Software and computing at the CMS experiment

56th Annual Fermilab Users Meeting

Mia Liu (Purdue University) June 28 2023

Outline

Brief intro to CMS trigger and data acquisition system

- Role of software and computing in CMS

Status of Current CMS software computing

Computing challenges at the High-Luminosity LHC (Run 4) starting in 2029

- Collider operation conditions, physics needs and changing computing hardware landscape
- R&D Highlights

Connection to other science domains and summary

CMS data collection

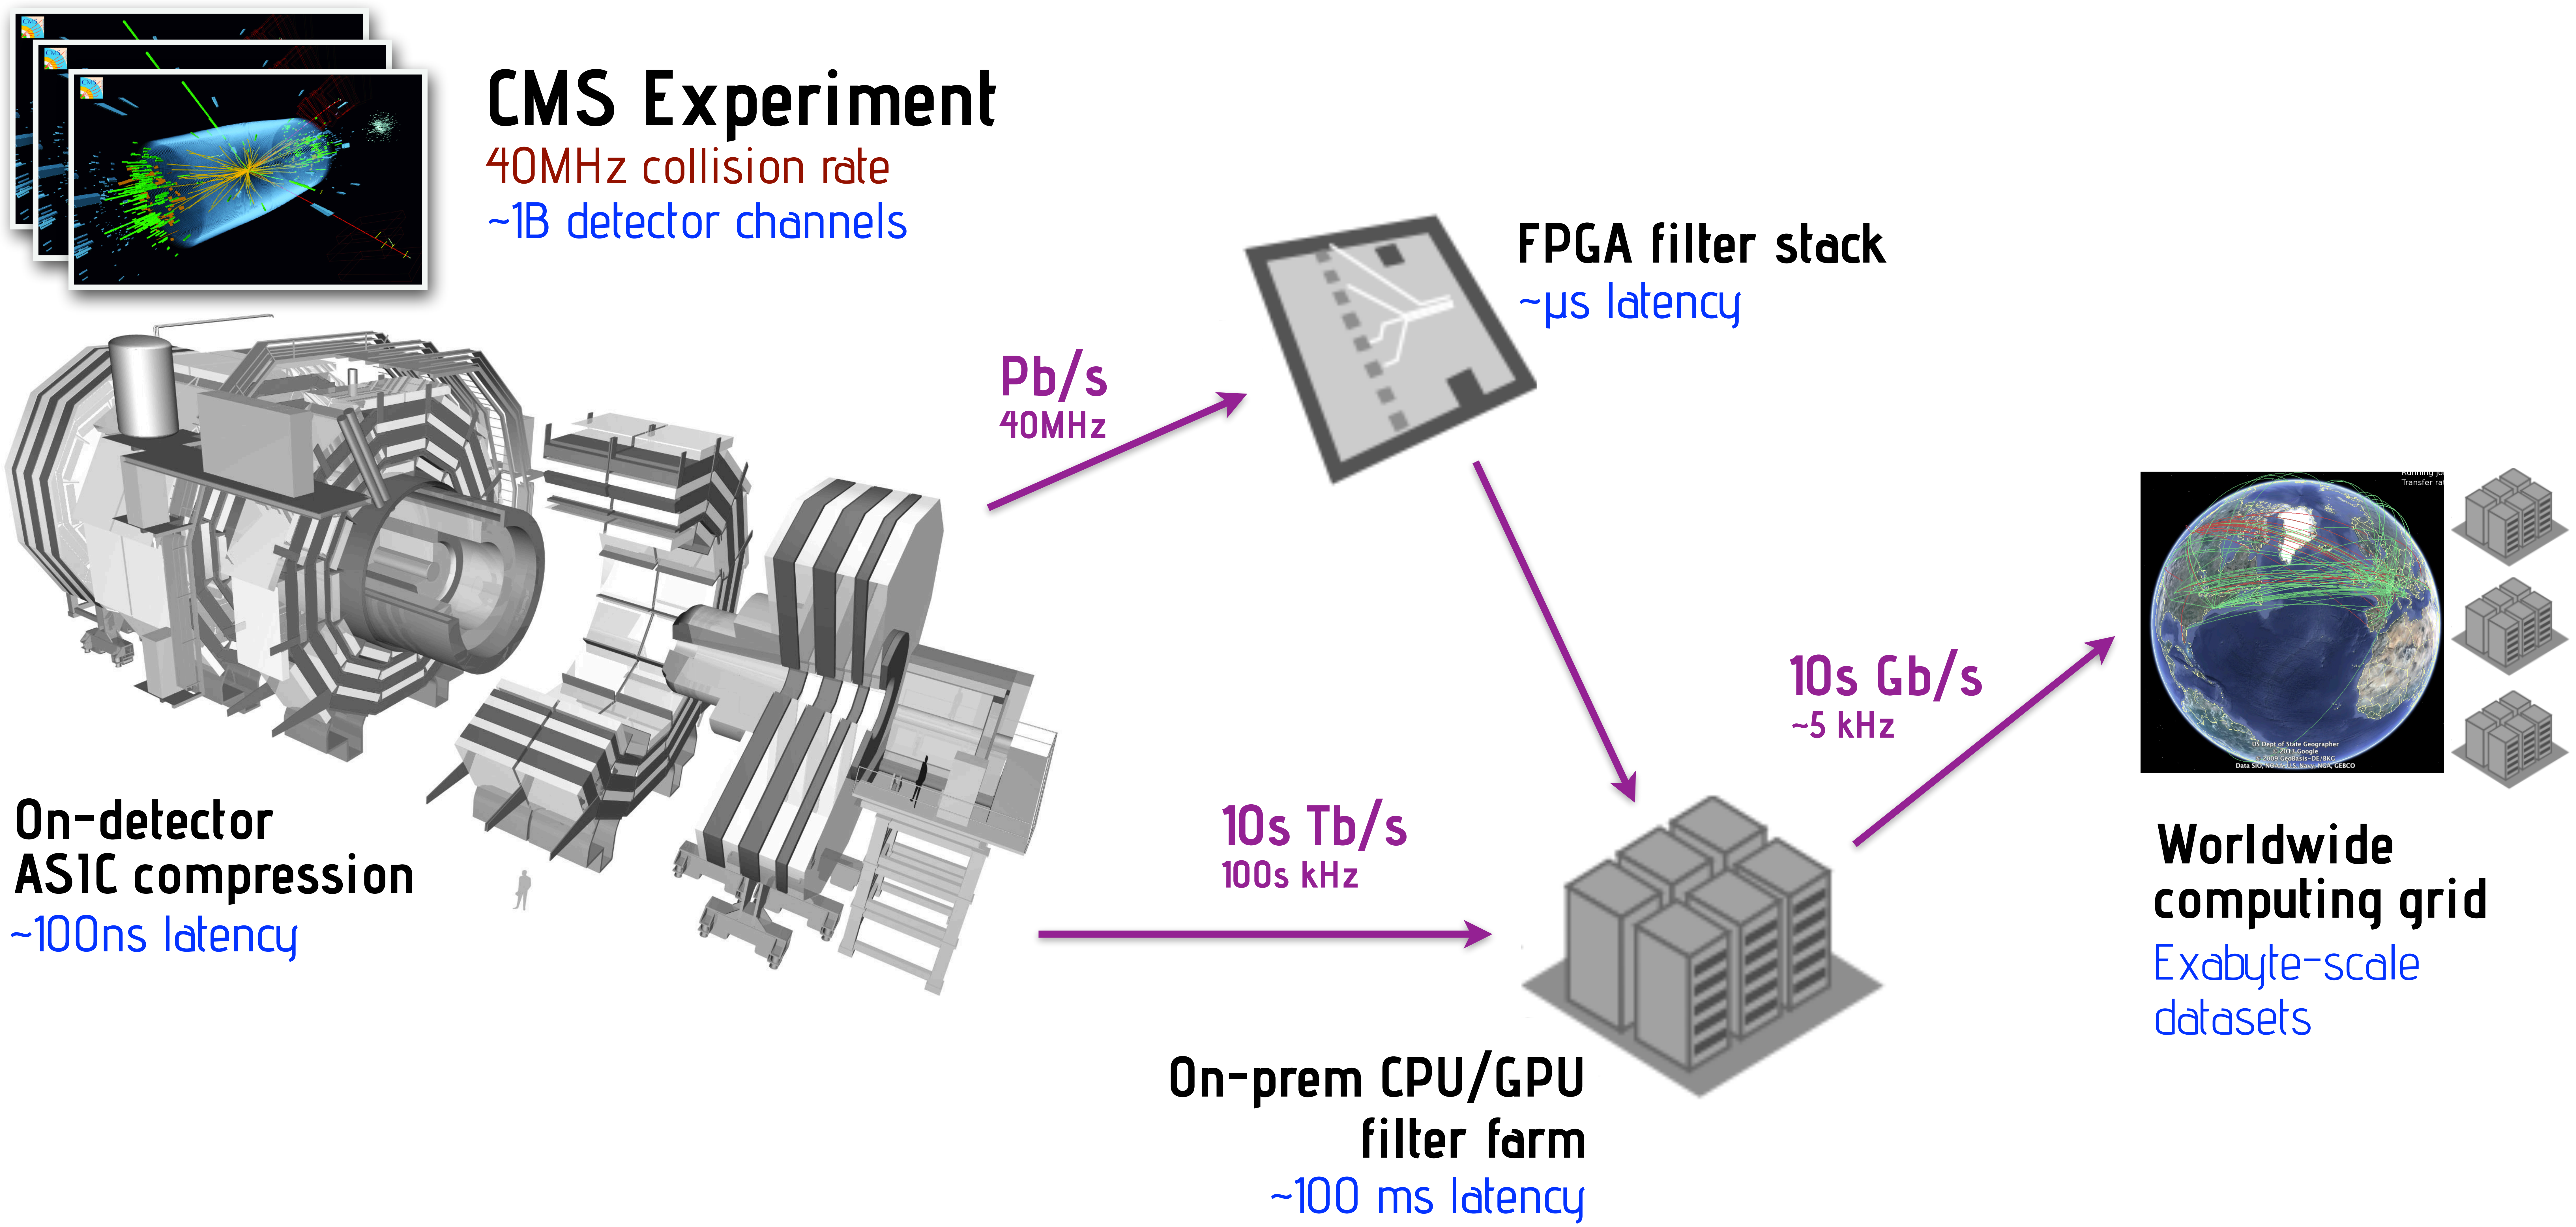

CMS data collection

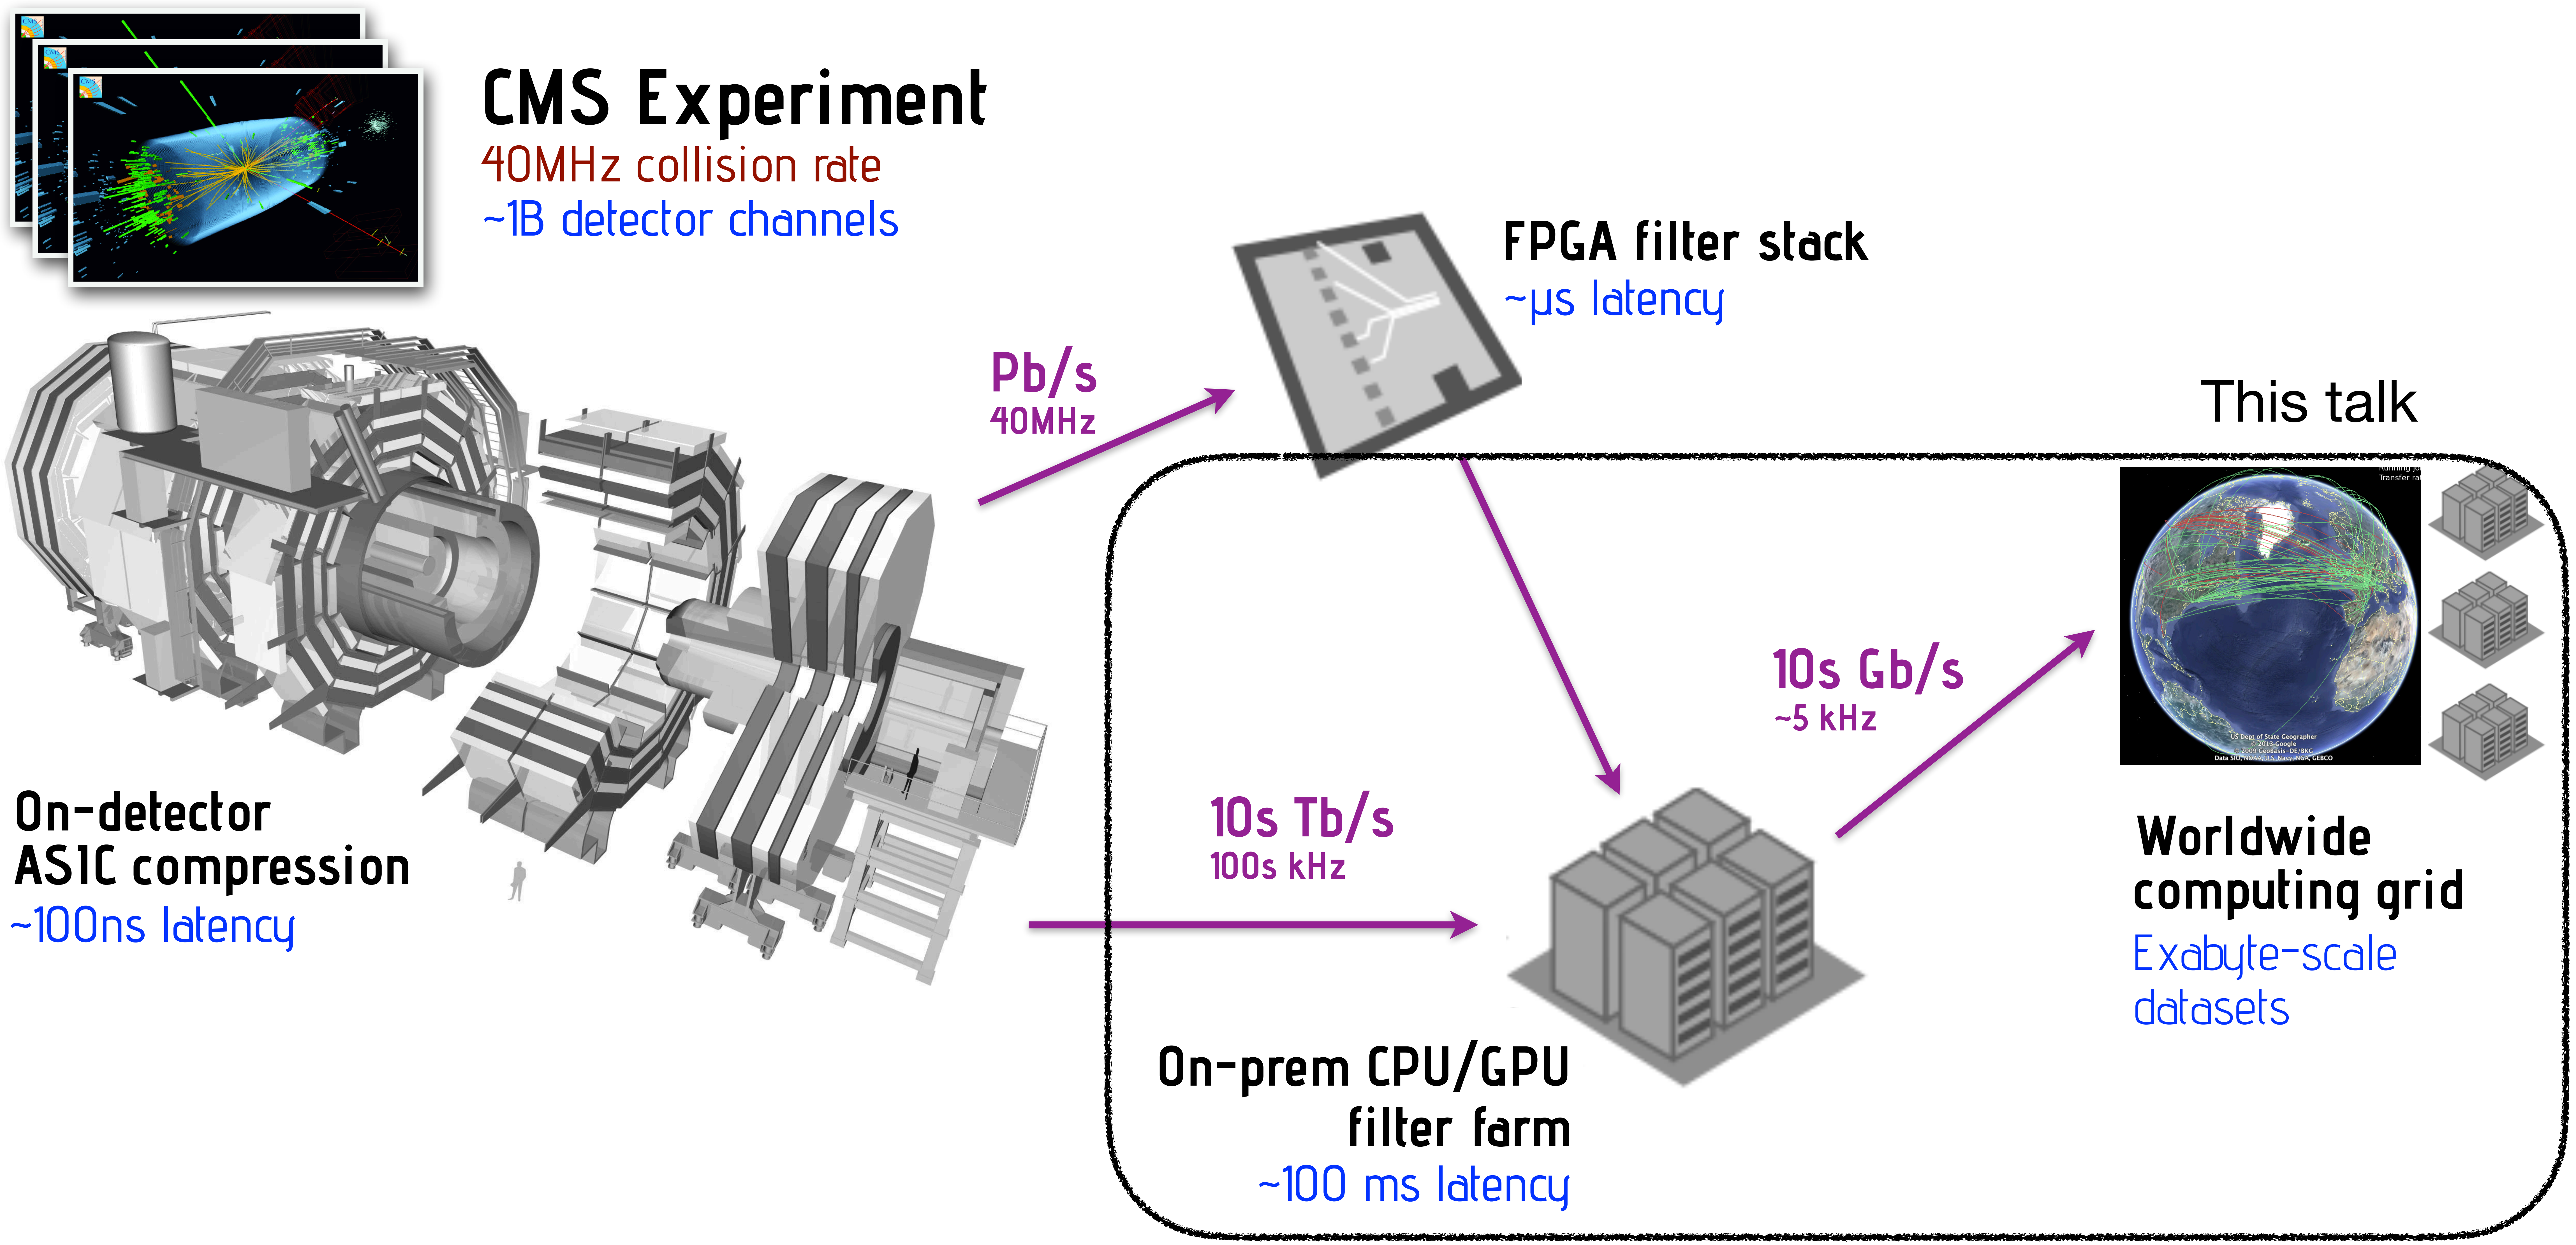

CMS computing and current status

Computing is crucial to CMS physics results

LHC Run 3 data taking is on-going, 2 times luminosity compared to Run 2

Offline distributed high throughput computing infrastructure

1 Tier-1 center at Fermilab. 8 Tier-2 sites in the US

Successfully using HPCs via HEPCloud

~1 B events produced in 2023. 5% of CMS total

Continue to onboard HPC sites: IACS at Stony Brook Ookami (First ARM HPC for CMS)

GPU running at the HLT in Run 3

Will come back to GPUs (and similar) later

CMS running cores on HPC sites (2022/01- Now)

