LCG France

FRANCE GRILLES

Perspectives d'utilisation pour le LHC

GRIF a 20 ans – 30/06/2025

J.-B. Sauvan

LLR CNRS / École Polytechnique

cnrs

ÉCOLE POLYTECHNIQUE

IP PARIS

Outline

- The past 10 years and current status of GRIF resources for LHC VOs
- The next 10 years and the HL-LHC: a challenge in terms of computing, storage and data management
  - HL-LHC expected requirements
  - Heterogeneous computing?
  - Analysis Facilities
  - Data Challenges
- More questions than answers 😅
- Many inputs taken from the May 2025 WLCG/HSF workshop

Talk 10 ans (Luc Poggioli)

LHC evolution in French sites

- CPU accounting for LHC experiments in the different French sites since 2015
- Factor >3 increase of GRIF CPU resources since 2015
  - With some stagnation in the recent years

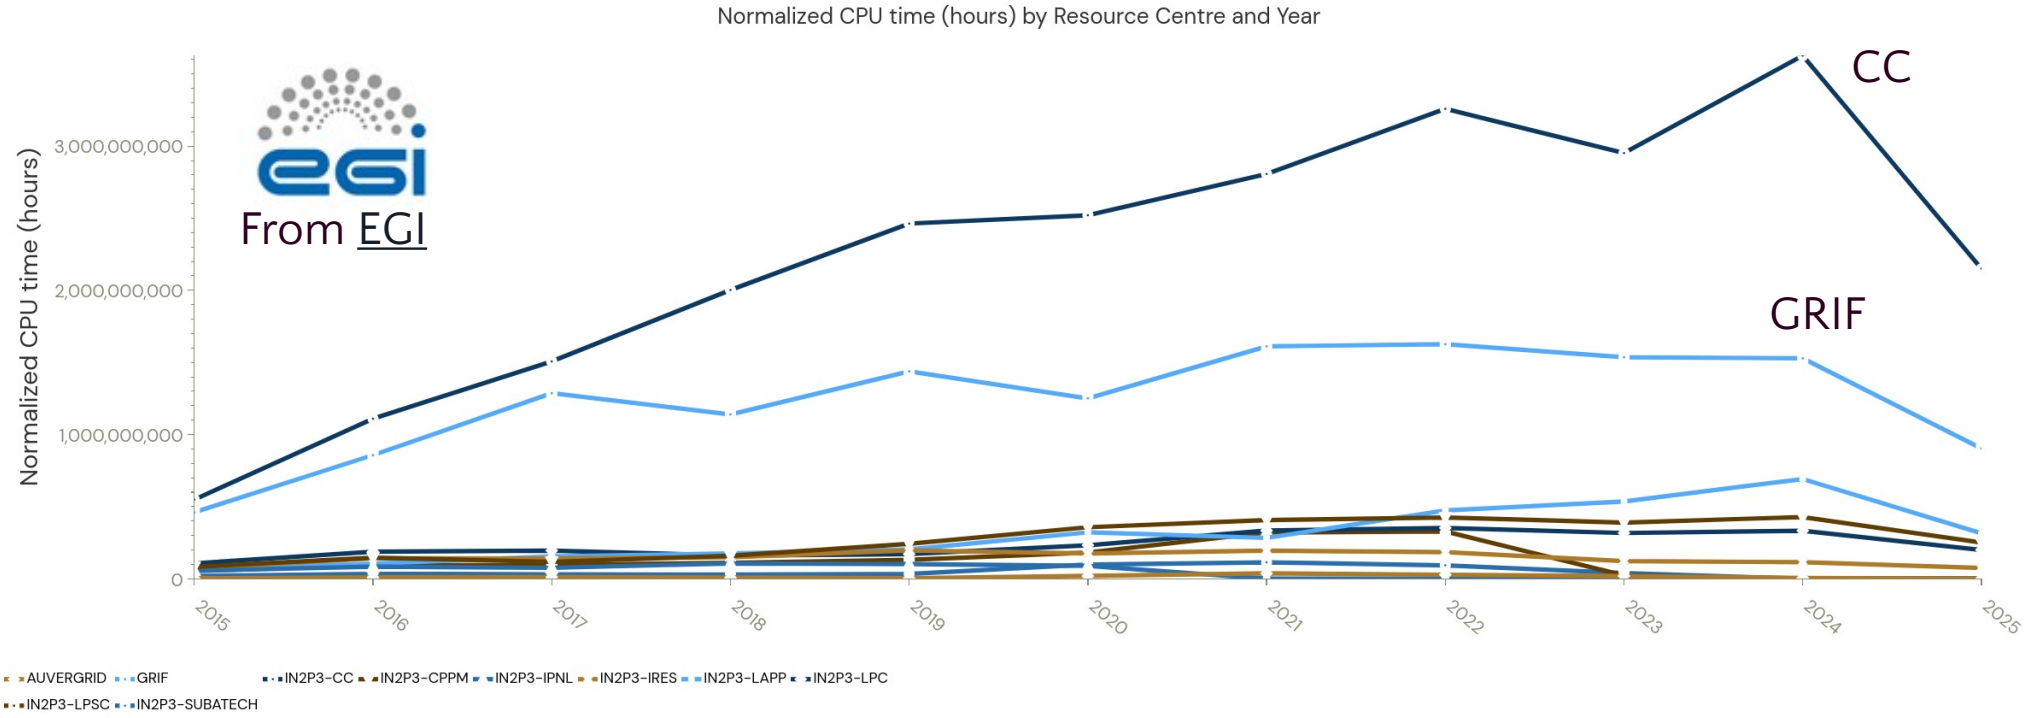

Current sharing between french sites

- Below is the current resource sharing of French sites for LHC VOs
- GRIF is the second largest French site after the CC, and first T2
  - About 20% of LHC resources (disk and CPU)

