CMS Status

Matthew Nguyen (LLR)

Journées LCG-France

May 22 – 24, 2019

Mostly based on Tommaso Boccali's presentation at HOW19
https://indico.cern.ch/event/759388/contributions/3302196/

CMS Computing at the end Run 2

- LHC Run 2 concluded at end of 2018
- CMS very successful in data taking and analysis operations, with unanticipated computing challenges
  - B Parking: additional 12 B events collected in 2018 to support B Physics; 20x more data than Babar and Belle!
  - Heavy Ions: 4.5 B minimum bias events collected in Nov – Dec 2018
- On top of that, standard pp operations (64 fb^{-1} recorded), analysis operations in full swing

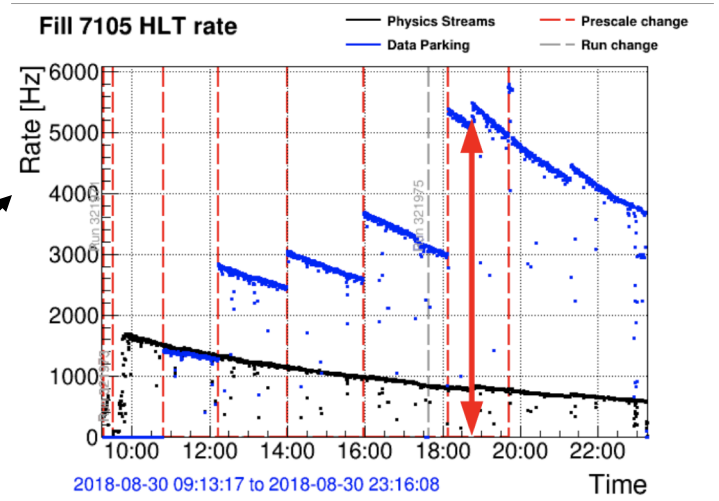

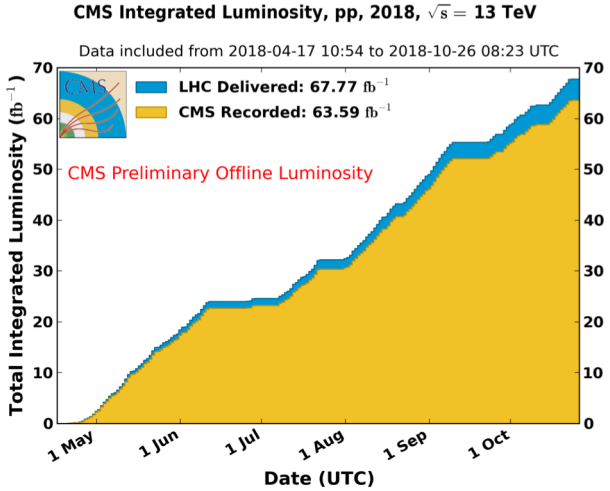

Resource Utilization

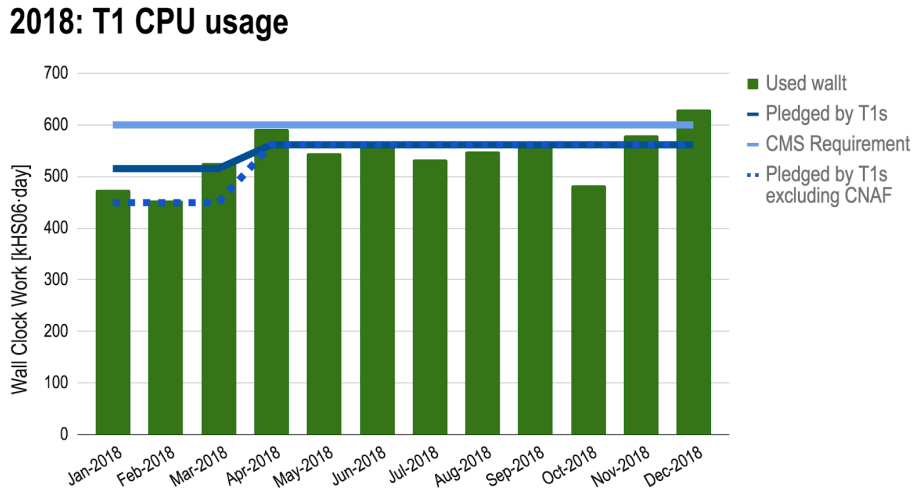

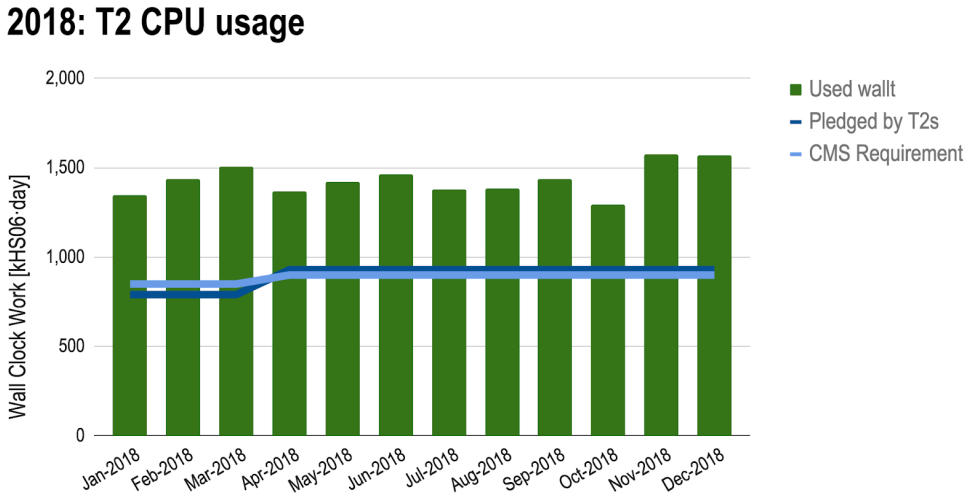

- T1: Mostly a “production” resource
  - Pledged to 90% of CMS requirement (as in previous years)
  - Loss of CNAF in first part of year compensated by other sites (both T1 & T2)
  - Pilot efficiency of 75% (+10% from 2017)
- T2: primary analysis resource
  - Pledged to 100% of CMS requirement
  - Utilized 160% thanks to opportunistic resources
  - Pilot efficiency of 66% (+6% from 2017)

