



Software and Computing report

LHCC September 2016

L. Sexton-Kennedy, D. Bonacorsi
(on behalf of Software/Computing)

LHCC - 20 September 2016

Resource utilisation


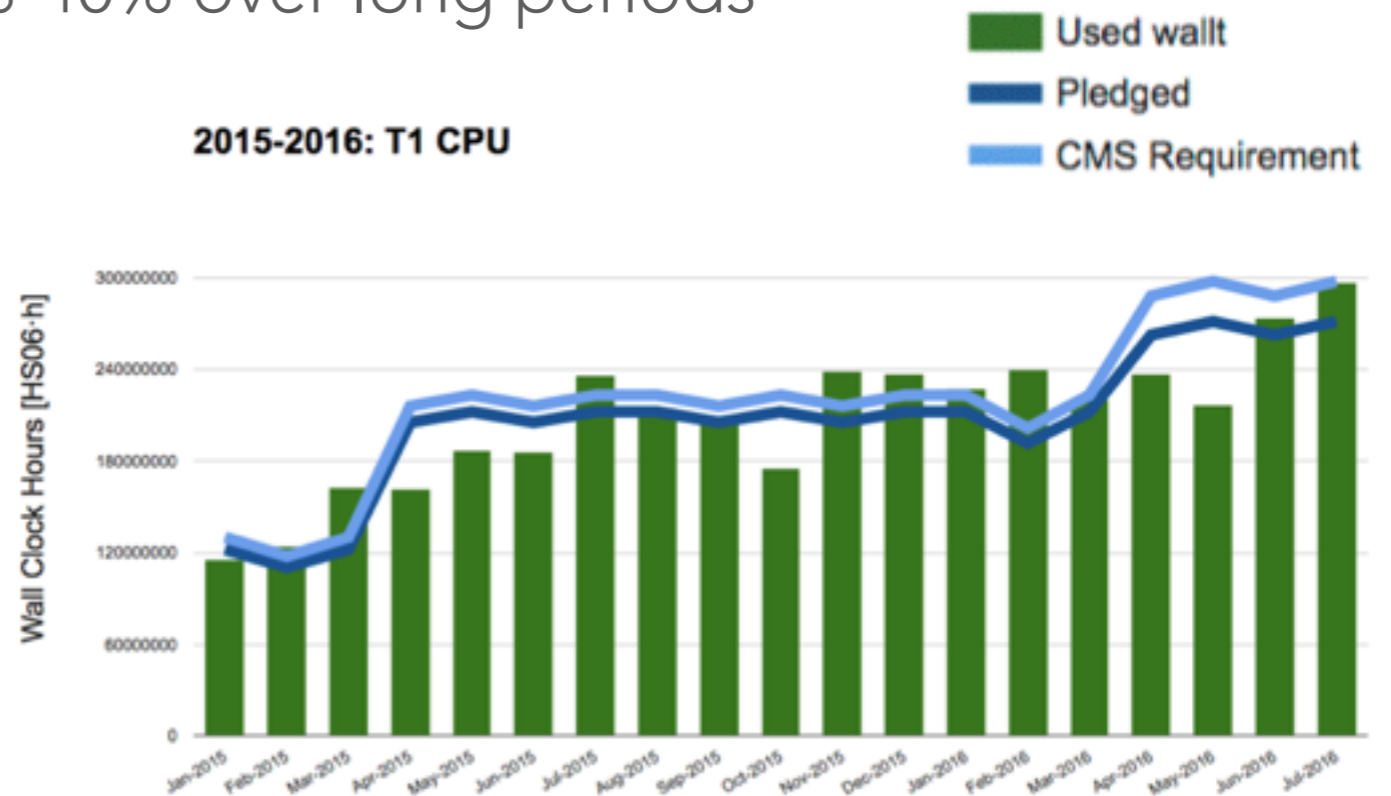
Resource utilisation continues to be very high

