




CMS Directions in CCE2

Lindsey Gray

CCE2 Roadmap Meeting


9 November 2023

Strategic Directions



CMS Public

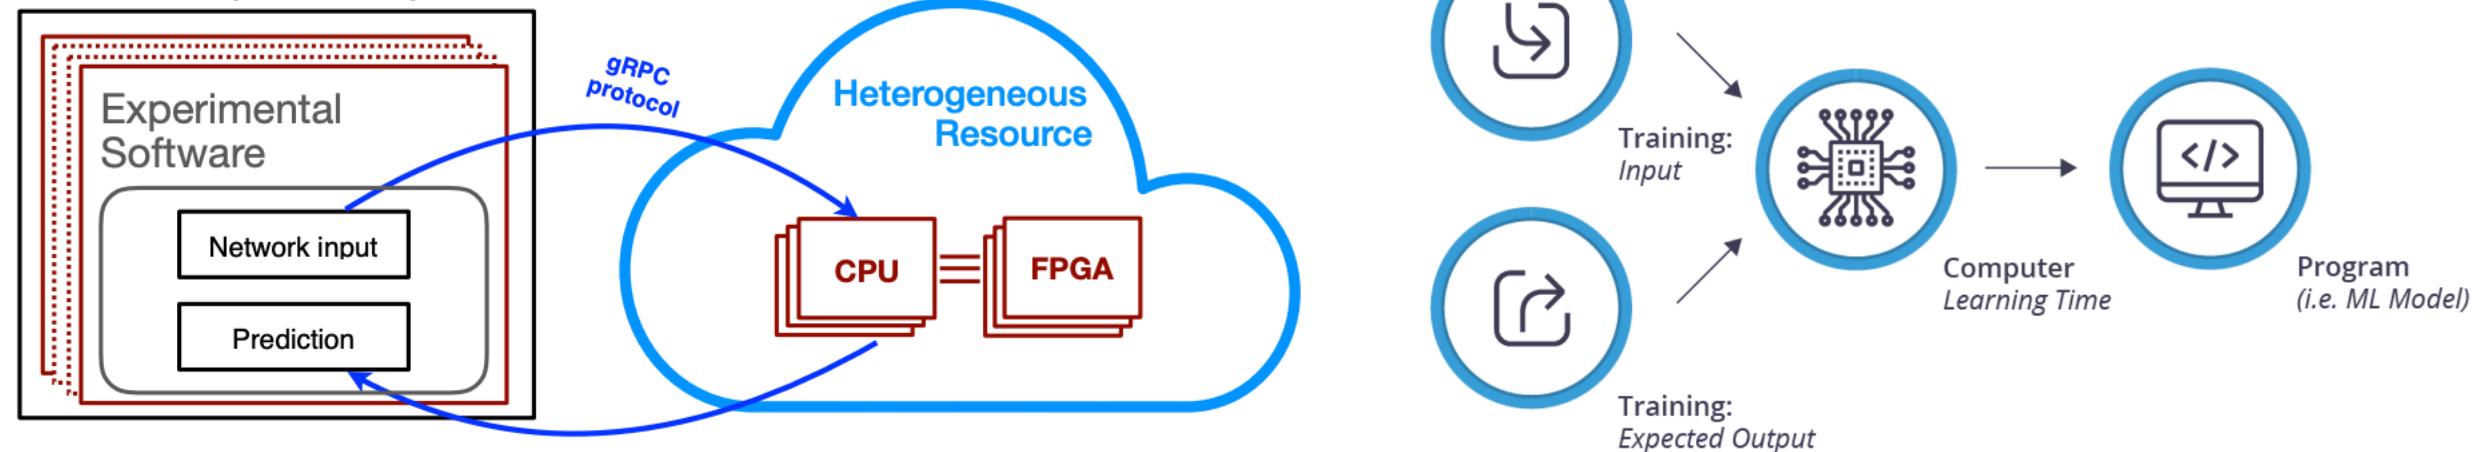
Number of Running CPU Cores on HPCs - Monthly Average