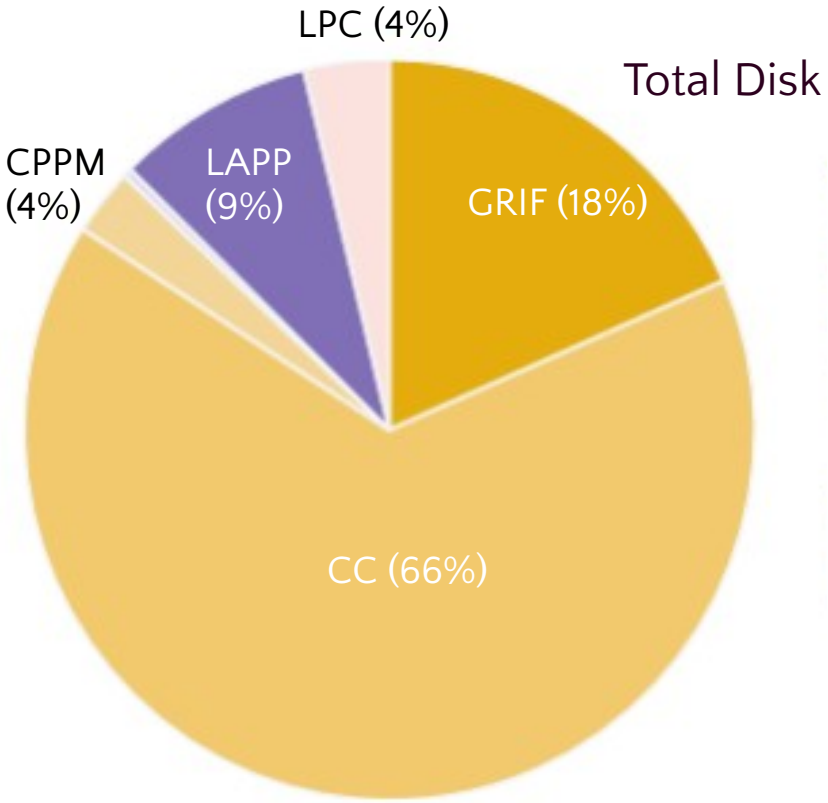

From monit-grafana

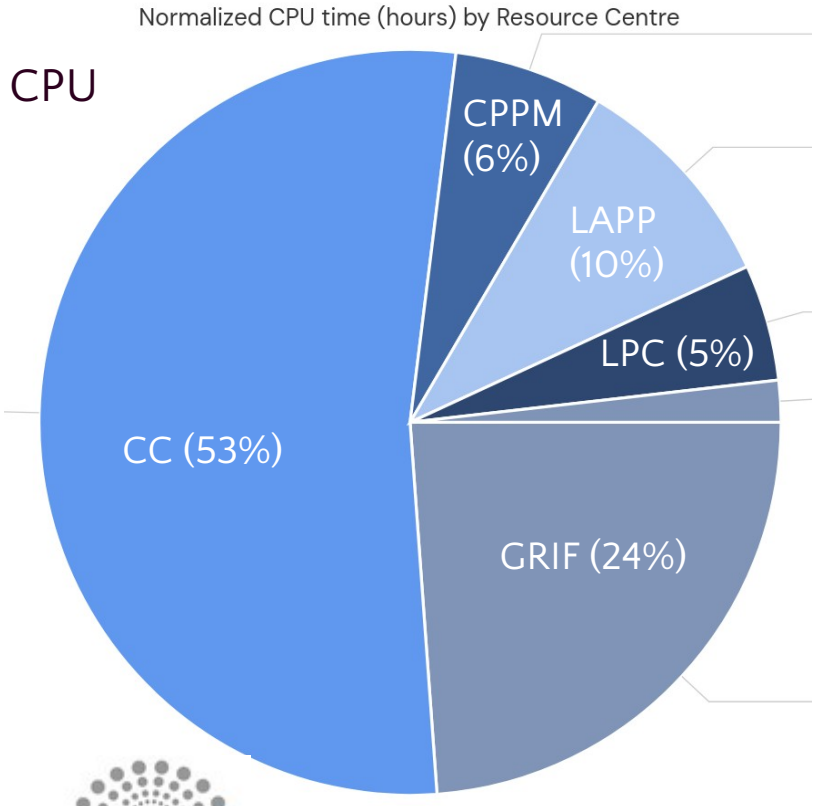

June 2023 – June 2025

VO LHC in France

- ATLAS: 50% of the LHC resources in France
- 10-20% for each of the three other experiments
  - Roughly 20% for CMS (both disk and CPU)
  - LHCb CPU share larger than disk share
  - ALICE disk share larger than CPU share

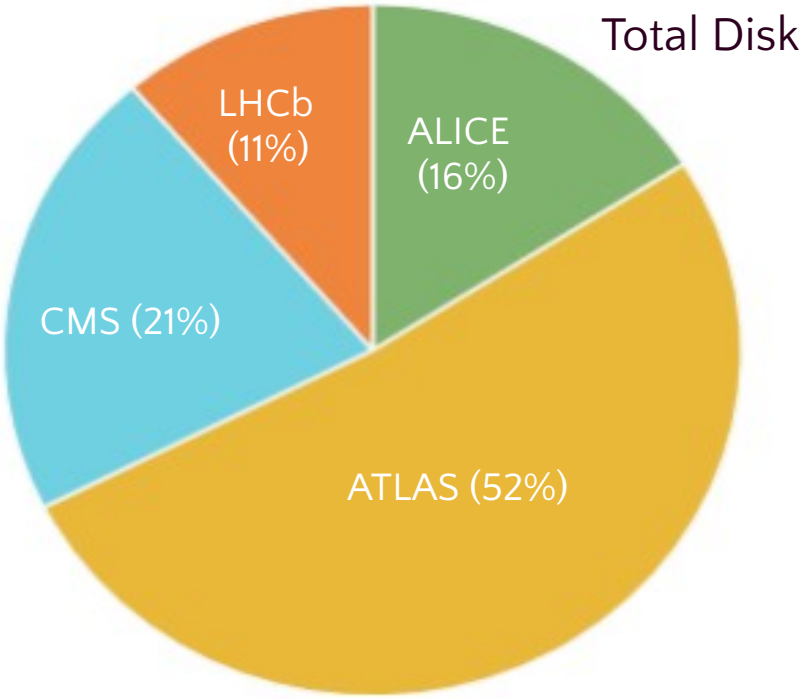

From monit-grafana

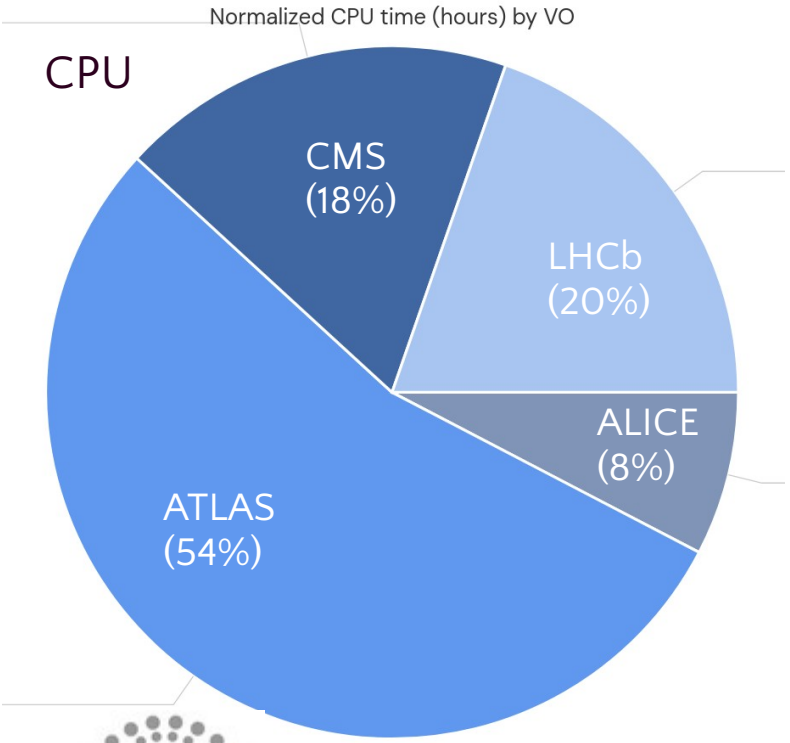

June 2023 – June 2025

LHC VOs evolution in GRIF

- Inside GRIF, constant increase of ATLAS and CMS
- Significant drop of ALICE after 2021 (linked to group sizes)
  - Which translates into GRIF in total being roughly constant after 2021 (see slide before)
- LHCb remaining roughly constant during the recent years

