# CMS Directions in CCE2

Lindsey Gray

CCE All Hands Meeting

18 December 2023

# Strategic Directions

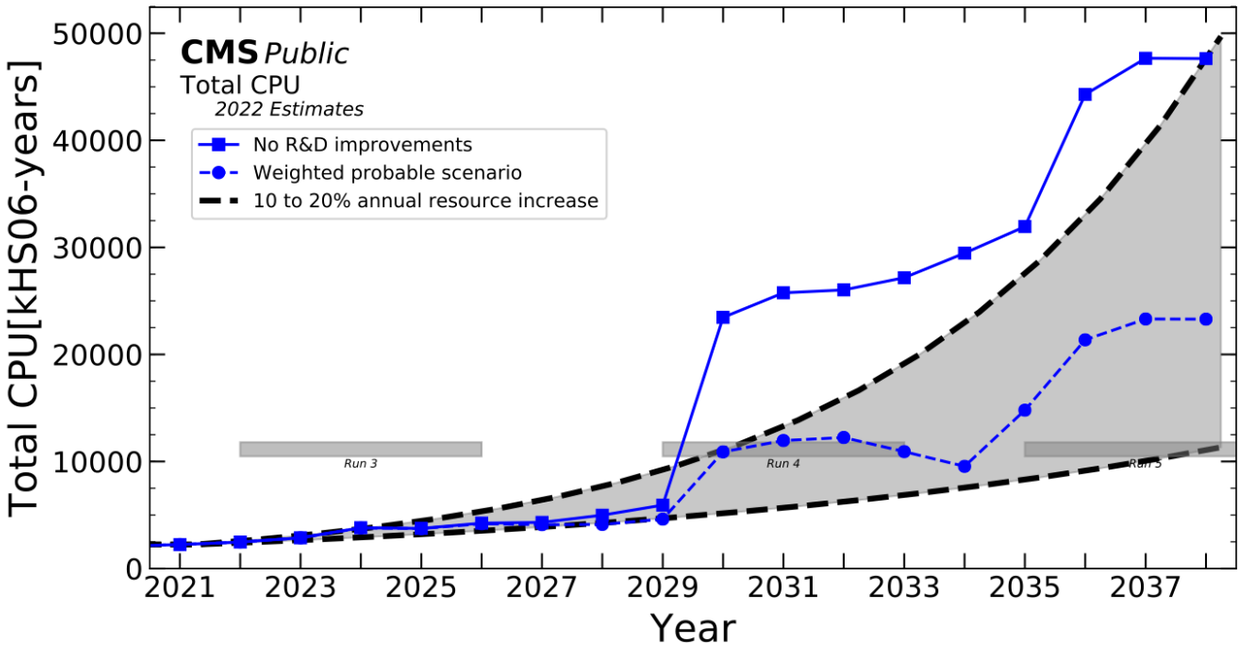


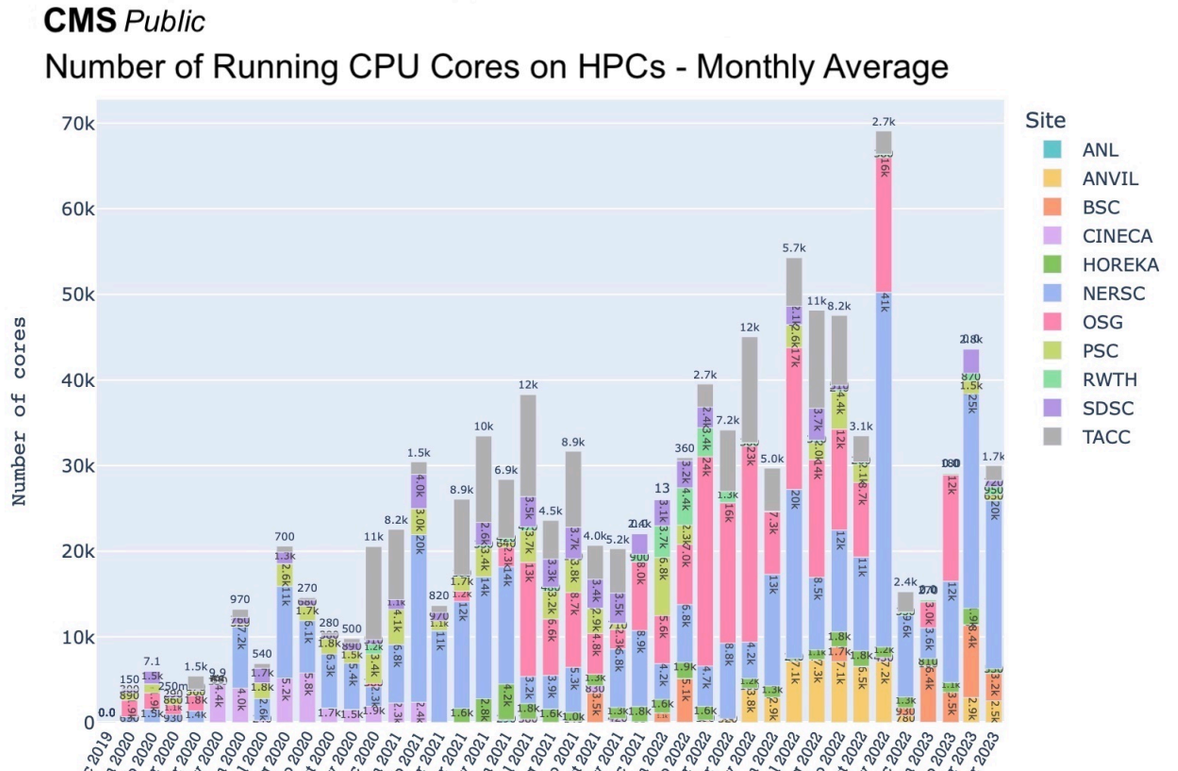


- USCMS S&C, in coordination with iCMS, is investing heavily in heterogeneous architectures and highly scalable storage for a wide variety of compute tasks to meet the challenge of the HL-LHC
  - From large scale MC production and data processing, to analysis-focused use cases
- Usage of HPC resources is key to this strategy
- We intend to achieve this strategy by
  - Easing and further ramping up usage of HPC resources for all compute use cases
  - Integrating heterogeneous resources and workflows into our software
  - Developing distributed, scalable, and efficient storage solutions