- 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

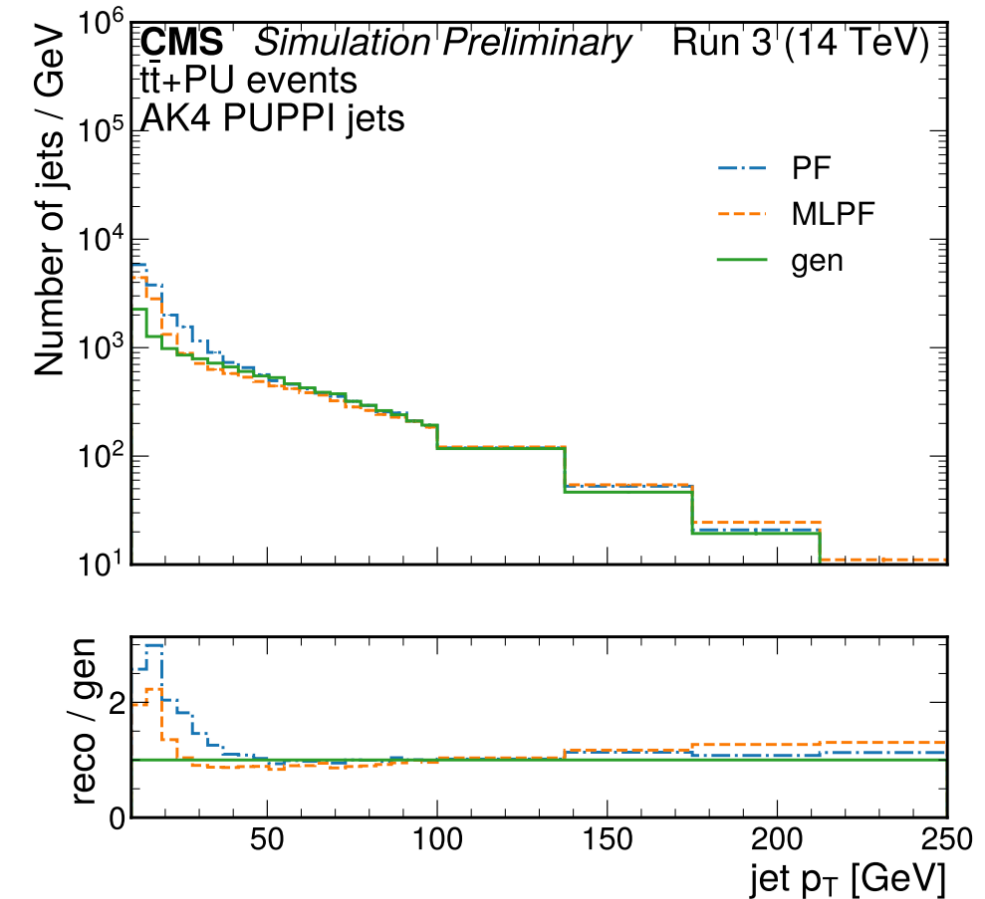
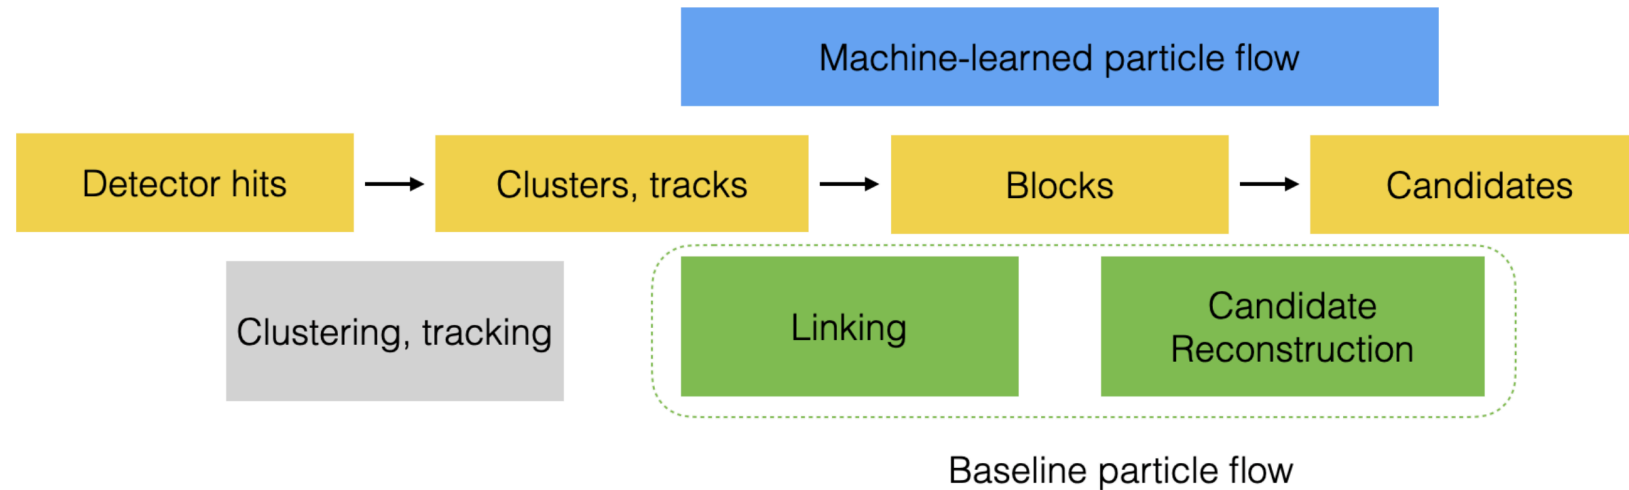
Effectively Utilizing HPCs

