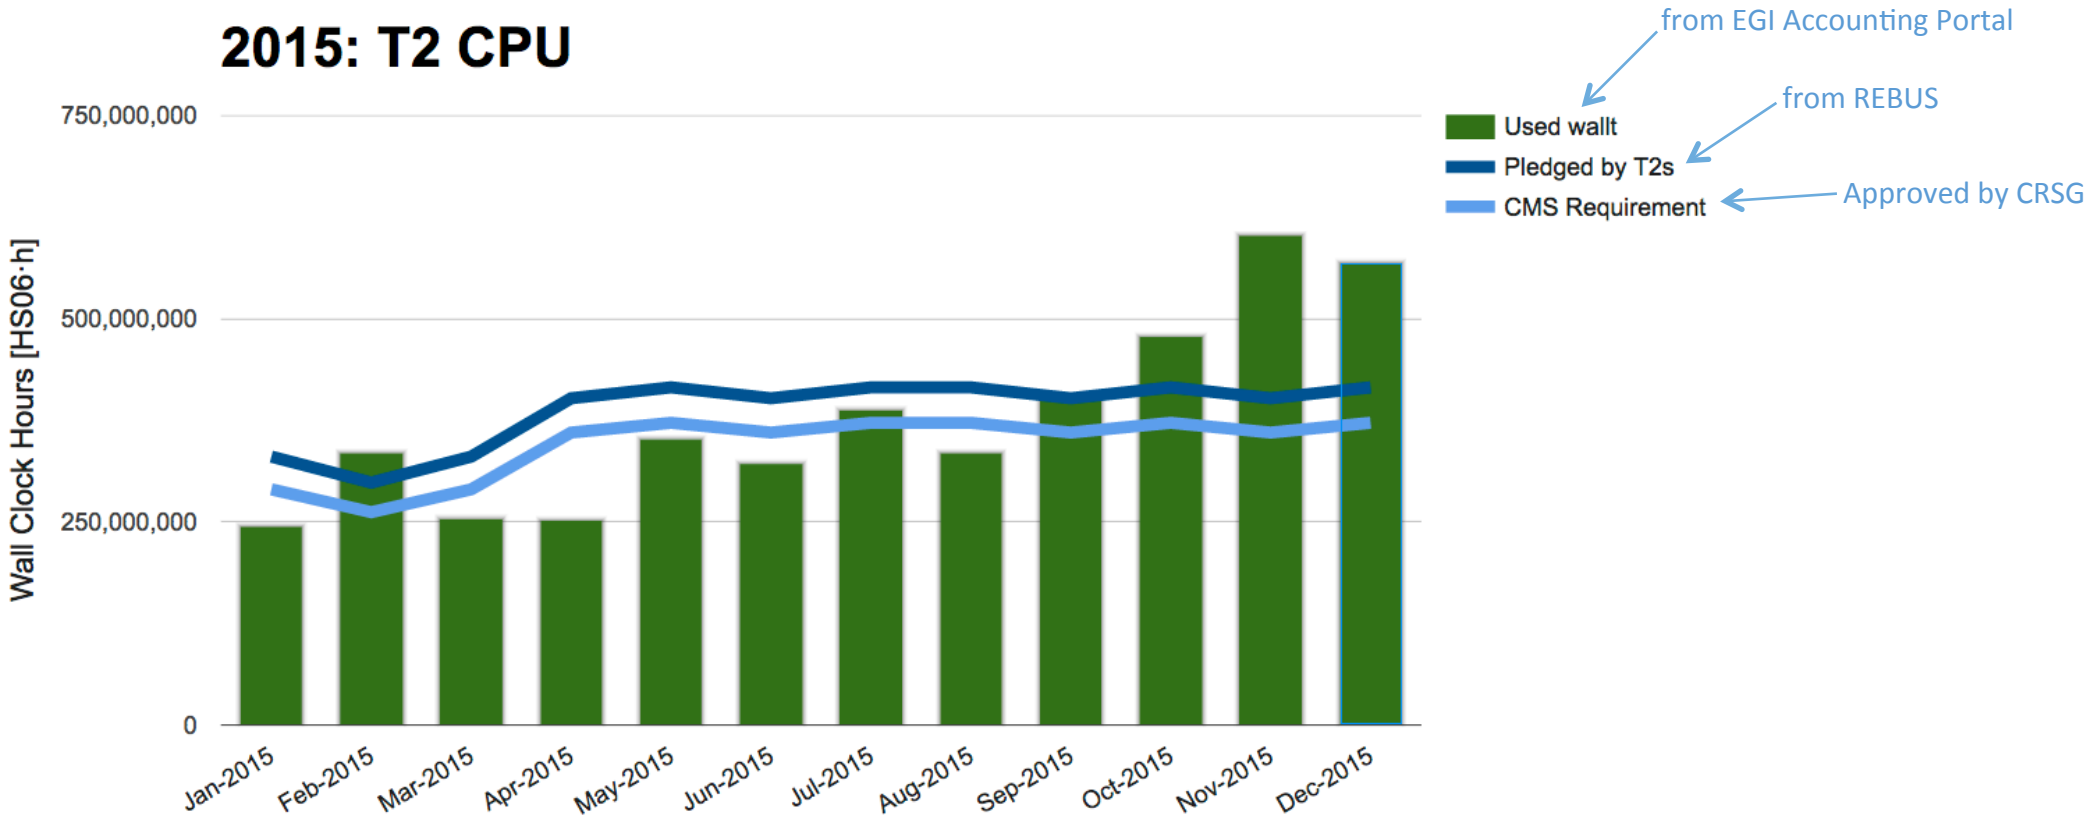

- The accounting information is extracted from the EGI accounting portal
  - No major issues observed so far, but in 2016 we are running MultiCore jobs in Tier-2 (need to check in details if things are ok – they should be)
  - We don't know the installed disk capacities in the Tier-2s (not reported)

Conclusions

- CMS makes this document twice a year and it typically takes 1 month to prepare it, discussed it and debug things
  - No major issues observed – except for CPU accounting issues related to MultiCore
- CMS has created a group of experts to produce plots and check the Resource Utilization monthly, and promptly react to issues
  - This will improve the quality of the CRSG report and spotting issues in advance
- WLCG should minimize having accounting pages with not accurate or conflicting information
  - Many plots are shown to funding agencies
  - CRSG rely on these tools – they should be full clear and consistent
- We need to work towards minimizing the complexity of the whole resource accounting system and checking the reliability of the data