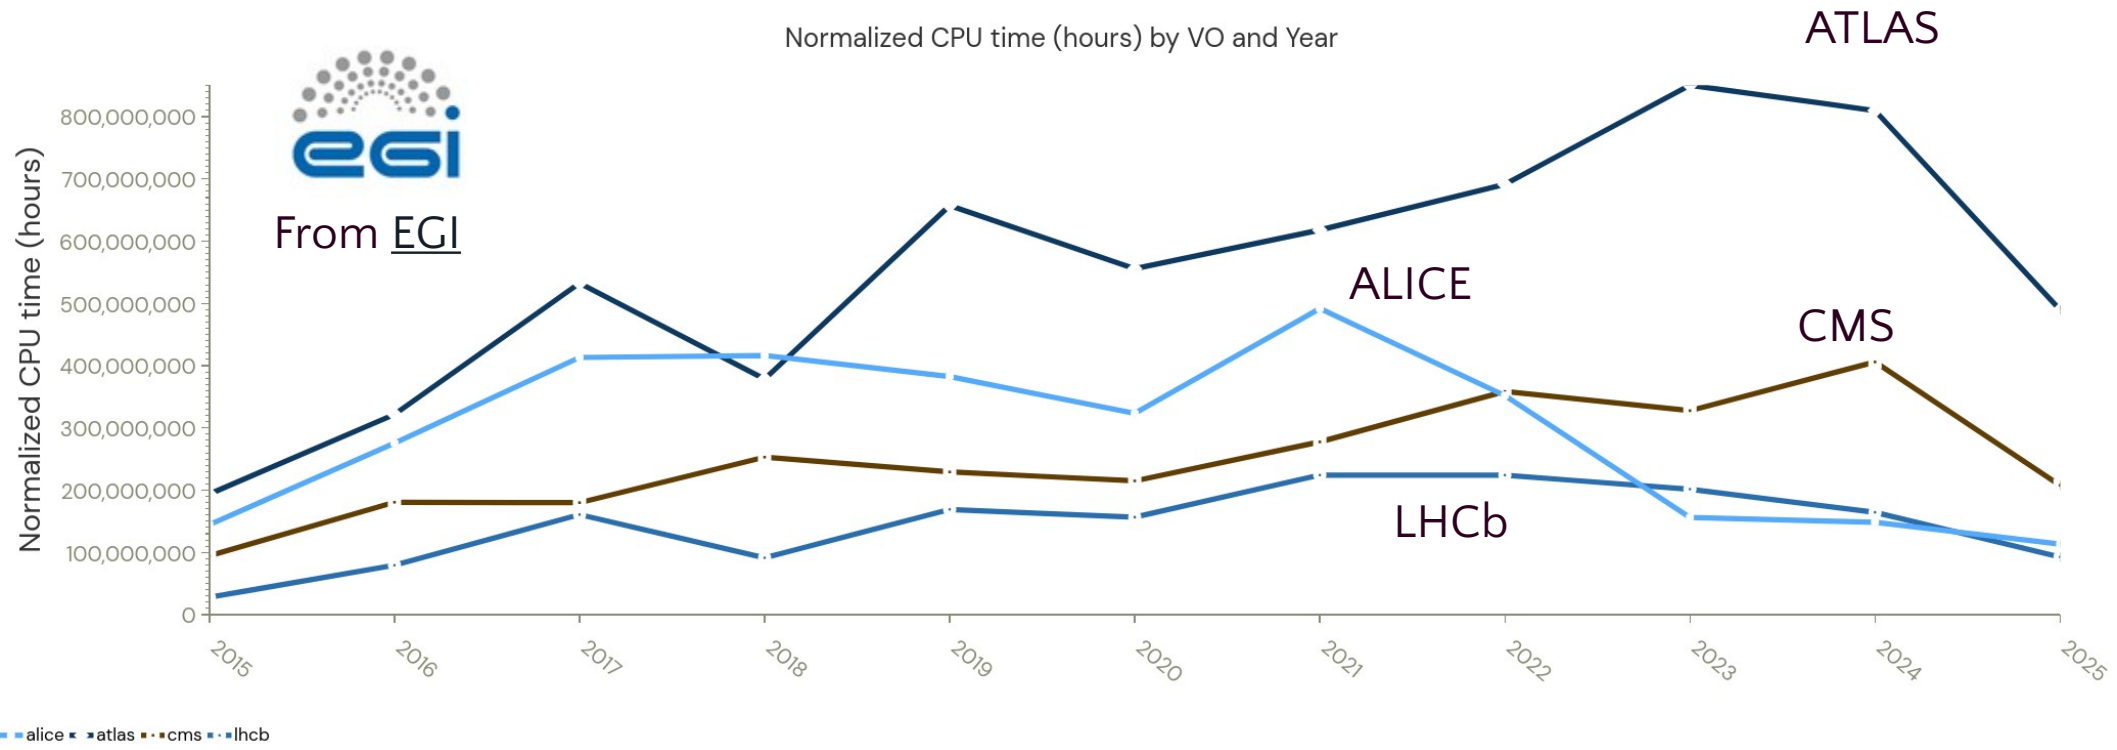

VOs LHC within GRIF

- Current snapshot of the LHC VO sharing within GRIF
  - ATLAS dominating with similar fraction as in France as a whole
  - CMS between 20-30% (with larger disk share)
  - LHCb CPU share larger than disk share
  - ALICE disk share larger than CPU share

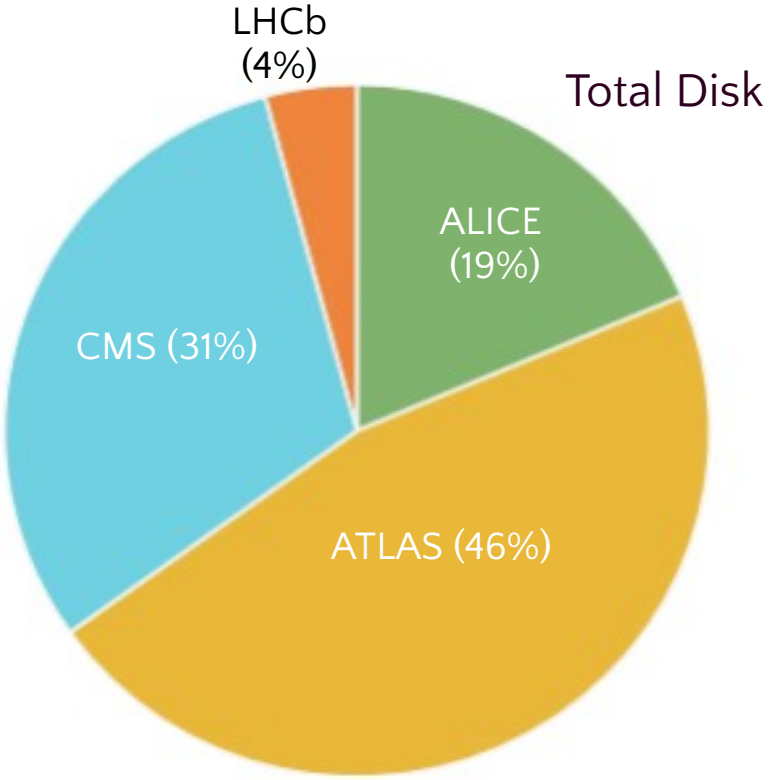

From monit-grafana

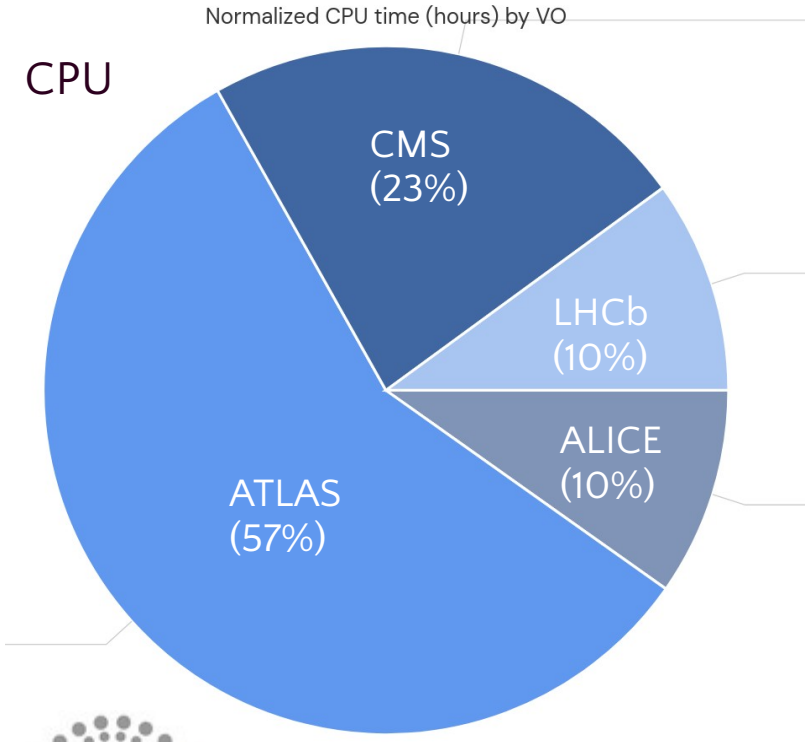

June 2023 – June 2025

Towards HL-LHC

- Revised schedule last Fall
- Run 3 ending in 2026 followed by LS3
- HL-LHC starting taking data in 2030
  - Run 4 2030-2033 → ATLAS and CMS upgrades
  - Run 5 2036-2041 → ALICE and LHCb upgrades (not covered here)