# Effectively Utilizing HPCs

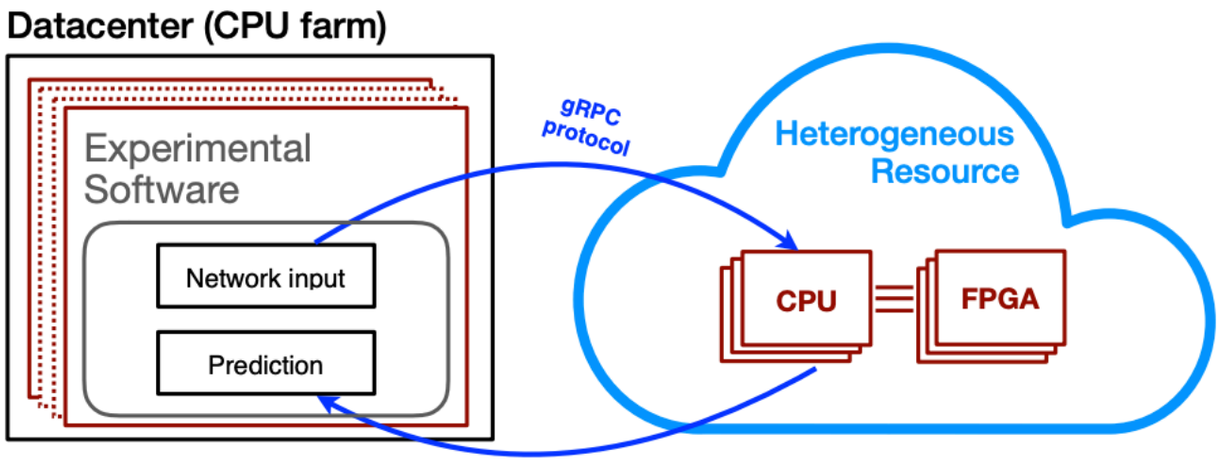


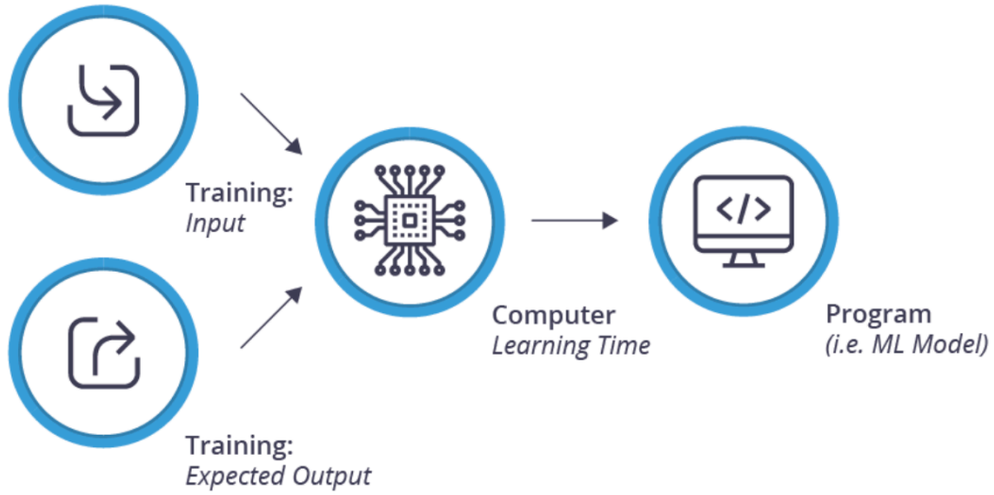


- Would like solidify and expand use of HPCs in CMS computing
- Need to take advantage of heterogenous resources for everything from production jobs to large scale ML trainings
- Pursuing this on five fronts
  - Already have demonstrations of on-host acceleration of reco. algorithms (patatrack)
  - Deploying offloaded inference services (a.k.a. SONIC) on HPC resources (0.25 FTE)
  - Integrating the HPCs through a common API, following the work of HEPCloud (0.25 FTE)
  - Developing a ML training interface that straightforwardly exposes scalable resources to users (0.25 FTE)
  - Data and networking optimization and management from HPCs to experiments (0.5 FTE)
  - Developing Physics-event Generator Workflows that scale on HPCs (0.5 FTE)

