



Meeting with LHCC

M. Girone, D. Lange, D. Bonacorsi

December 1st, 2015

Outline

Last LHCC (Sep'15):

- ◆ CMSSW threaded framework, code improvements (sim, reco)
- ◆ flexibilities in the Computing model as from the work done during LS1
- ◆ status of resources in Run-2, projections for 2016-17 (C-RSG was scrutinising)

Today:

- ◆ data taking in Run-2
- ◆ few highlights on improvements since last meeting

Next (2016):

- ◆ shifting more and more towards ideas and plans for Computing beyond Run-2



Offline + Computing

Merging into a joint Offline + Computing project

- ◆ working more closely together
- ◆ an opportunity to refresh, rethink, streamline

Maria / David / Daniele in charge from Sep 2015 to Dec 2015

David / Daniele in charge from Jan 2016 to Aug 2016

Daniele / X in charge from Sep 2016 to Aug 2017

Maria wants to thank the LHCC referees team for the useful discussions and input over the years!

pp physics @ 13 TeV



- ◆ 2.7 fb⁻¹ validated for physics @ 13 TeV with B=3.8 T

pp reference run @ 5.02 TeV

- ◆ 28 pb⁻¹ of data recorded (94% efficiency) all with B=3.8 T in 5 days
- ◆ Crucial for H1 physics run

HI run **Pb-Pb** @ 5.02/nucleon

- ◆ collecting data now at full swing





Preparing final 2015 CMS datasets

Final 2015 CMS datasets for MC and re-reco of data with:

- ◆ Final alignment, efficiencies, calibration corrections, reco algorithms, ..

Major **MC production requests for Winter conferences** have already been submitted

- ◆ Situation may change, but presently >3 billion events in the queue

On track for launching **re-reconstruction** of Run II data before Christmas

T1/T2 resources utilisation


Average use of Tiers over Jan-Oct 2015:

- ◆ T1s: **98%** of the pledge
- ◆ T2s: **82%** of the pledge
(wider variations at the T2 level)


Still quite high utilisation.

*CAVEAT: T1 data from the accounting dev portal,
T2 data from the accounting prod portal
(so, most recent may be less reliable)*

2014/15 T1



2014/15 T2




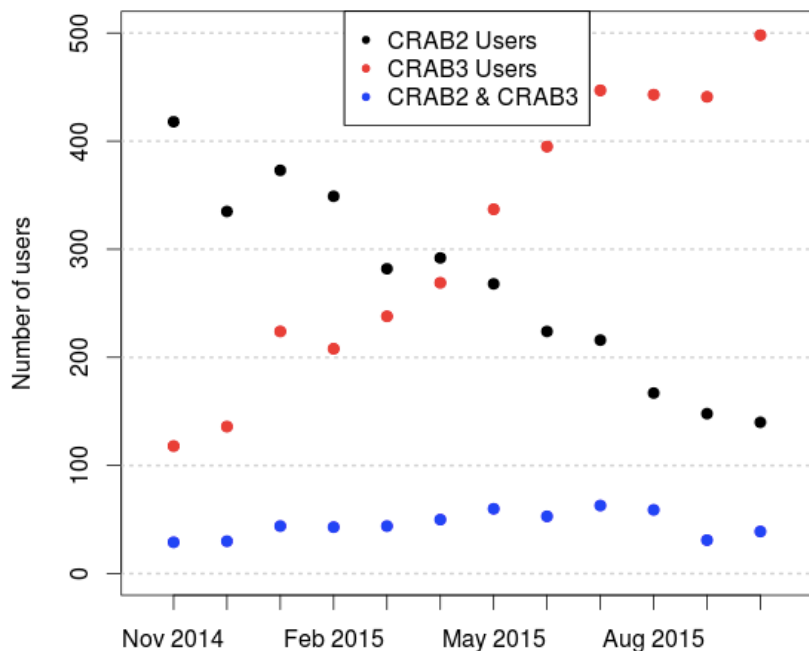
Distributed analysis

CRAB3 is the latest version of the CMS distributed analysis toolkit

- ◆ improvements in client performance, more automation in the resubmissions, a better centralised handling of user outputs, and more

In the second half of 2015, CRAB3 has become the version adopted and used by the vast majority of CMS analysis users


- ◆ the scale of analysis jobs has doubled in the last few months, up to ~2.5 Mjobs/week or ~350k/day, and the success rate is slowly improving towards ~70%
- ◆ now we can phase out CRAB2 soon, focus on CRAB3 support as well as further evolutions



Global Pool

Global Pool for resource provisioning via glideInWMS

- ◆ Reached ~**150k** jobs running in parallel (mostly scheduled processing jobs)
- ◆ Can operate all T1/T2/opportunistic resources in a single pool
- ◆ Allows central control of job priorities




Workflow execution improvements

In addition to better prioritisation and allocation across activities and sites by the global condor pool:

Latency for completing digi-reco reduced by almost a factor 3

- ◆ **many improvements to WMAgent** (the CMS scheduled processing tool)
- ◆ **higher automation** in the workflow execution and monitoring phase



Continue to explore all possible optimizations..