Disk/Tape Utilization

Disk usage fixed near 90% by Dynamic Data Management (DDM), at both T1 and T2

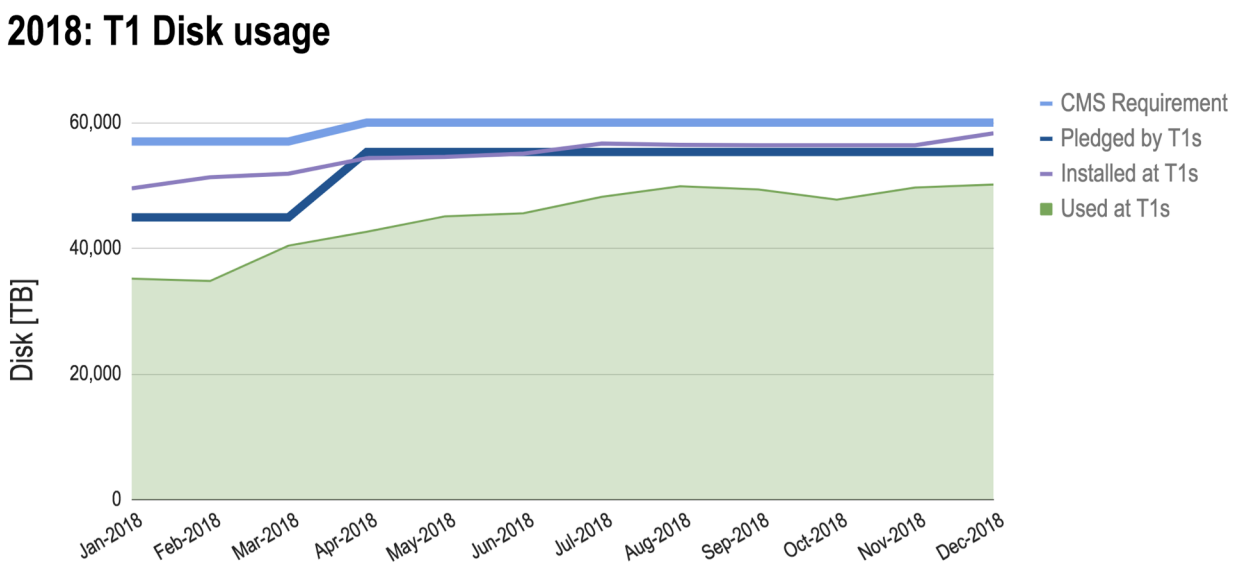

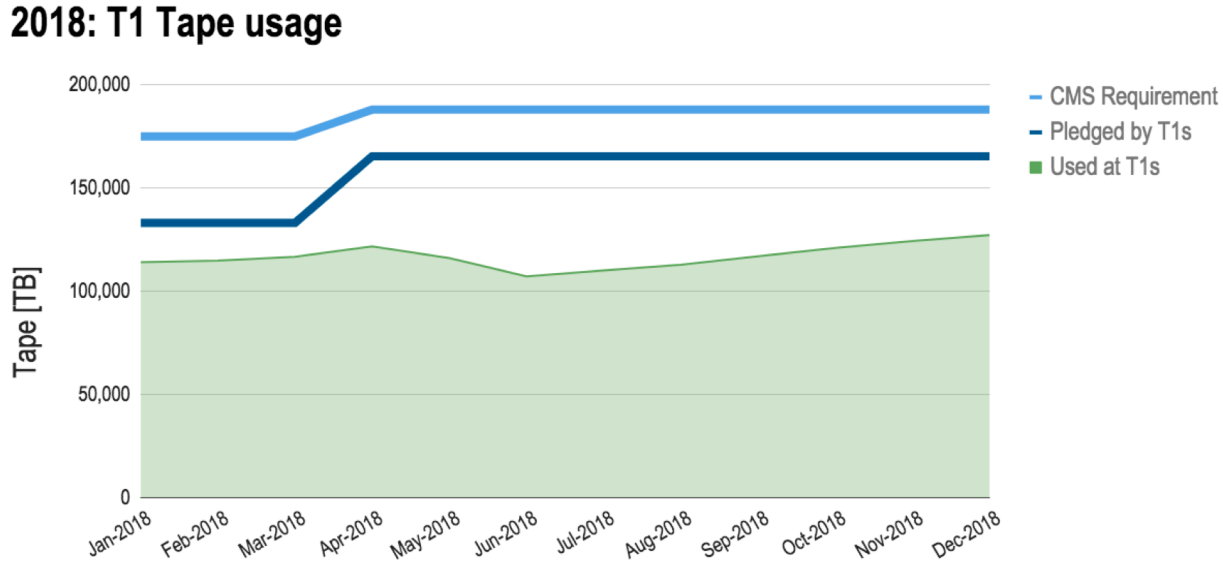

Pledged to 92% of request, 97% installed by EOY

Pledges were 85 – 90% of request
Aggressive deletion campaigns were needed to free up space for 2019 productions

Key activities during LS2

Pursuing parallel activities in three areas:

1. harvest of run 2 results
2. preparation for data taking & analysis in Run 3
3. preparation for HL-LHC