# Heterogeneous Workflow Integration

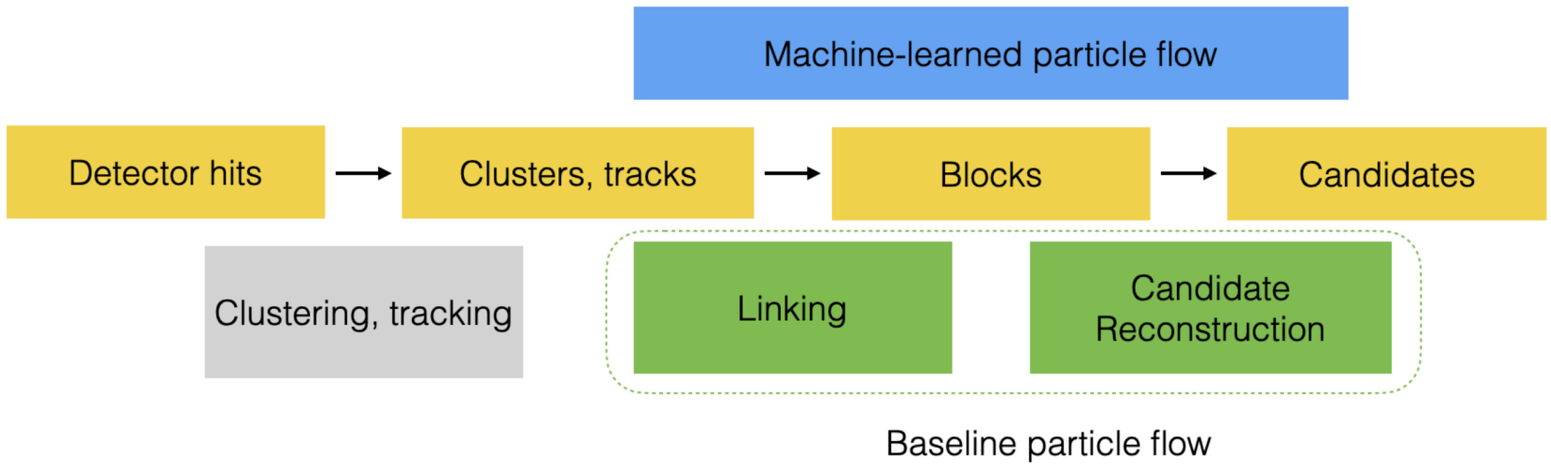


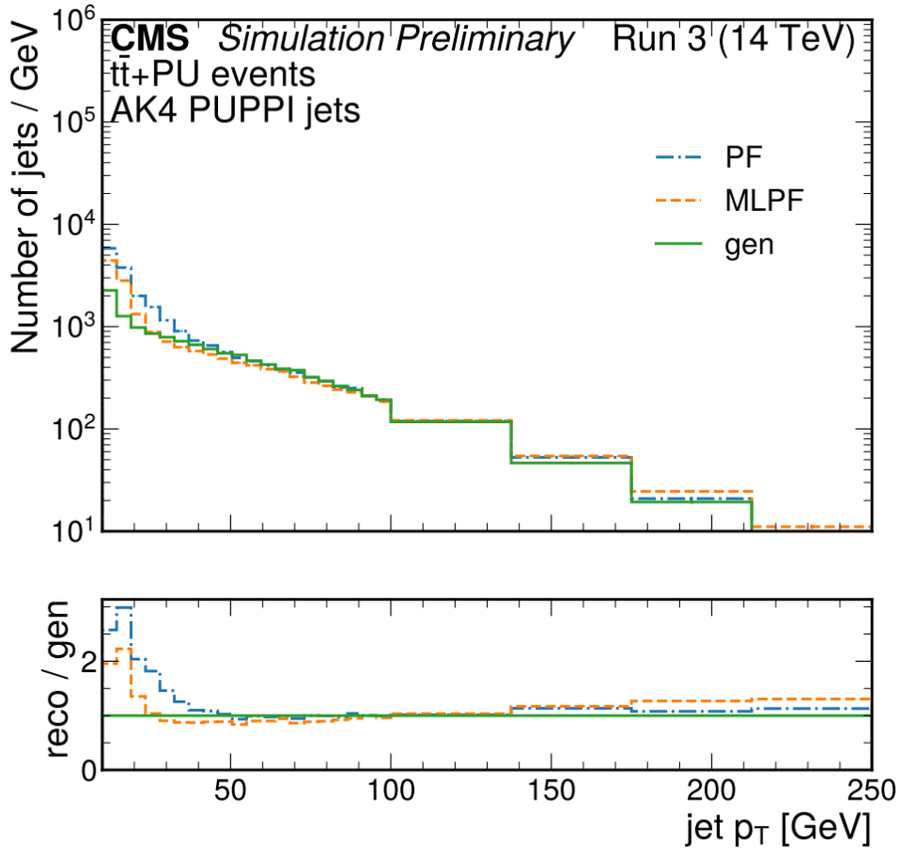


- Using MLPF (https://arxiv.org/pdf/2303.17657.pdf) as a testbed (0.25 FTE)
  - Next generation reconstruction algorithm, based in end-to-end ML techniques
- Excellent choice to exercise heterogeneous resources
  - Significant heterogeneous resource training needs
  - Inference also benefits substantially from running on accelerators
  - Significant performance improvement after hyperparameter optimization
  - Long training times when restricted to using a single GPU
- Clear synergies with lines of work in utilizing HPCs and requires solid integration of those resources with standard CMS software
  - MLPF lets us create a workflow from training to deployment all at scale with clear, physics-driven FOMs