Datacenter (CPU farm)



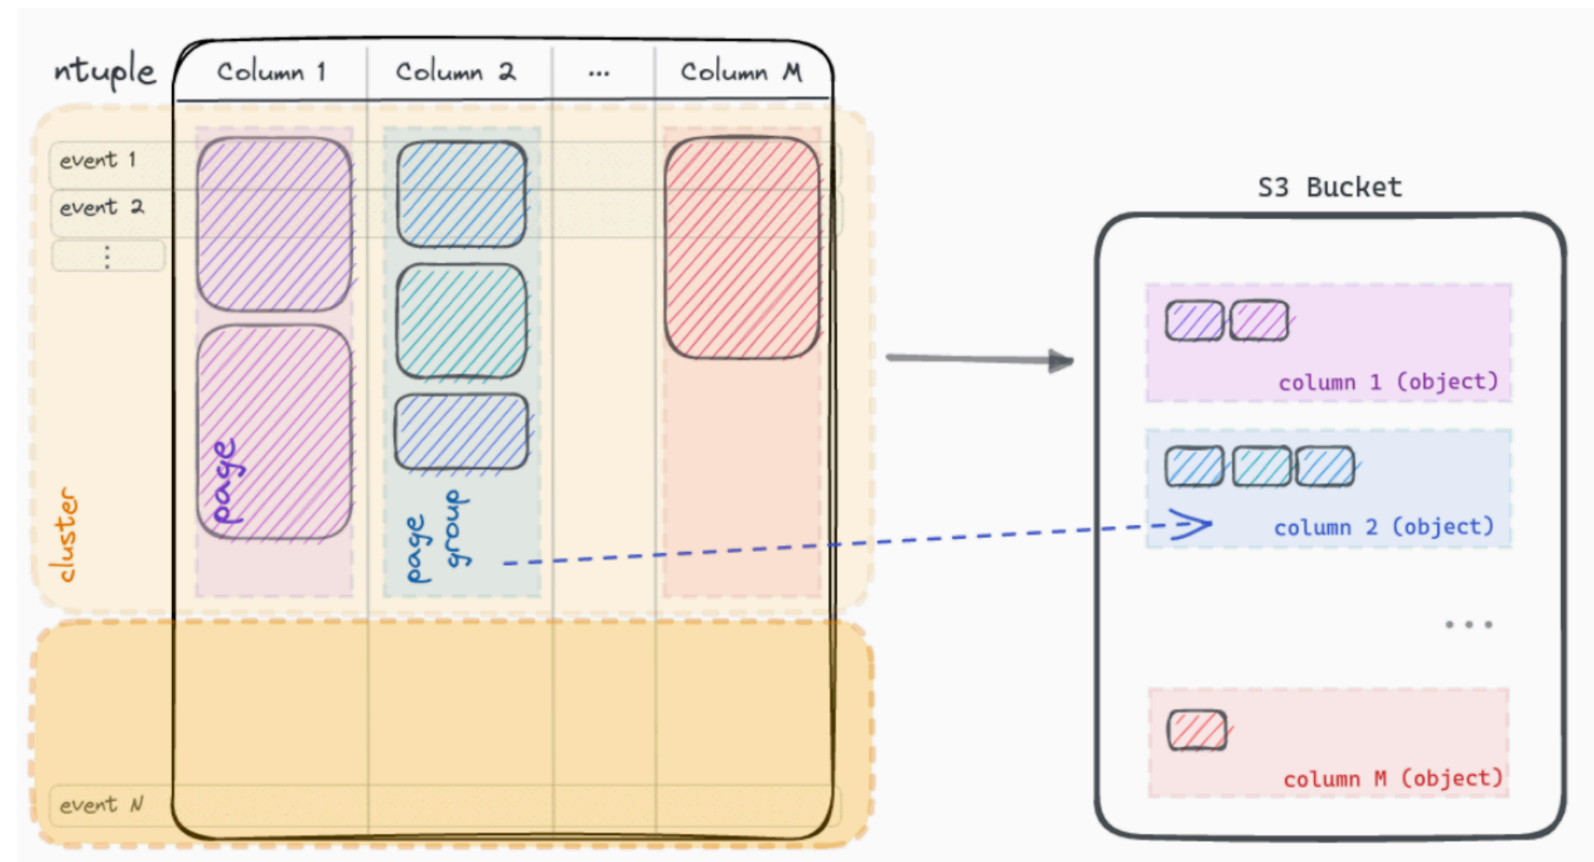
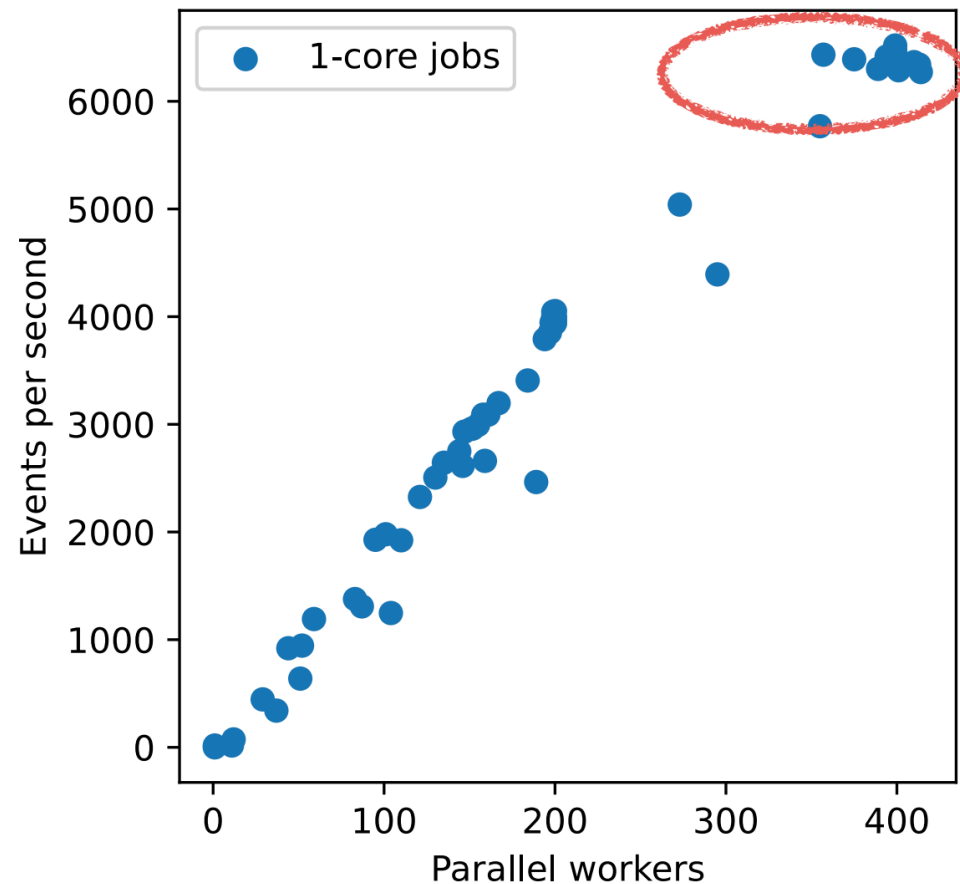
- Would like solidify and expand use of HPCs in CMS computing
- Need to take advantage of heterogeneous resources for everything from production jobs to large scale ML trainings
- Pursuing this on five fronts
 - Deploying offloaded inference services (a.k.a. SONIC) on HPC resources (0.25 FTE)
 - Integrating 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

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