CMS Overview

Chris Brew, Katy Ellis

19th July 2021

Science and Technology Facilities Council
Particle Physics

Physics Goals

- Runs 1 and 2 150 fb^{-1}
- Run 3: additional 250 fb^{-1} (~1.7x gain from Luminosity and energy)
  - Consolidate Run-2 observations of rare SM processes
  - Experimental improvements (trigger, analyses, systematics) in preparation for…
- HL-LHC: 3000 fb^{-1} and superior detectors
  - Standard model: discovery through precision
  - Higgs: H(125) properties at the percent level
  - Direct searches: discover new physics or close a few chapters

Challenges

- Budget-per-luminosity falling faster than computing costs
- Flat cash projections indicate a 4x shortfall in CPU and a 2x shortfall in disk by the end of Run4
- Mature experiment means a lot of history and reduces the ability to take off in radical directions

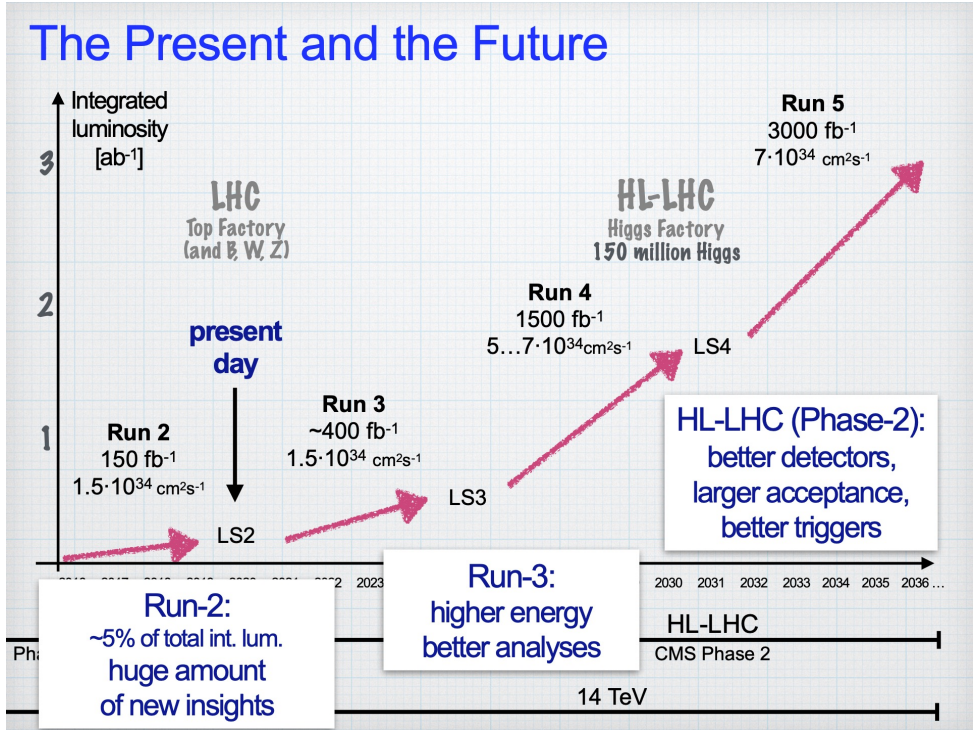

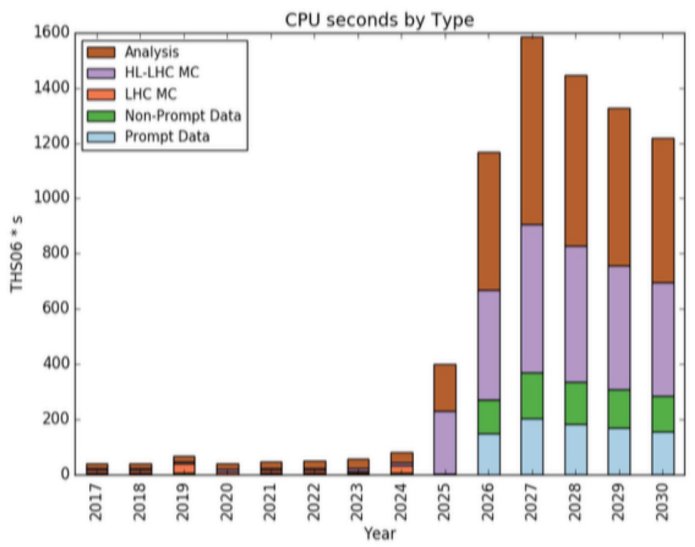

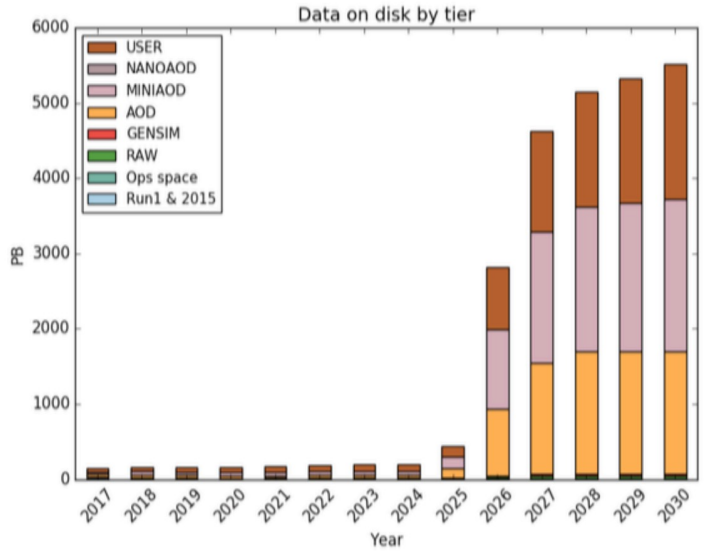

Science and Technology Facilities Council

Particle Physics

Trigger/Online

- Higher Luminosity means more triggers and more complex events.
- Improved detectors and readout means more information available to the trigger, same fixed time budget to make a decision.
- Increasing use of accelerators in the trigger:
  - FPGAs in readout/L1 → Tracks and Particle Flow at L1!
  - GPUs in HLT
  - More interesting physics, less background
- Readout less – Scouting saves greatly reduced events (1.5kB) at high rate