Data Federation

AAA allows CMS applications to read data efficiently over WAN, thus relaxing the constraints on datasets location and workflow execution

Stable operations in Run-2 so far

- ◆ some of the instabilities observed 6 months ago have largely been resolved

Introduced the concept of a transitional federation

- ◆ A necessary step towards guaranteeing high quality of service to the users
- ◆ procedure, metrics and documentation in place
 - results of AAA scale tests used as starting point to add a site to the production federation, and SAM, HC success rate and no GGUS ticket about data access for the past 2 weeks to keep it there

HLT as a cloud resource

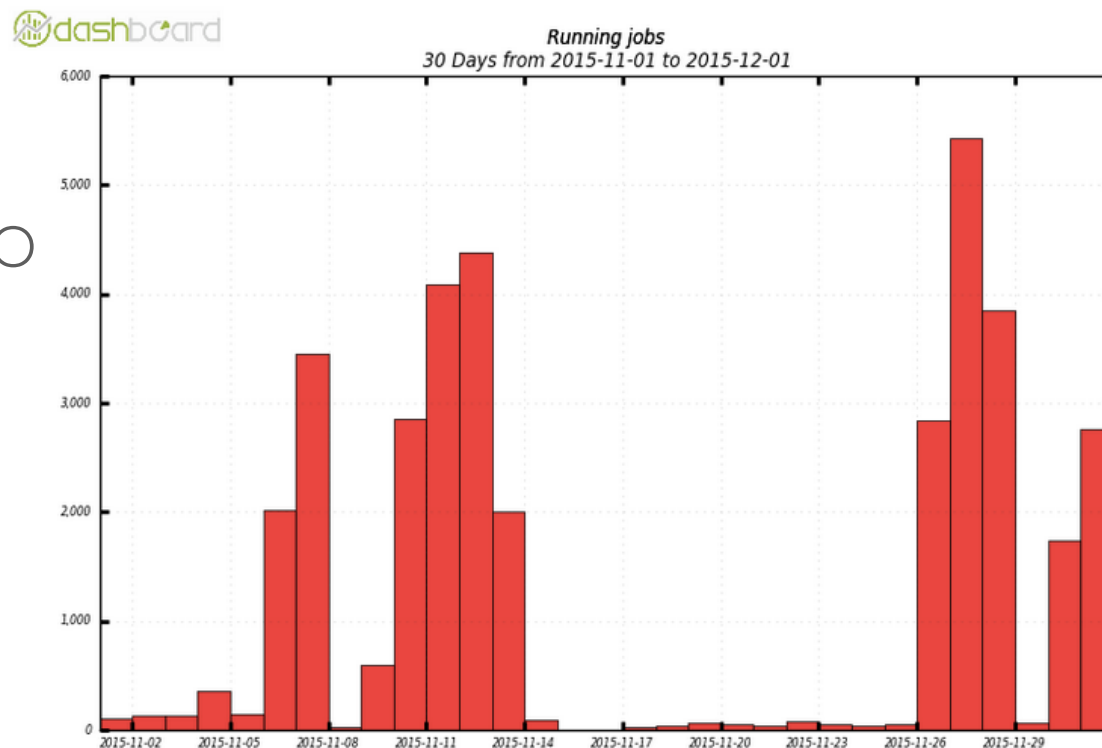
Another addition in Run-2 was the use of the HLT farm for offline processing

- ◆ Large computing resource
 - low RAM, requiring 2.5 GB restrict to ~5.5k cores. Network link from P5 to the CC is at 4*40 Gbps
- ◆ based on OpenStack
 - DAQ starts/kills the cloud, VMs join automatically the predefined HTCondor pool