- ♦ production vs analysis share at 60%-40% over long periods

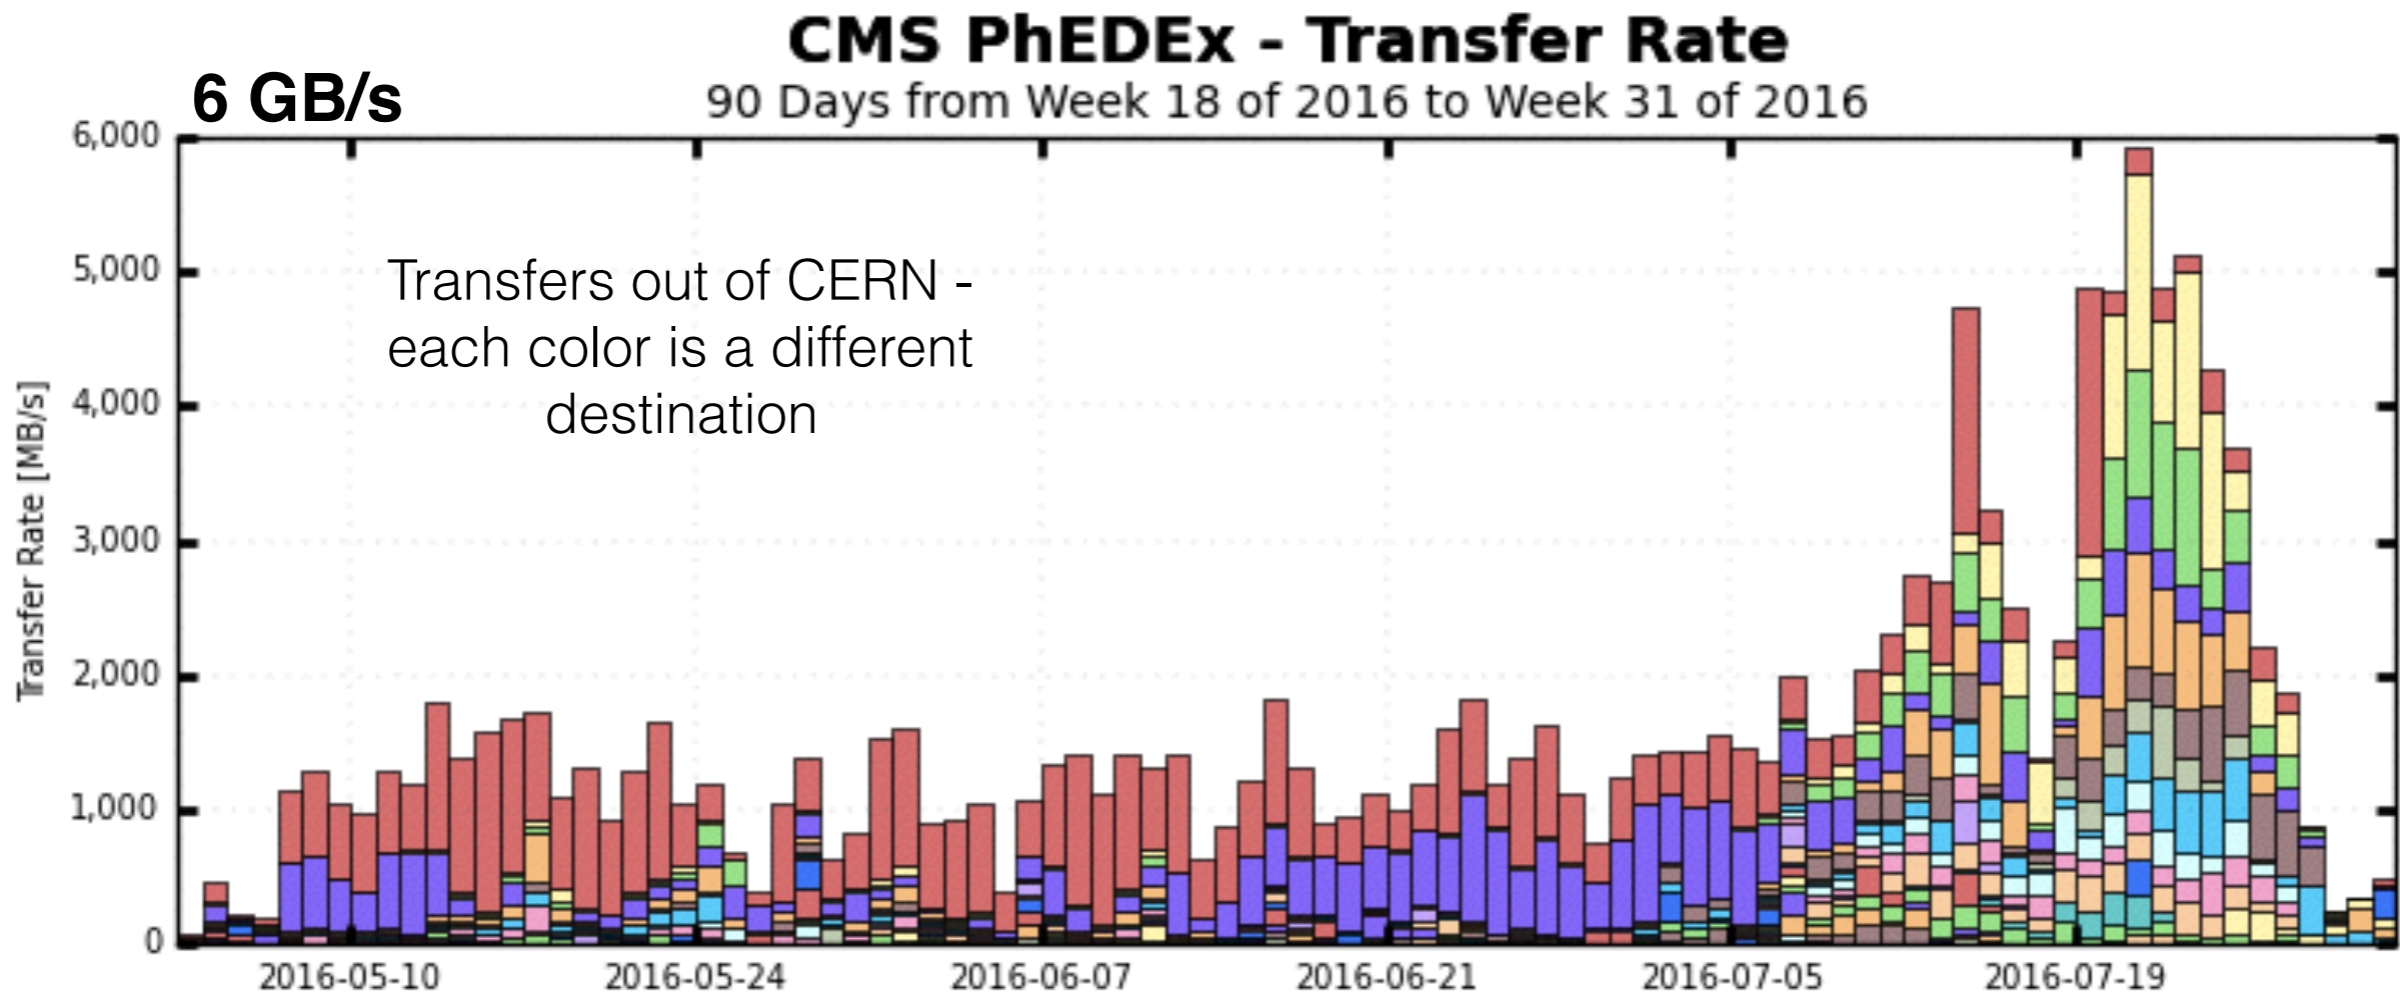
Average CPU usage at T1 level over last year was **~103%**