- Reconstruct less – Parking saves raw for later reconstruction

HLT

- In Run 2 the CMS HLT farm was limited by offline storage and processing power
- Phase-II computing needs are currently forecast to be unaffordable unless solutions such as GPUs or FPGAs are adopted to reduce cost
- Run3 needs no major change beyond a modest capacity increase to support things such as data scouting, however being used a test bed to gain experience of heterogeneous farm
  - Goal - to add a T4 GPU to each node
  - Currently 30% of the menu is offloaded to the GPU
- Major challenge: maintaining separate GPU and CPU code bases
  - Use a portability library such as Alpaka to enable "write once, run anywhere" code
  - Ensuring that the code is still performant is a challenge
  - Secondary challenge is ensuring that we support and gain expertise to program in a GPU/FPGA friendly way

Offline and analysis

- CMS Offline Core Usage:
  - 41% Production
  - 36% Processing
  - 22% Analysis
- Reduced data formats
  - AOD (400-500kB/evt) → MiniAOD (40-50kB/evt) → NanoAOD (1-2kB/evt)
  - Saves disk space but also network bandwidth and processing time
  - Use still not universal, still gains to be made
- Production
  - Ongoing efforts in improvements to generators, simulation and reconstruction
- Multithreading
  - Many central processing jobs are multi-threaded
  - Lots of generator production still single threaded
  - User analysis code mostly single threaded