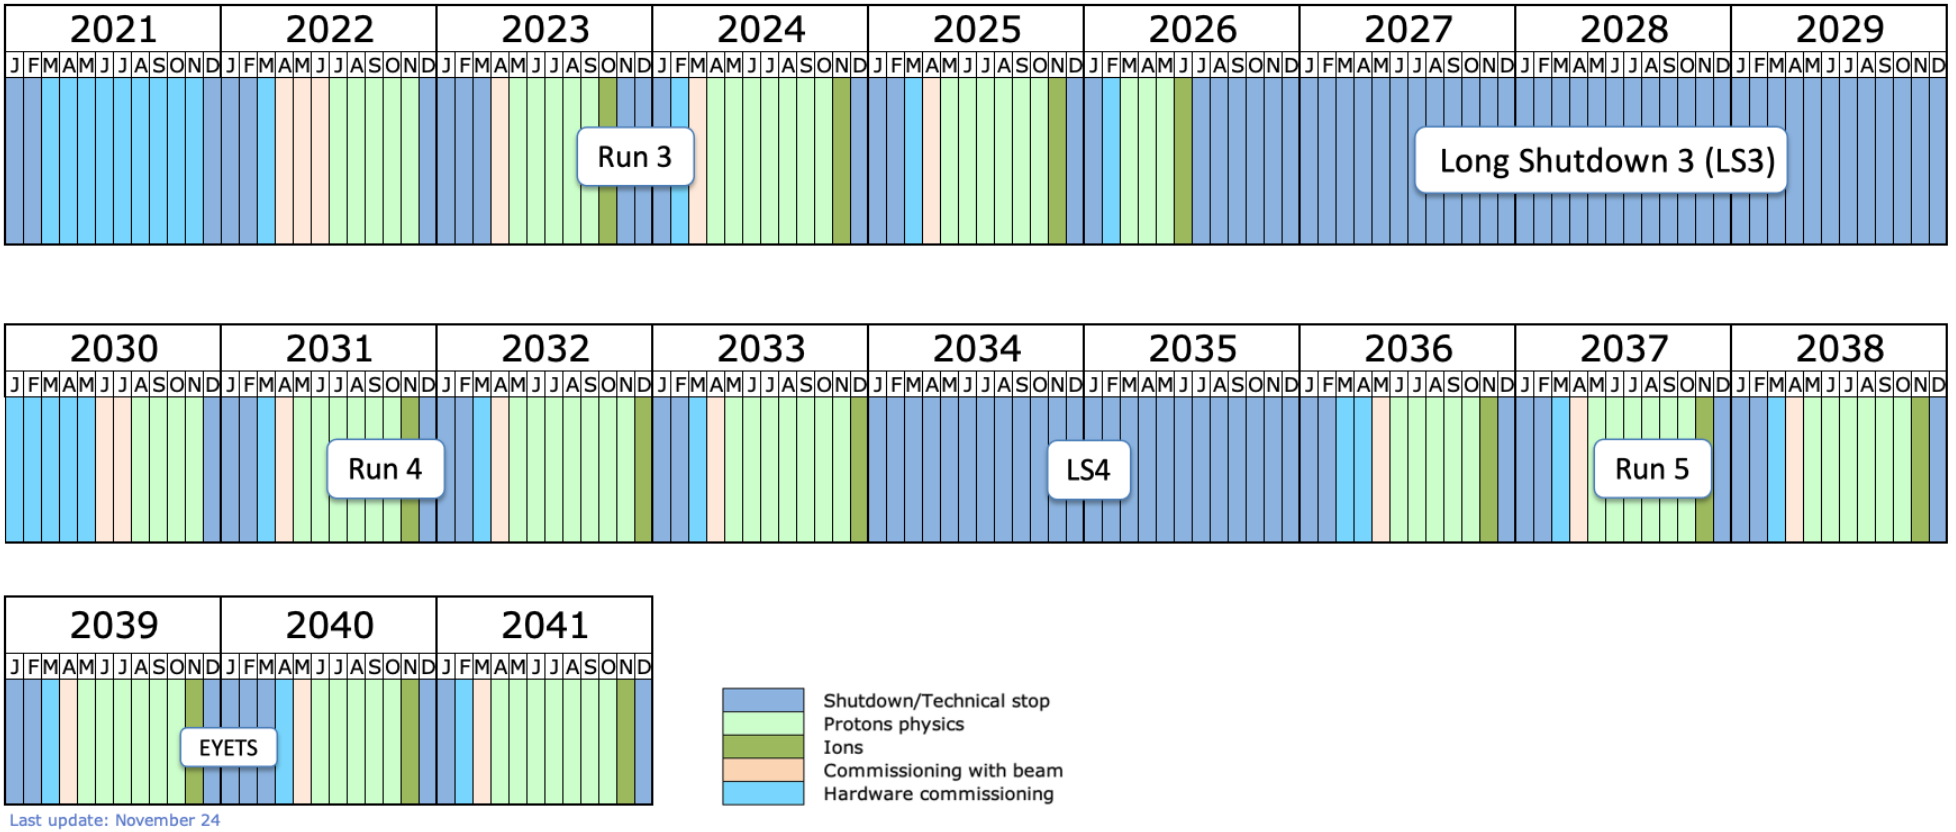

CPU requirement forecast

- Forecast from 2022, before the HL-LHC schedule revision
- Both ATLAS and CMS foresee that the increase in computing needs will not be sustainable without R&D improvement
- With R&D improvements would require a yearly increase of 10-20%
  - Rapid increase at the LHC restart