- ♦ note that T1s are under-pledged in CPU/Disk/Tape, situation to be addressed with the help of the CRSG

Over last year, CMS used **~122%** of the pledges at the T2 level



Improvements in data transfers



All experiments experienced issues with Tier-0 data export in mid 2016


- ◆ became evident when starting to write data faster than we were able to migrate away
- ◆ data clogged and caused processing backlog accumulation on Tier-0 resources

Problem solved - experienced rates higher than ever (up to 6 GB/s)

- ◆ CERN EOS infrastructure was improved, networking to Wigner was fixed

Very good example of collaboration among all exps and CERN experts

Simulations


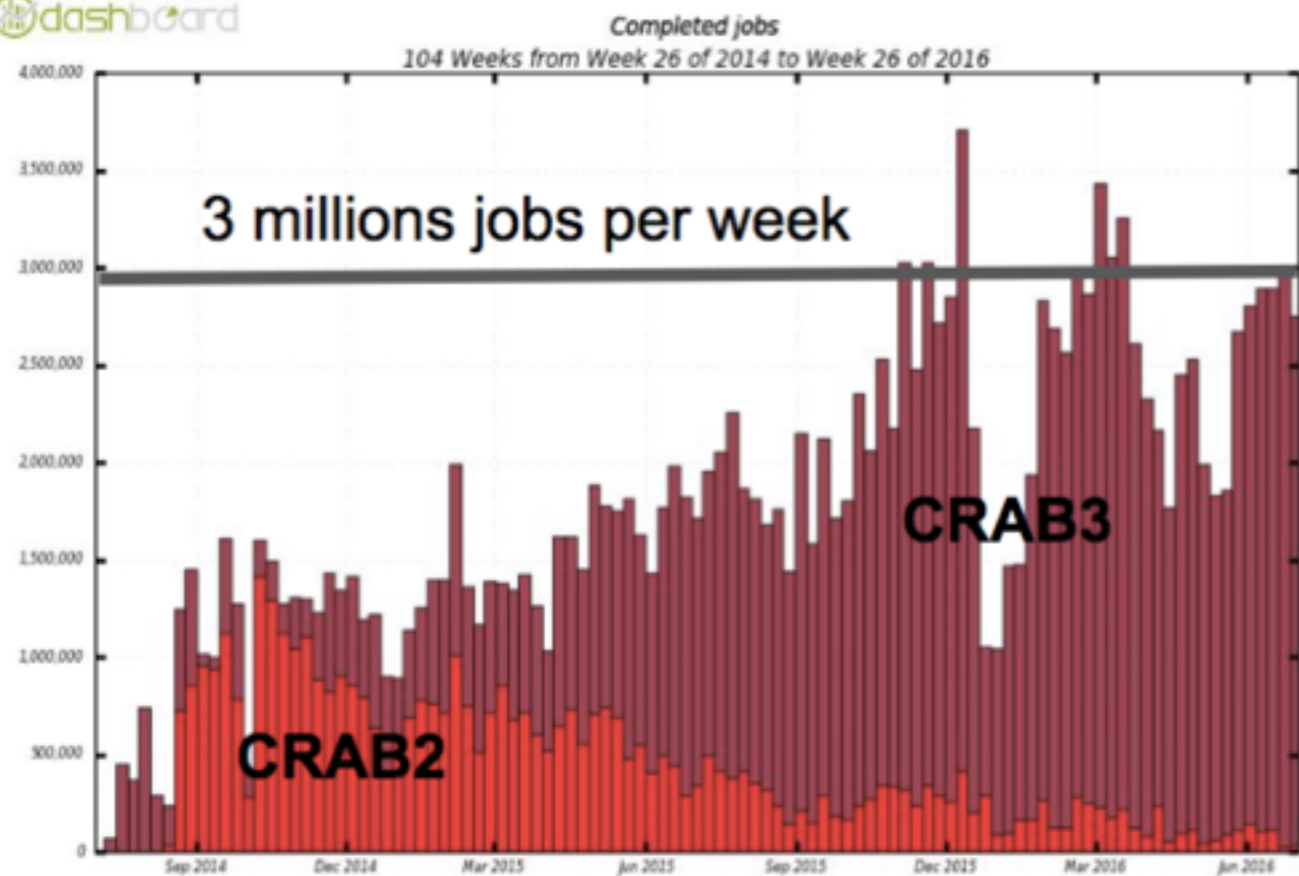


Massive Monte Carlo productions, focus on providing lighter reduced data formats

- ◆ large number of AODSIM / MINIAOD evts produced compared to GEN-SIM-RECO

Distributed analysis on Grid (1)

dashboard




Successfully completed migration to "CRAB3"