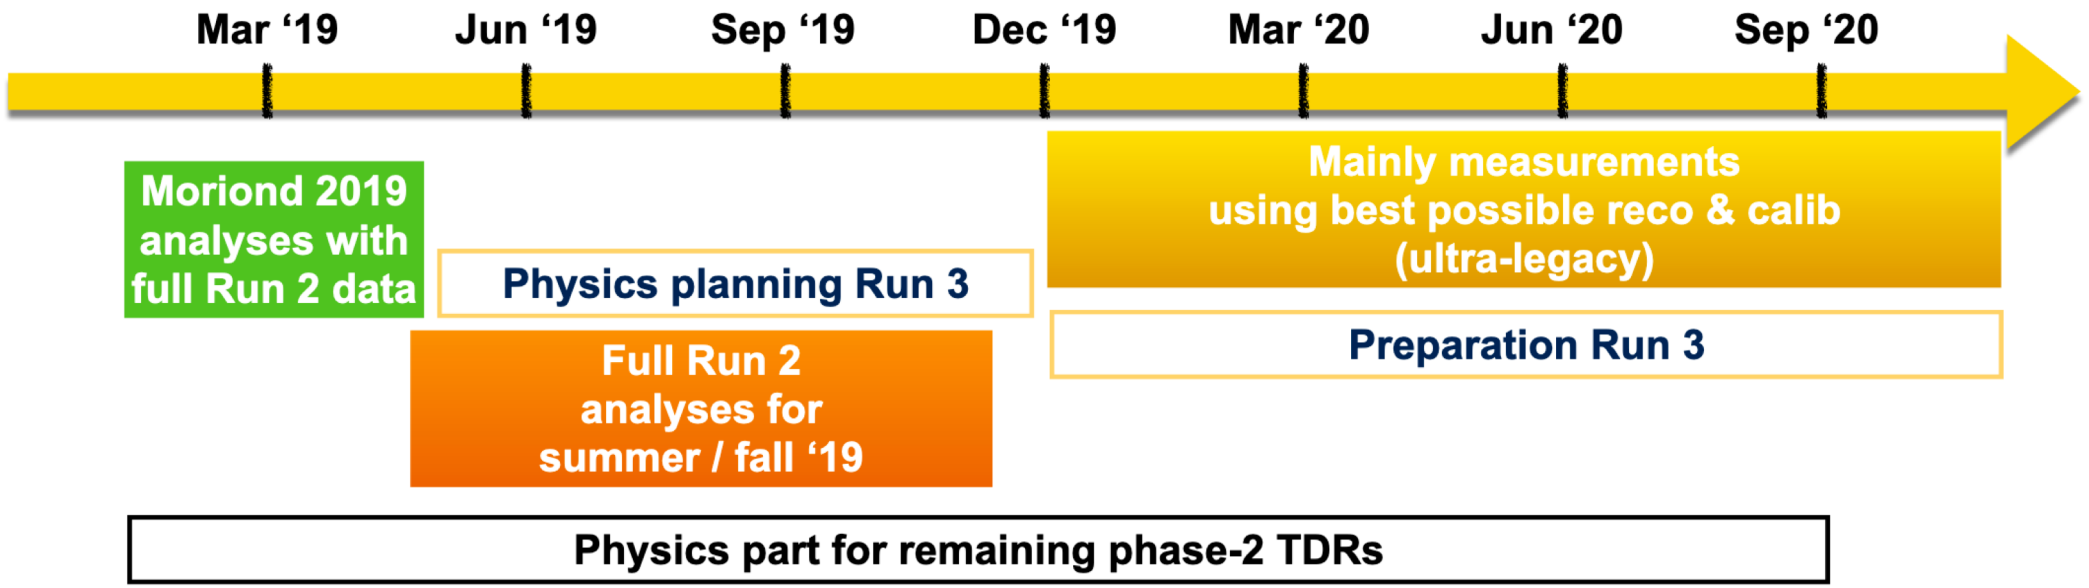

Run 3

- CMS resource requests where flat for LS2 years 2019 & 2020
- CMS currently anticipates no major change to running scenario:
  ~ 1kHz trigger rate + 500 Hz parking
- Increased PU from 35 in Run 2 to 60 for nominal Run 3
- Various gains are taken into account in model, e.g.,
  - 20% gain in simulation from GEANT & CMS sides
  - Increased adoption of nano-AOD
- Projects to 20 – 30 % increase for resources in 2021 & 2022
- Request is still under discussion w/ scrutiny group

Current Phase II Projections

- Initial projection 50-100x wrt current resources: 6x current PU, 7.5x HLT rate
- Optimistically, 4x from technological improvements, still leaves 10x to account for
- Last public version of our 2027 estimates
  - CPU: 44 MHS06
  - Disk: 2.2 EB
  - Tape: 3 EB
  - i.e., 22x, 13x & 15x wrt 2019 pledges
- NB: storage decrease by 2x from use of nanoAOD for 50% of the analyses, reducing the processing and storage of larger data formats