# Object Stores - RNTuple integration

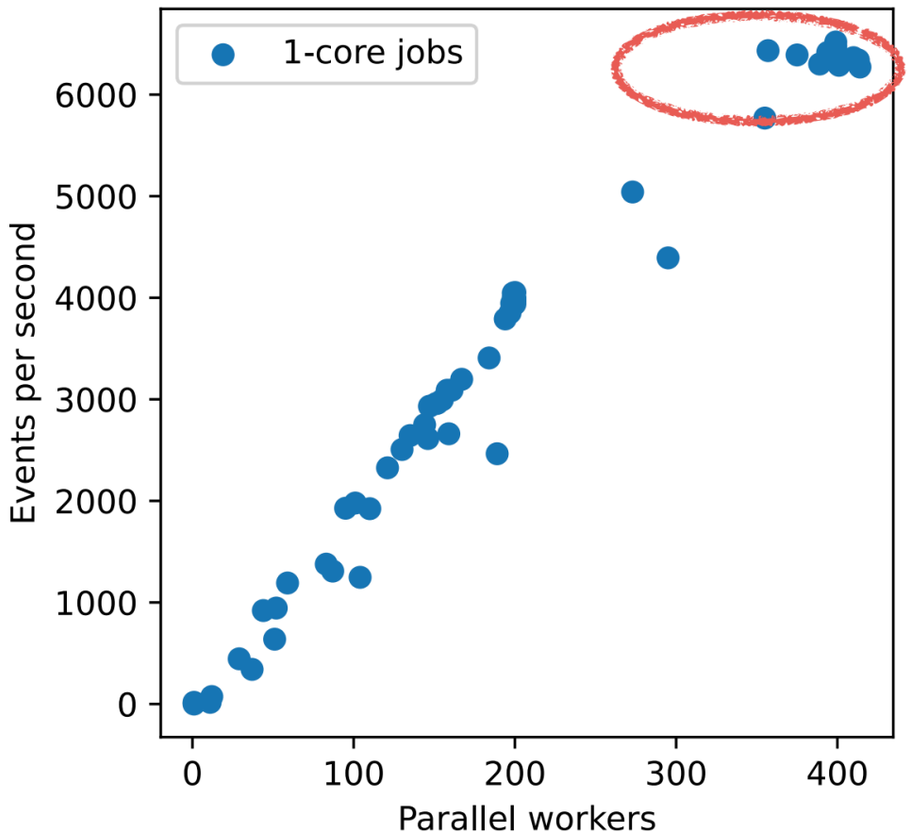


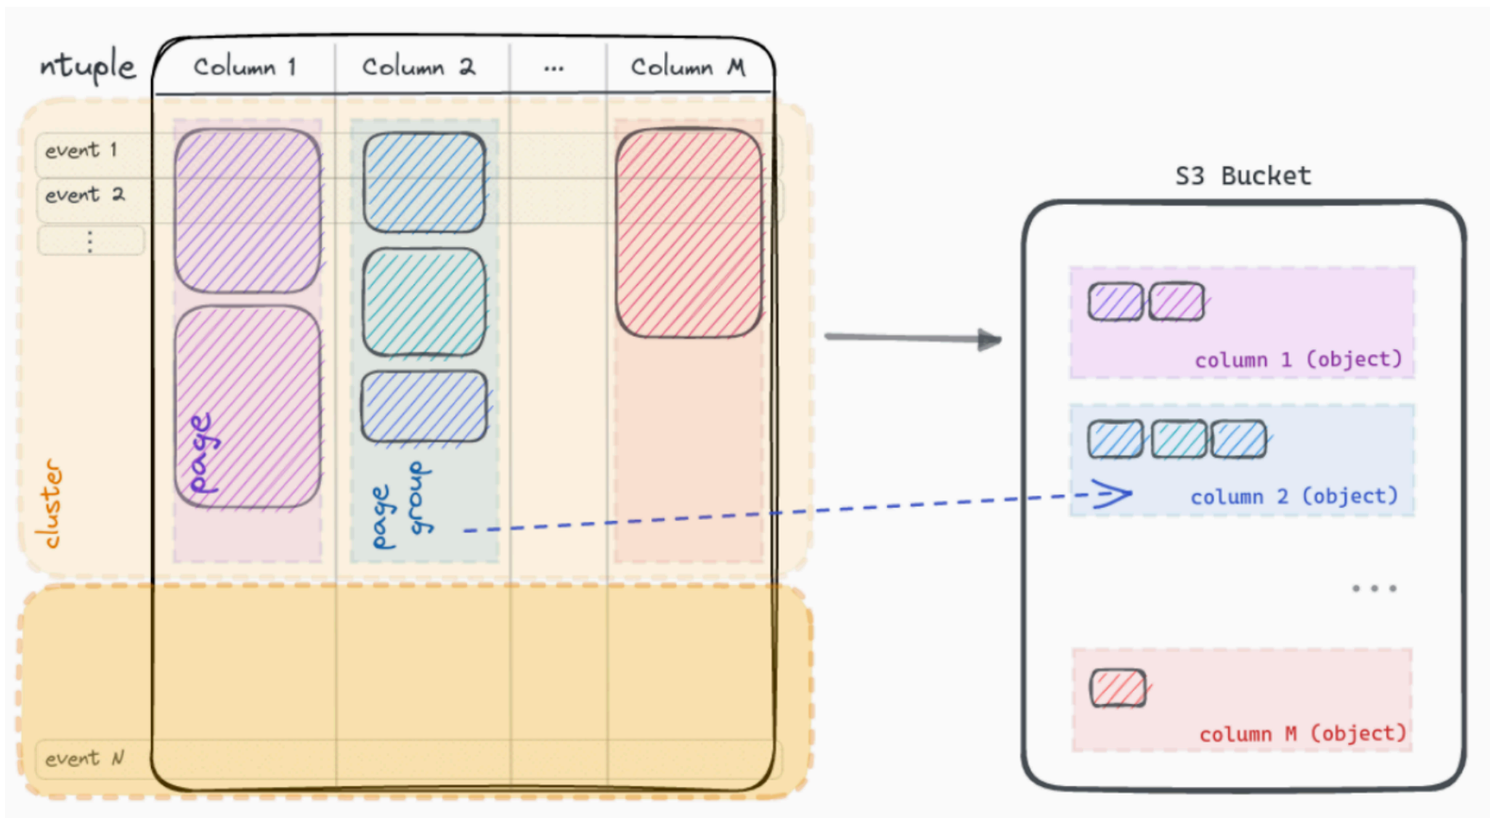


- Investigating object stores as part of radically different data-sharing model
  - Possibility to replace the idea of data tiers with queries and tiered caching structures
  - Initial studies (N. Smith) demonstrate excellent scaling of s3 protocol backed by Ceph object store
- ROOT team expected RNTuple, next-gen ROOT serialization scheme, to be integrated with object stores - requires community investment to achieve required performance
- Expect 0.5 FTE of junior scientist to lead project
  - Possibility for synergy with DUNE in 2025

# Commentary and Concerns on Draft Plan

- Intro
  - We worry that the cookbook may not be very useful for CMS
    - CMS GPU algo. dev. already on par CCE1 PPS testbeds
    - Contributing to cookbook would require working with CMS developer teams directly and negotiating with CMS for their time
  - Patatrack is only true representation of CMS reconstruction performance in test beds
    - Both on-host offload and remote offload are of interest to CMS
    - FastCaloSim and propagate-to-radius testbeds not reflective of actual CMS computing needs
    - However, Patatrack not in CCE2 plans, could add line-segment tracking?
    - Hardware / Software evolution monitoring definitely useful
- PAW
  - Will need Alpaka / CMSSW experts for to make CCE hackathons useful for CMS
  - API and workflow work of clear beneficial use in time
- SOP
  - All RNtuple and object store work is important for CMS
  - R&D on SoA structures < - > RNtuple layouts would be extremely beneficial

# Commentary and Concerns on Draft Plan Cont'd

- SIM
  - Broadly approve of GPU-enabled generator, and Celeritas development
- SMiLe
  - Multi-GPU training of interest, and very much interested in making such workflows easily accessible
  - Concerned by lack of plans for inference as a service (maybe it's in workflows?)

# Conclusions

- CMS has a three front plan to effectively work within CCE2
  - Focuses on the utilization of HPCs and heterogeneous resources
  - Also addressing storage needs in HL-LHC

- The tasks we have chosen suit the needs of CMS and work within the CCE2 mandate
  - There is also significant possibility for crosscutting given the very similar needs of DUNE computing

- Some strategic concerns from CMS on CCE2 goals given past work
  - We can figure out together how to address them