- ◆ the latest version of the CMS Grid analysis toolkit

Average number of job slots used for analysis at Tier-2 sites **increased by 56%** over the 2015 average

- ◆ similar failure rates as in the past

Distributed analysis on Grid (2)



Consistent load from distributed analysis users as Run-2 resumed

Average total number of individual submitters per month at Tier-2 sites during the first half of 2016 ramping up to Run-1 levels

- ◆ during recent months about 20% of the collaboration submitted at least one Grid job



Computing resources in 2017/18

Higher LHC live time and increased $\langle \text{PU} \rangle$ pose challenges

- ◆ similar situation for all LHC experiments
- ◆ increase of requests with respect to the Spring'16 requests

Resource requests docs submitted to the **CRSG** on **Sep 5th**

- ◆ scrutiny in progress now, in preparation for the October RRB

All possible margins of flexibility in the Computing Model have been carefully evaluated and exploited to mitigate the requests

- ◆ see next

CMS is short of resources in 2016 already

- ◆ e.g. -12% underpledged in tapes
- ◆ issue attacked by aggressive tape deletions at T1s



Mitigations

We worked hard to exploit all possible handles in the model to mitigate requests

- ◆ updated LHC input parameters: July 29nd
- ◆ CRSG documents delivery deadline: Sep 5th

Few examples:

Reduction of # replicas (and # versions) of AOD and MiniAOD

- ◆ Impact: data availability, need to rely more on xrootd remote access

Decrease amount of RECO and RAW on Disk

- ◆ caches clean-up, ~50%/65% available for 2/3 months respectively, RECO to T0, ..
- ◆ Impact: inconvenient for detector commissioning; RAW on tapes reduces efficiency in data reprocessing

Decrease Run-1 data and MC availability on Disk

Launch a massive (~30 PB) **tape deletion** campaign (in progress)

All this will have some impact on Physics, but necessary to achieve an increase in resource requests of **+20% overall instead of +40%**

- ◆ the latter would have been the relative increase w.r.t Spring'16 if we had just applied LHC params

Software Plans

Development focus is on preparing for 2017 detector upgrades

1. All new pixel detector which will have a 4th layer of pixels

- ◆ Especially important as pileup increases. The upgrade addresses data flow limitations in current detector, and removes dynamic inefficiencies.
- ◆ Biggest impact for software is new track seeding capabilities given 4th layer. The new cellular automata algorithm is looking very promising.

2. Upgraded electronics in HCAL endcap

- ◆ This upgrade was previously planned for LS2 but has been partially moved to the EYETS given aging effects observed during data taking.
- ◆ Use the new TDC capability in local reco. and exploit HE multiple depth read-out to improve energy measurement and suppress pile-up.



Software Milestones

- Create a release capable of generating data for the 2017 detector up to the level of simulated hits. We will also update the G4 release. This is targeted for 10 Oct 2016, CMSSW_8_1_0
- Create a release digitize and reconstruct the upgrade MC in order to produce “high-level” calibration samples for 2017 detectors. May also need to be the Phase2 TDR release. Feature freeze in Nov. CMSSW_9_0_0
- Create a release that will be the final digitization/reconstruction release for 2017. Schedule TBD. It needs to be in time to create final MC samples for the target conferences. CMSSW_9_1_0



Summary

CMS Resources are being utilized at the predicted levels. Given the shortfalls in pledging to CMS this is really hurting. CMS would like to revisit the pledging system to make it more flexible.

Distributed analysis continues to ramp up to Run1 levels.

We have been forced to delete non-legacy Run1 data from tape earlier than originally planned (~30PB) and will run short with no contingency in 2017.

We have an aggressive schedule for software, but these early milestones are needed to fit in the computing resources and the timelines needed by CMS.