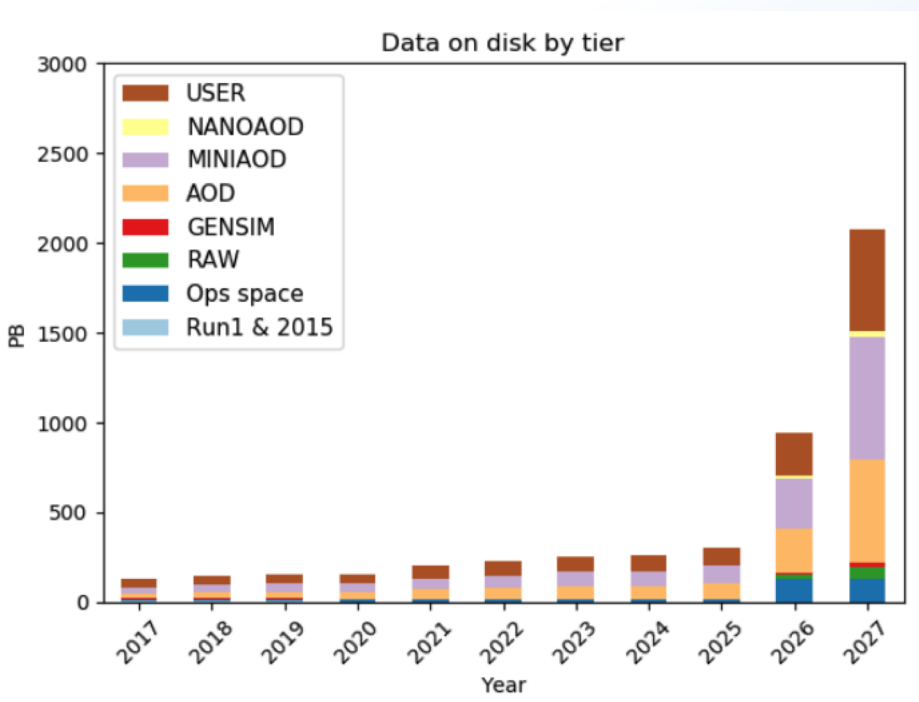

Opportunistic resources

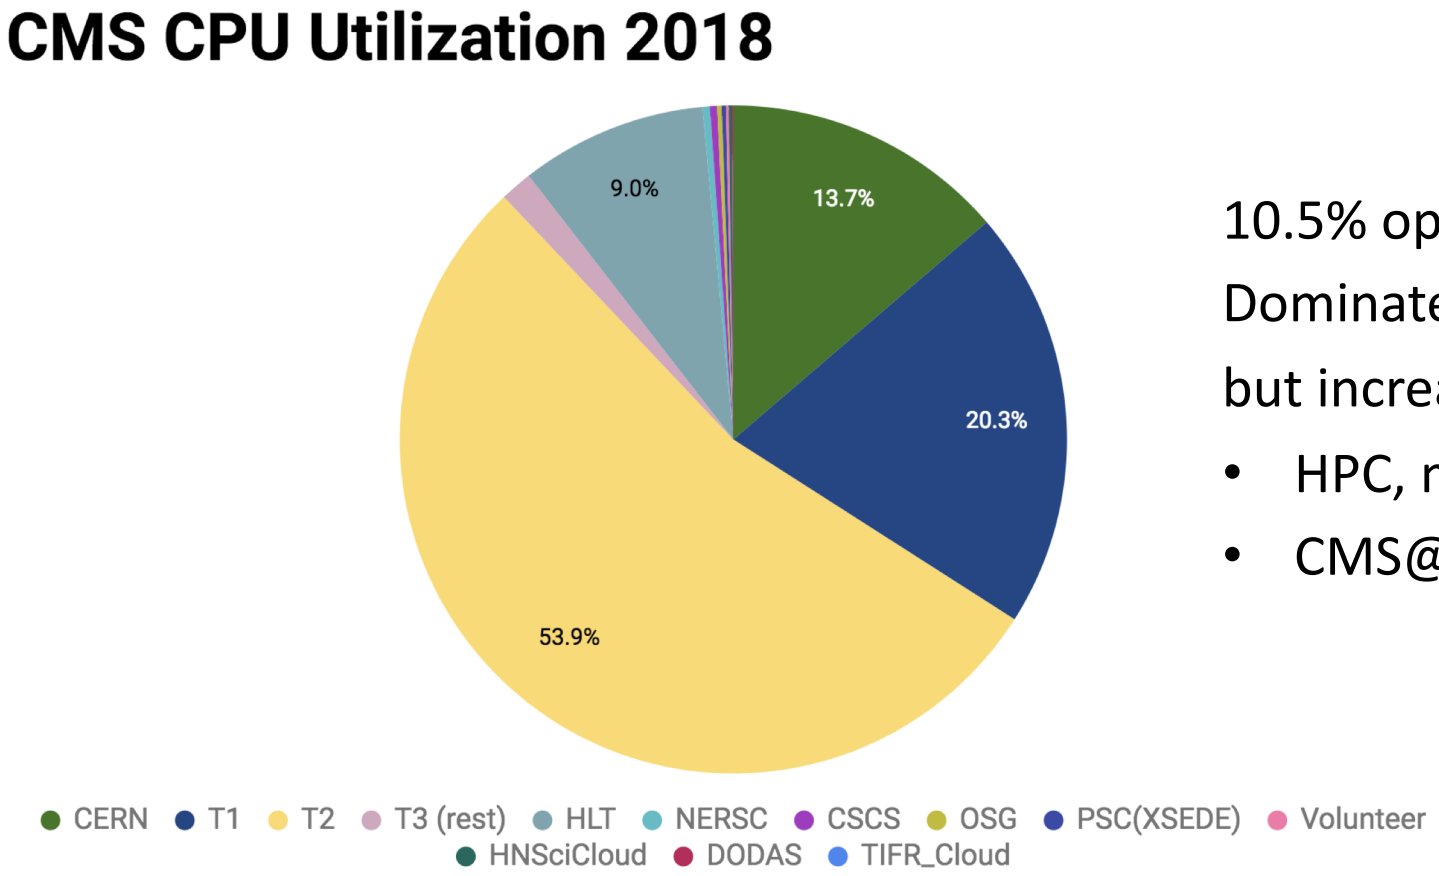

10.5% opportunistic CPU

Dominated by HLT (9%), however small but increasing use of:

- HPC, mostly NERSC in the US
- CMS@Home

HLT

- Provides up to 30k cores during interfill periods and during shutdowns
- Complemented Tier-0 for heavy-ion prompt reco (yellow line)
- Contributed 240 kHS06 to offline computing, +60% compared to 2017 (44% of 2018 CPU at T1)
- Hard to predict during Runs, but inserted into resource model during LS2
- Experience with HLT has been useful to exploit other opportunistic resources