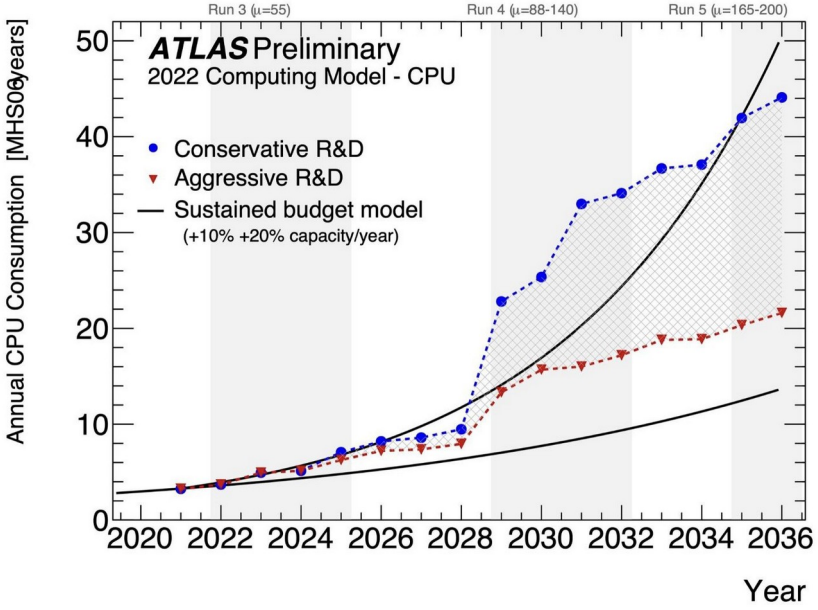

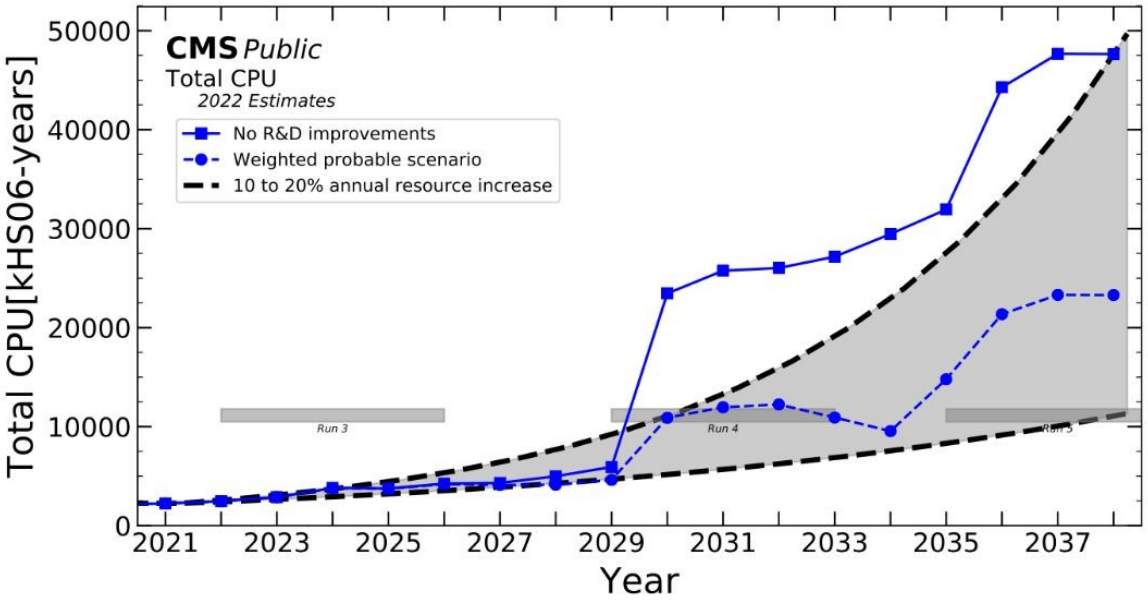

Disk requirement forecast

- Less stringent pressure on disk requirement
  - Although still needs R&D developments (smaller data formats)
- But again foresee a need of 10-20% yearly growth
- The storage limiting factor will actually be the storage on tape
  - Dominated by RAW data size

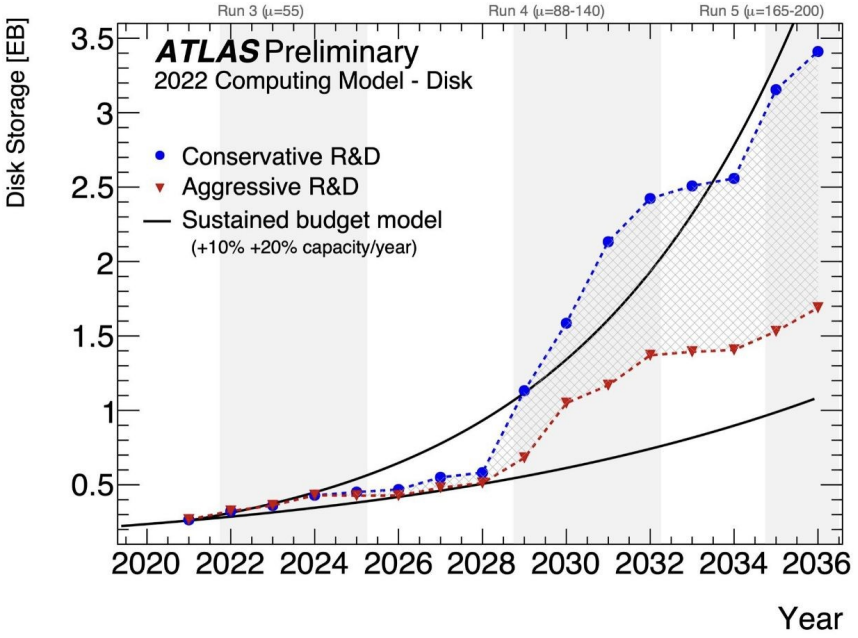

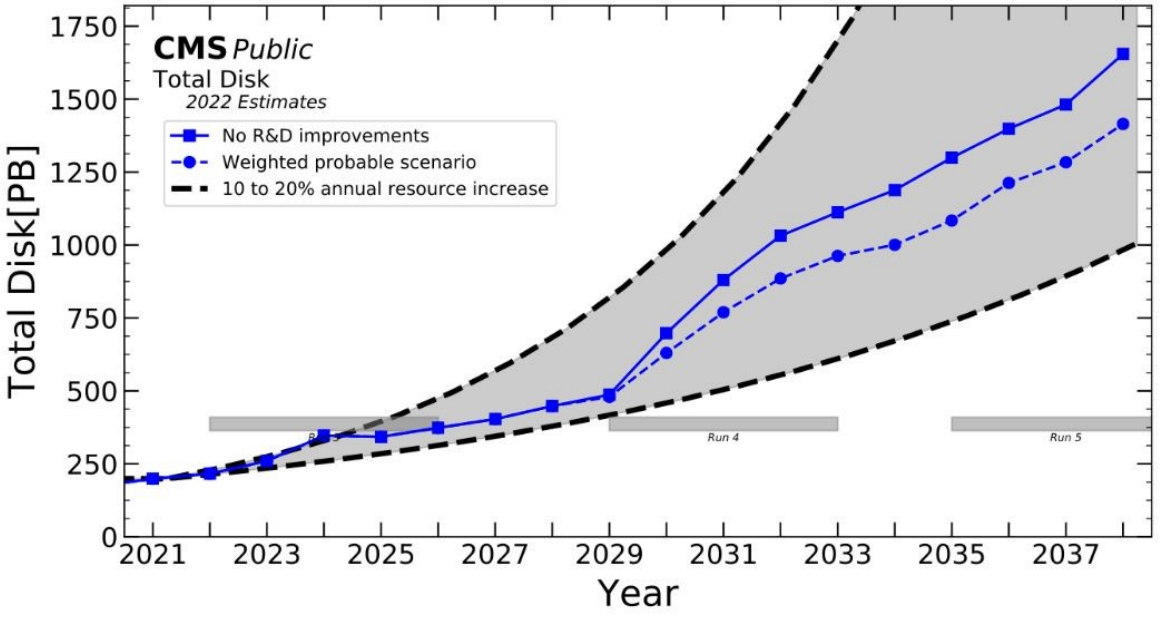

CPU price/performance evolution

- Realistically considering flat budget, what can we gain solely from technology improvement?
- Last 5 years improvement of 28% yearly, but more stagnant recently
- Could possibly expect 10-15% yearly improvement in coming years (source)