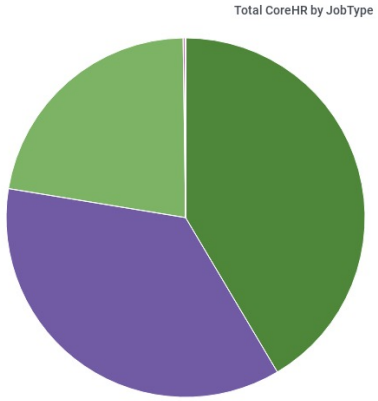

	total	percentage
Production	782 Mil	41.5%
Processing	681 Mil	36.1%
Analysis	419 Mil	22.2%
Merge	3 Mil	0.2%
LogCollect	498 K	0.0%
Cleanup	247 K	0.0%
Express	144 K	0.0%
Repack	50 K	0.0%
Harvesting	18 K	0.0%
FarrukhTest	79	0.0%
Skim	32	0.0%
Unknown	19	0.0%
analysis	0.3	0.0%

Offline and analysis

- Better algorithms:
  - Still scope for deep delves into the code to tease out inefficiencies
  - Small gains in core tools/algorithms have widespread effect
  - Could be much easier for end physicist
- New platforms:
  - Hardware – GPUs are now a mainstream, expert tool being deployed widely with central software starting to have the option of making use of them. FPGAs a few years behind apart from specialist applications
  - Software – Pythonic analysis frameworks and big data style analysis tools. Potential to lower barriers to better code and hardware accelerators

Computing operations

- Efforts during LS2 to make software more modern, performant and scaleable
- Joining of community projects
  - e.g., replacing PhEDEx with Rucio for data management
- Increased automation to free up expert time in Run 3 to prepare for the more significant challenges of Run 4.

HPC

- CMS already uses significant HPC resources
  - On a good day we use cores equivalent to a large Tier-1
- Future HPC resources likely to be accelerator based rather than CPU based

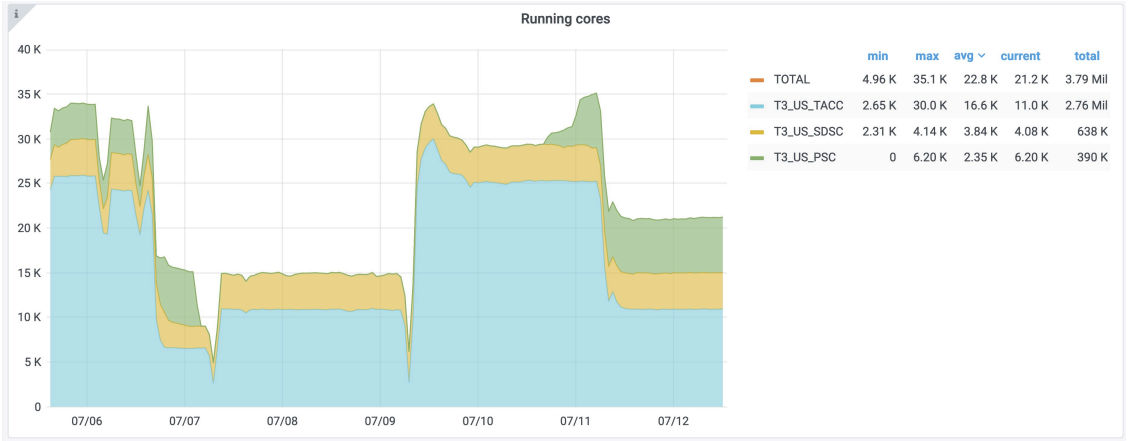

Work continues to integrate HPC sites with the CMS job submission infrastructure

UK strengths

- Significant expertise in use of FPGAs for trigger
  - Hoping to expand to more generic use as accelerators
- Coordination of GPU usage in the HLT
- Exploration in use of GPUs for User analysis code
- Efficient user computing
- Coordination of computing systems and operations

Summary

- CMS computing is in a good position for Run 3.
- However, it faces significant challenges in Run 4 when the amount of storage and processing nodes increases hugely to take advantage of greatly increased LHC luminosity.
- Several changes already in progress
  - e.g., reduced data formats, GPUs in the HLT, opportunistic use of HPC
- Lots of areas to address over the next few years to optimise use of resources