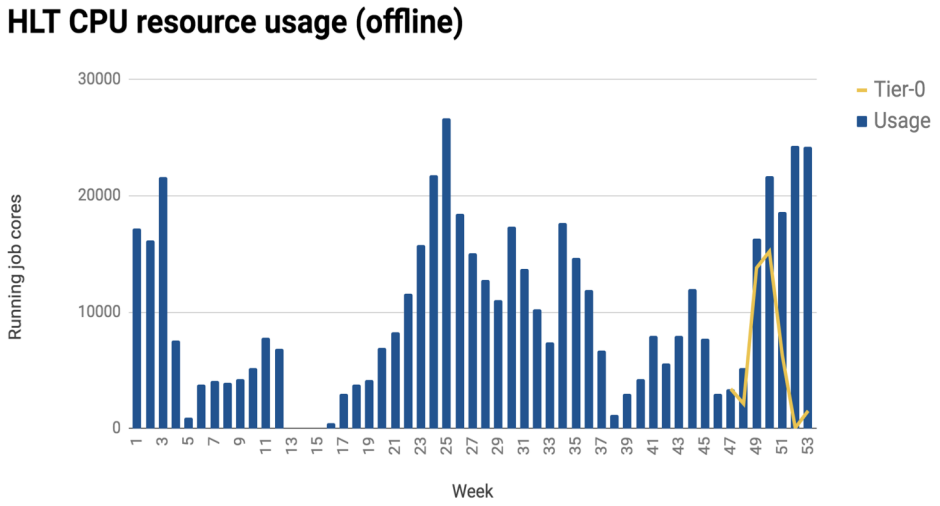

Utilizing HPC resources

- CMS (spurred by funding agencies) requested feedback on use of HPC sites
- Call was met with mixed success, with certain countries supporting active initiatives (Italy, Spain, Germany), while others are lagging behind (e.g., France)
- Clearly there are many challenges both technological & political
- CMS prepared documents addressing both set of challenges
  - Technical document highlighting our needs
  - Political introduction explaining why we seek collaboration w/ HPC centers
- CMS will facilitate, but “boots on the ground” needed at national level

Heterogeneous architectures

- There is general consensus that the best performance/$$ will not be obtained with standard CPUs
- Testbeds active on GPUs, FPGAs initially as standalone exercises
- In the last year, CMS has attempted to systematically include these into the standard CMS Software Framework:
  - Allow multiple versions of “equivalent” modules, deferring decision on which to use until last moment
  - Allow the best communication between modules exposing different interfaces (for example, chain GPU modules without moving data back to the host)
  - Have CUDA as an external tool in CMSSW, for native utilization
  - Next step (in collaboration with other experiments?) is to try and have automatic code translation in place (is it even possible?)

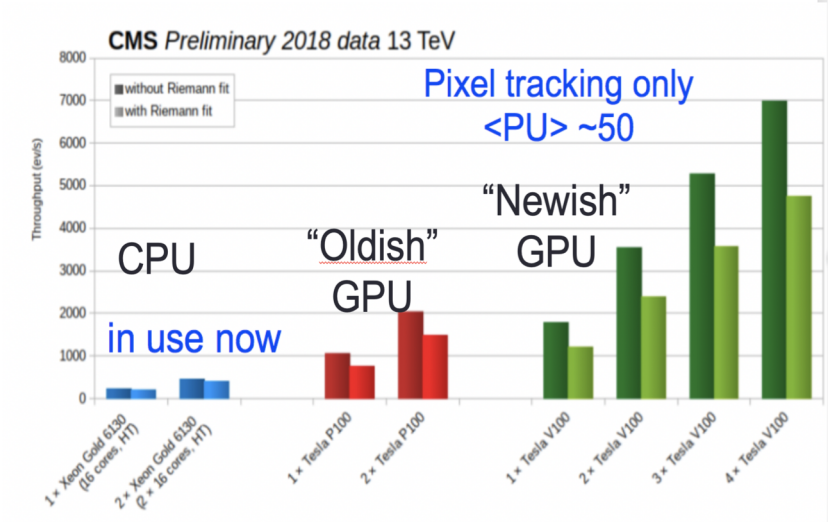

- CMS software has made steps to benchmark different architectures
- Potential gains are large, but still quite a few implementations & technologies to choose from

Common solutions

- The CMS SW stack and Computing Infrastructure were **adequate for CMS needs in RunII**, and then some.
- We have **no real hint that RunIII would pose irresolvable problems either**; but, since RunIV is a different story, CMS has planned to **try and test any disruptive technology already in RunIII**
- Among the software tools, the biggest worries in the RunIV time scale are about **software support and sustainability**. **Common solutions with other experiments are a way to mitigate the support cost**

- CMS identified **3 initial areas** where we can benefit from existing open source SW:
  - Geometry description: **testing DD4HEP from AIDA2020**; if testing is positive, transition in ~1 y
  - **CRIC from CERN** as a replacement for the Information System - already in place for the first use cases
  - **Rucio** (initially from ATLAS) as the **Data Management solution** - transition and then large scale test in ~ 1 y