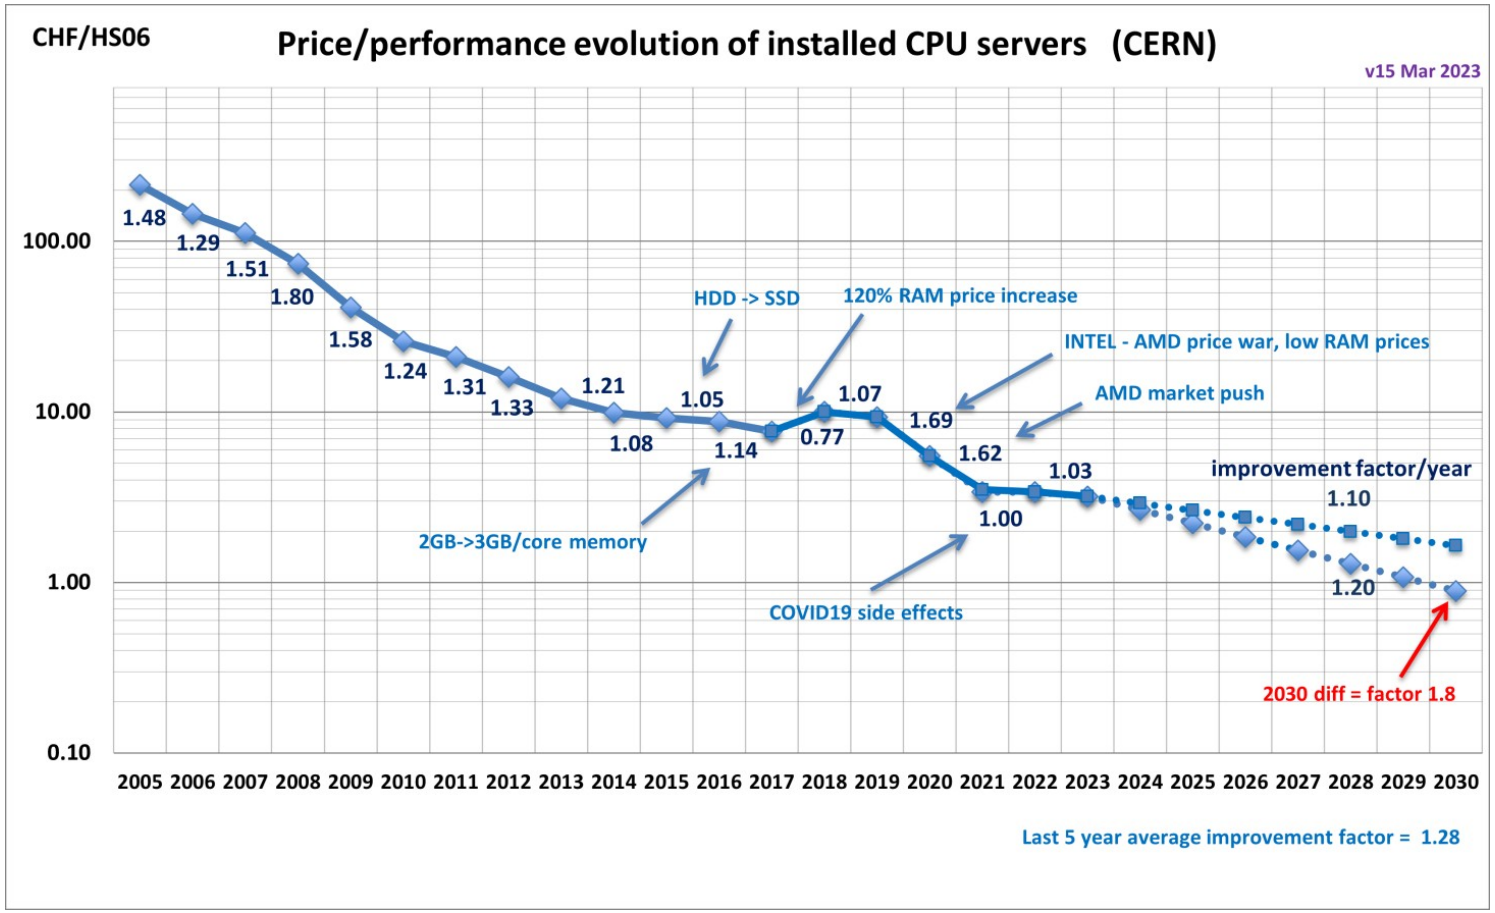

Disk price/performance evolution

- Last 5 years improvement of 9% yearly, but also more more flat recently
- Could possibly expect 0-10% yearly improvement in coming years (source)

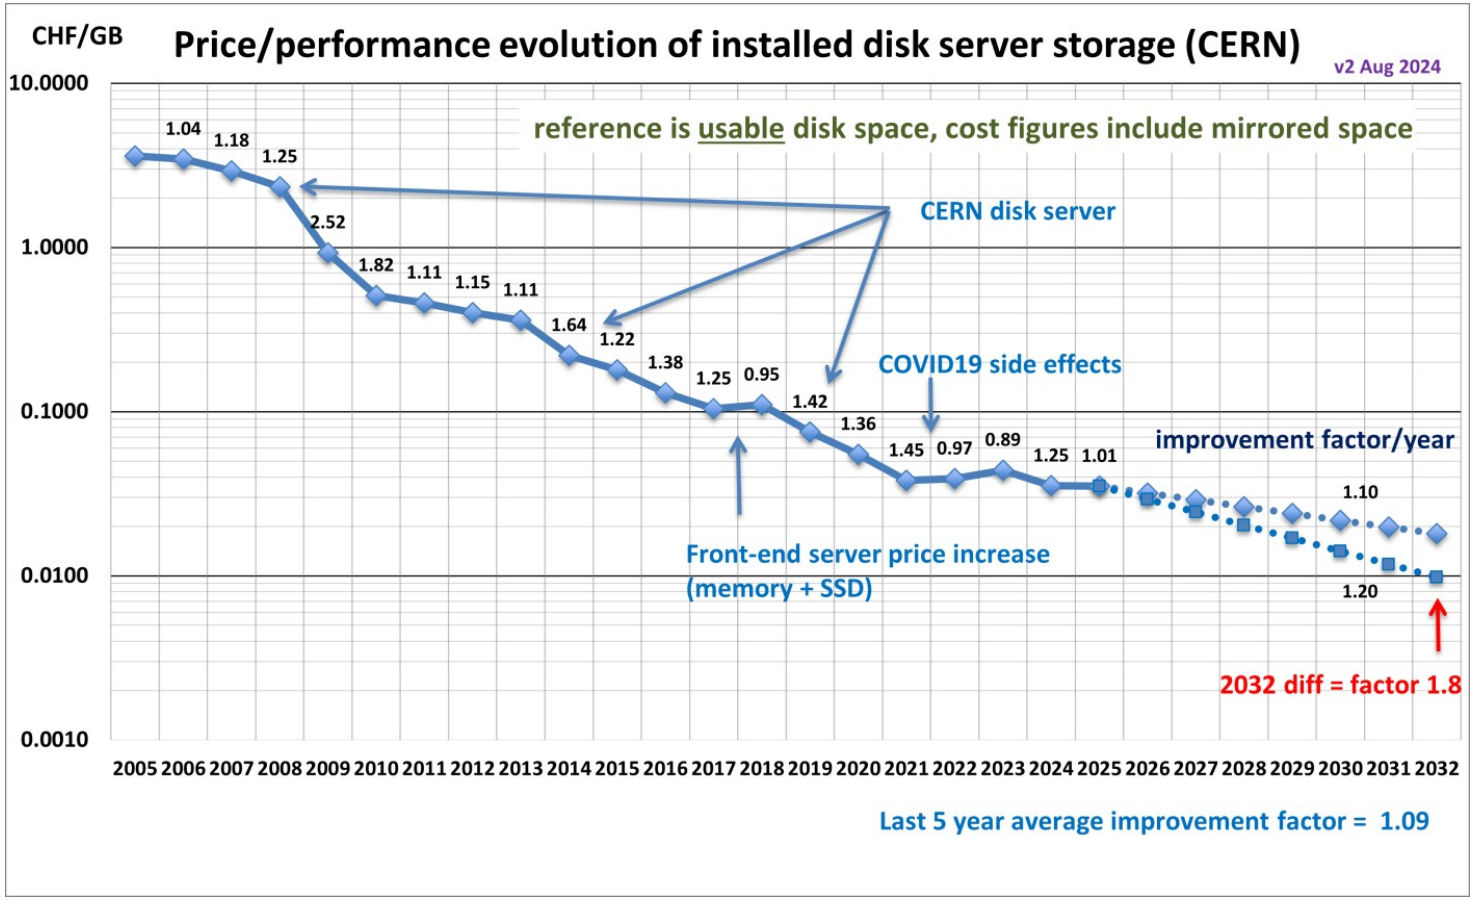

Heterogeneous computing?

- Forecast (done in 2022) of the CPU sharing between the different processing steps
- Simulation and reconstruction are the largest consumers
- This is also where most of the R&D happens
- But a lot is being done on parallelization and running on GPU
  - Heaviest reconstruction (e.g. tracker, HGCAL) ported on GPU
  - Fast simulation kernels with ML in Geant4 (again, GPU acceleration)
- Would we be able to get enough improvements without GPU acceleration?