Offline workflows


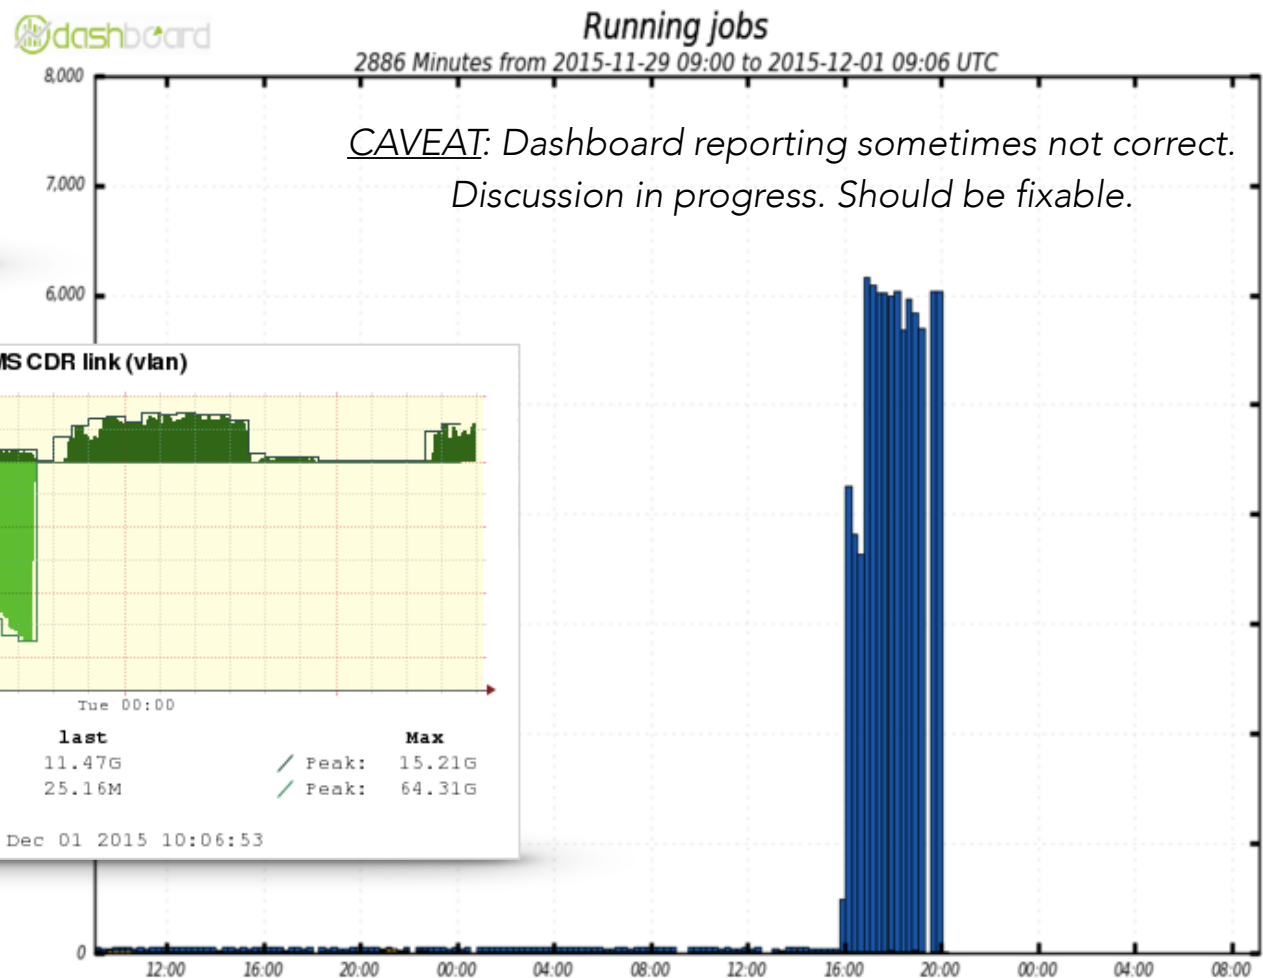
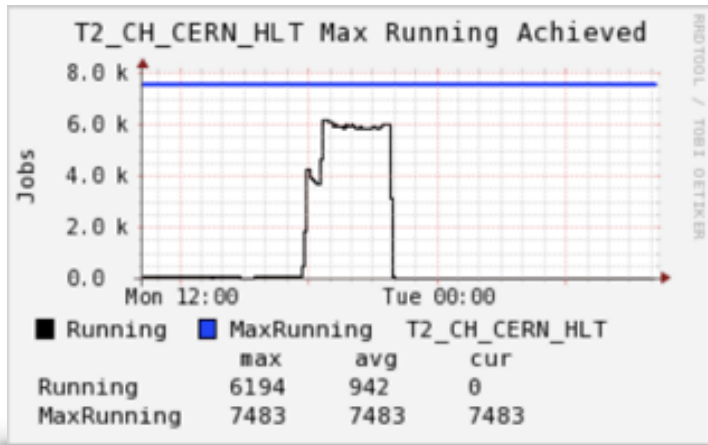
- ◆ GEN-SIM, DIGI-RECO, data RECO

Resumed tests to exploit it also in the inter-fill periods



Zoom on one recent example

E.g. we used the HLT in the inter-fill yesterday evening for 5 hrs





AWS project status

Goal: demonstrate the ability of executing on AWS any of the CMS centrally-organised workflows

- ◆ Project duration: June 2015 - March 2016

AWS credit to be used in a large-scale workflow that would be visible to the CMS collaboration

- ◆ demonstrate the value of dynamically provisioned external computing resources for GEN-SIM, DIGI-RECO and data RECO

Target scale is 56k simultaneously running cores for 1 month

- ◆ ability to ingest data through Xrootd at a sufficient rate to maintain a high CPU efficiency, plus ability to export the data produced

Project Status: CMS workload and data management infrastructure adapted to run on S3 and AWS

Scale tests starting soon, followed by large scale production

- ◆ 1%, 5%, and 10% of the major organised processing workflows

Quite active in exploring:

- ◆ New architectures, low power
- ◆ Intelligent Networking, Integrated Storage and Networking
- ◆ Specialised computing (FPGAs, GPUs, ..)
- ◆ Big Data analytics, Machine Learning techniques
- ◆ ...

Some activities will also be explored in synergy with CERN open lab and connected labs/experts