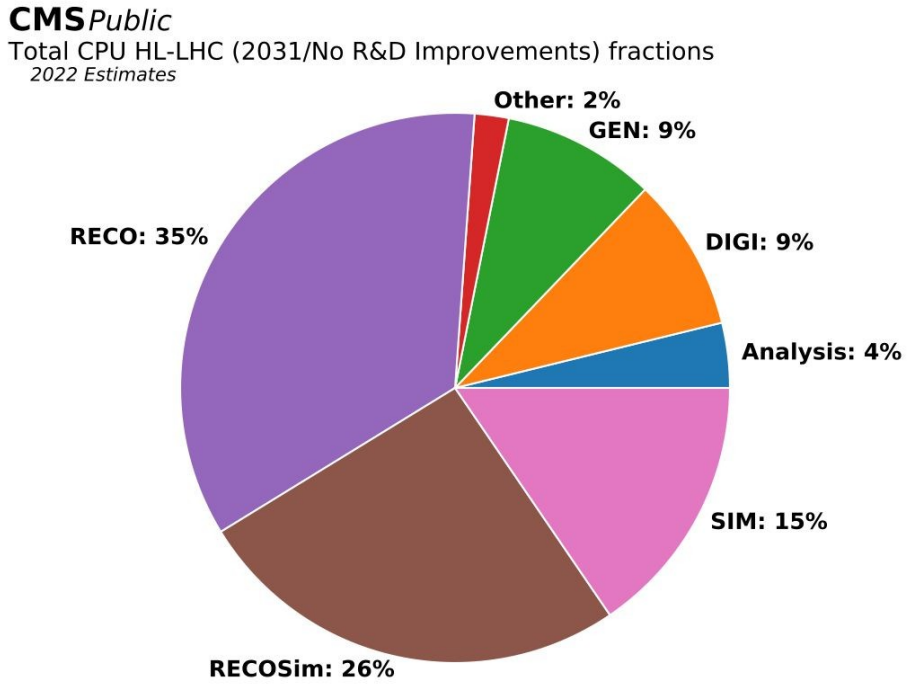

Heterogeneous computing?

- In addition, GPUs are more energy/carbon efficient than CPUs
- Also recent developments in porting and evaluating simulation and reconstruction software to ARM architecture
  - ARM also notoriously more power efficient than x86
  - Although recent x86 processors are reducing the gap
- Would we have heterogeneous T2s in the future?

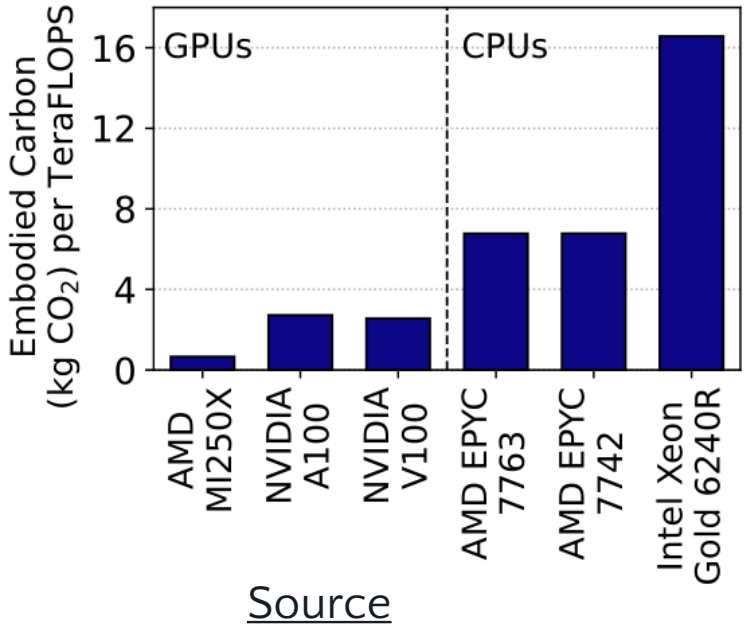

Source

Efficiency for different x86 and ARM CPUs

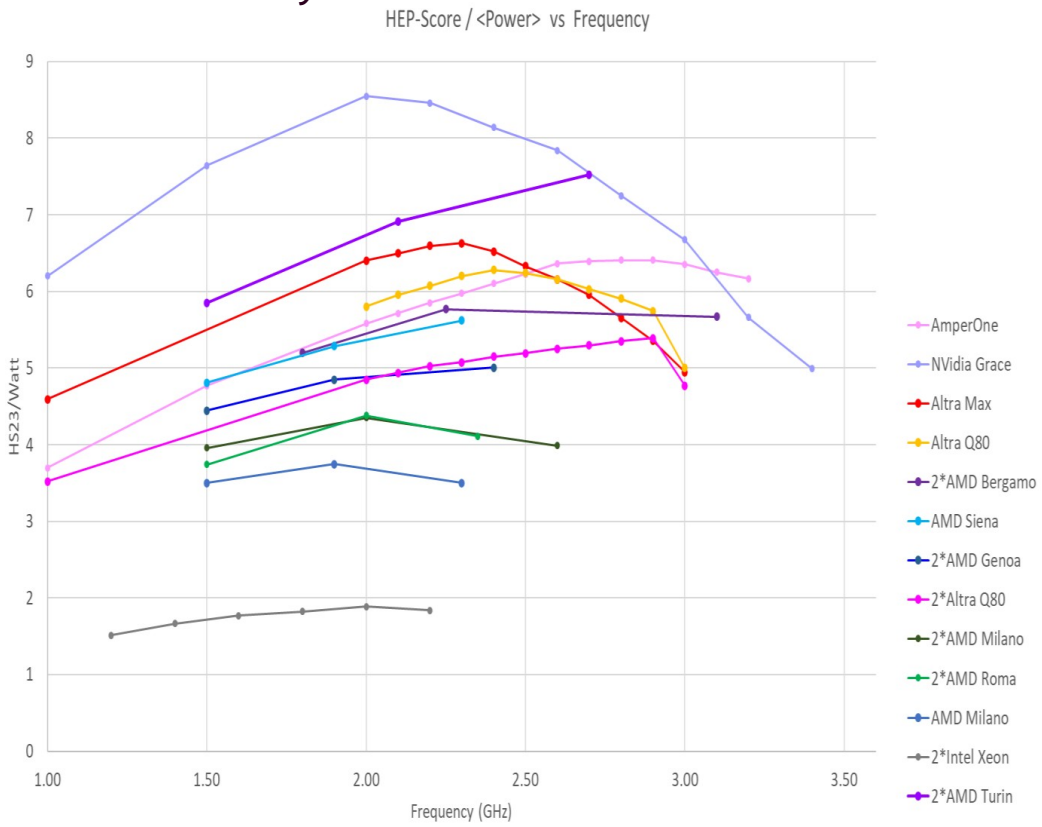

Analysis Facilities

- Set of questions asked by the LHCC to guide the assessment of the analysis infrastructure evolution
- Summary of these questions
  - What are typical Run 3 analysis workflows and how are they run?
  - How will it scale to Run 4&5 and what won't scale?
  - How analysis infrastructures will need to evolve? And will this require new features to expand those of the Grid?
- (some of the) Identified requirements in HSF Analysis Facility White Paper (2024)
  - Fast research iterations
  - Easy switch between interactive and batch
  - Efficient training of ML models
  - Ability to collaborate in multi-organisational team
  - Ability to move analysis to new facilities
- How such infrastructures will look like and where would they be implemented?
  - Something to be integrated in (some) T2s?

Analysis Facilities

- Testing analysis workflows at scale with the Analysis Grand Challenge
- Relying on columnar analysis concept and ML workflows
  - Memory demanding
  - Requiring GPUs
- Need infrastructure beyond the Grid or beyond the Grid specs (in particular in terms of memory)

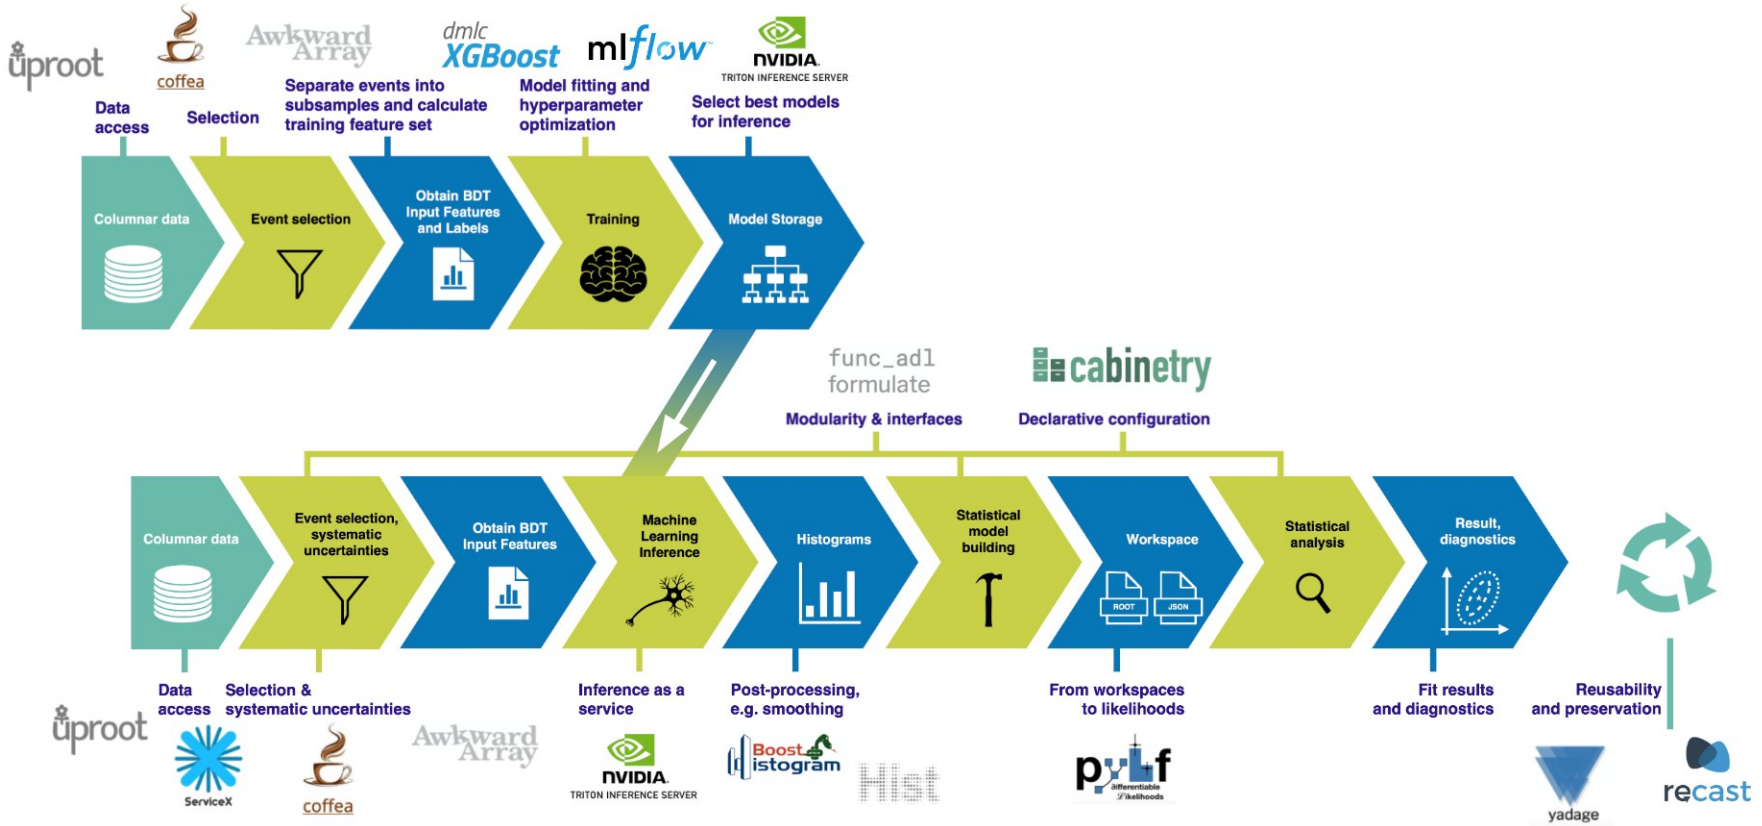

Data challenges

- HL-LHC data will not only increase the requirements in terms of processing and storage
- It will also increase requirements in terms of data transfer and management
- Several data challenges (DC) done and planned before the HL-LHC start
- Goal to exercise the complete transfer and storage chain with increasing data rates (from 10% to 100% of those expected at the HL-LHC)
  - From T0 to T1s
  - But also from T1s to T2s and between T2s

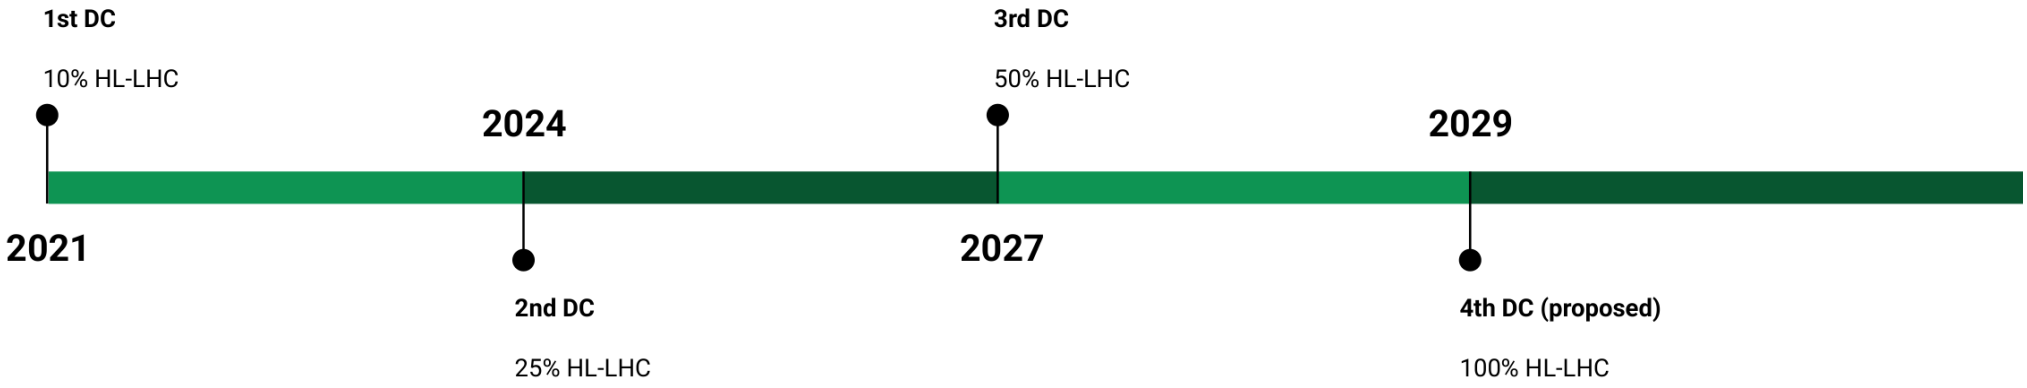

Conclusions

- HL-LHC data will be demanding for the Grid infrastructure
  - In terms of computing, storage an data transfer & management
- Will it be possible to continue doing things as before?
  - i.e. can we simply rely on hardware (e.g. CPU and disk) improvement?
- The next 10 years will probably see significant evolutions
  - But difficult to know in advance the exact shape the computing model will take
- The problem is becoming more multidimensional
  - Questions of power efficiency and carbon footprint in addition to just growing resources every year
  - Heterogeneity of hardware that can be used for simulation and reconstruction
  - Question of expanding the computing infrastructure for analysis use cases. What is the impact on the existing Grid model?