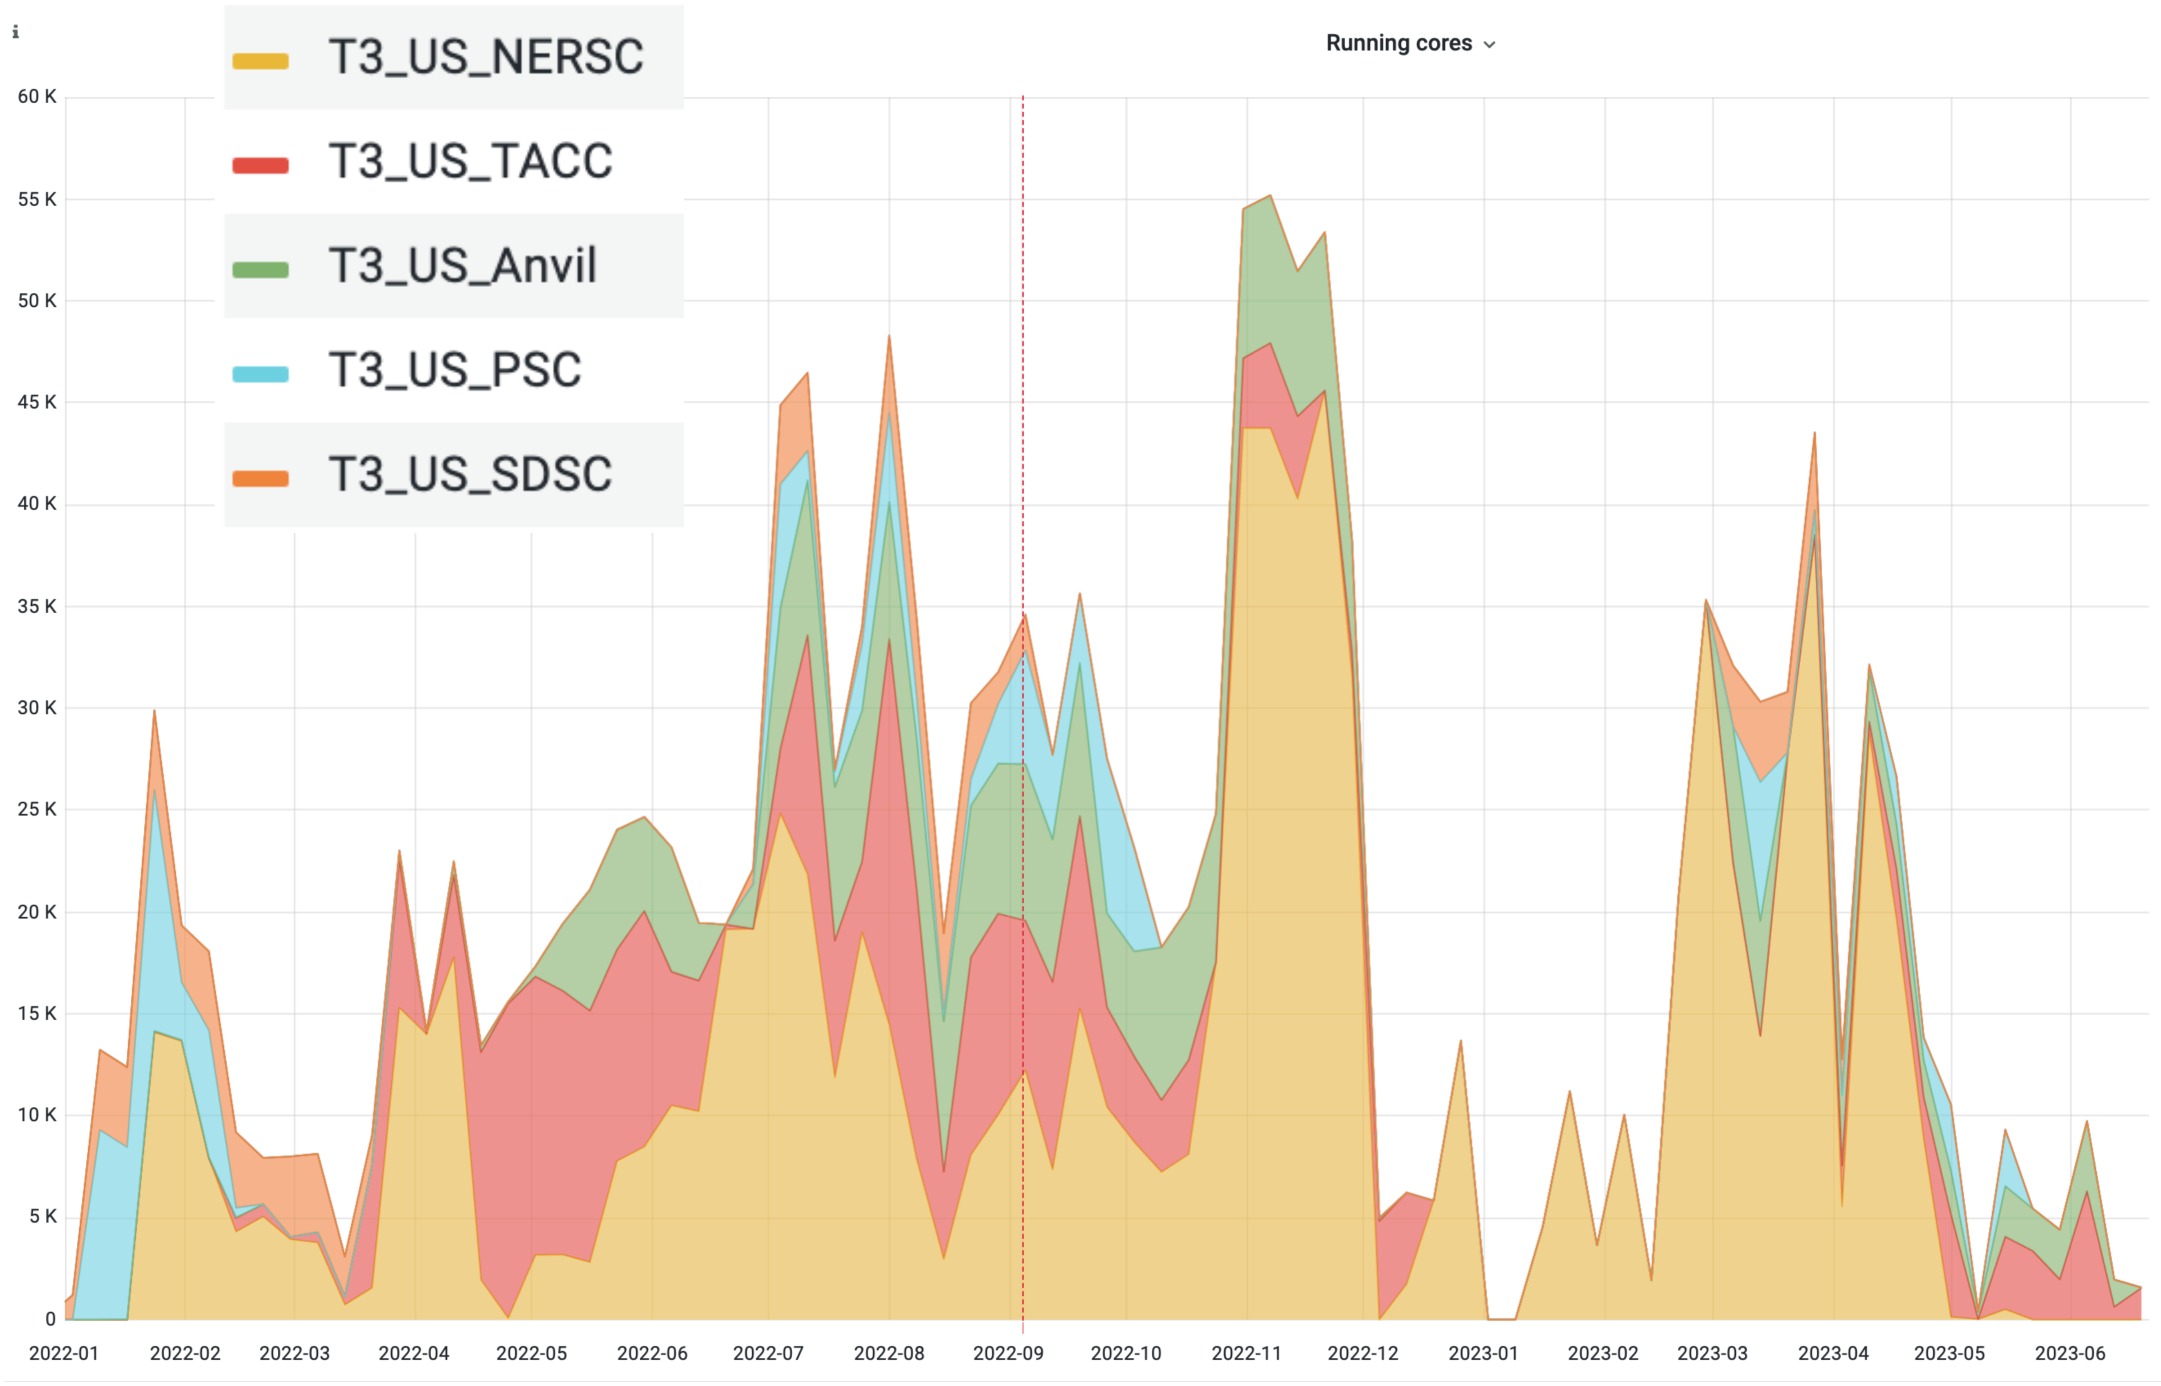

Computing challenges in CMS at the HL-LHC

Data volume: 10 X more integrated luminosity

Data complexity: 10 X Detector channels (O(100M)->O(1B))

Trigger rate increases by 10x (to ~10kHz)

Physics needs

- Advanced sophisticated algorithms
- Alternative readout schemes: B-parking (already in Run 3)/L1-scouting
- Enhanced MC stats needed for precision measurements

Increased needs in storage, processing power, user analysis support

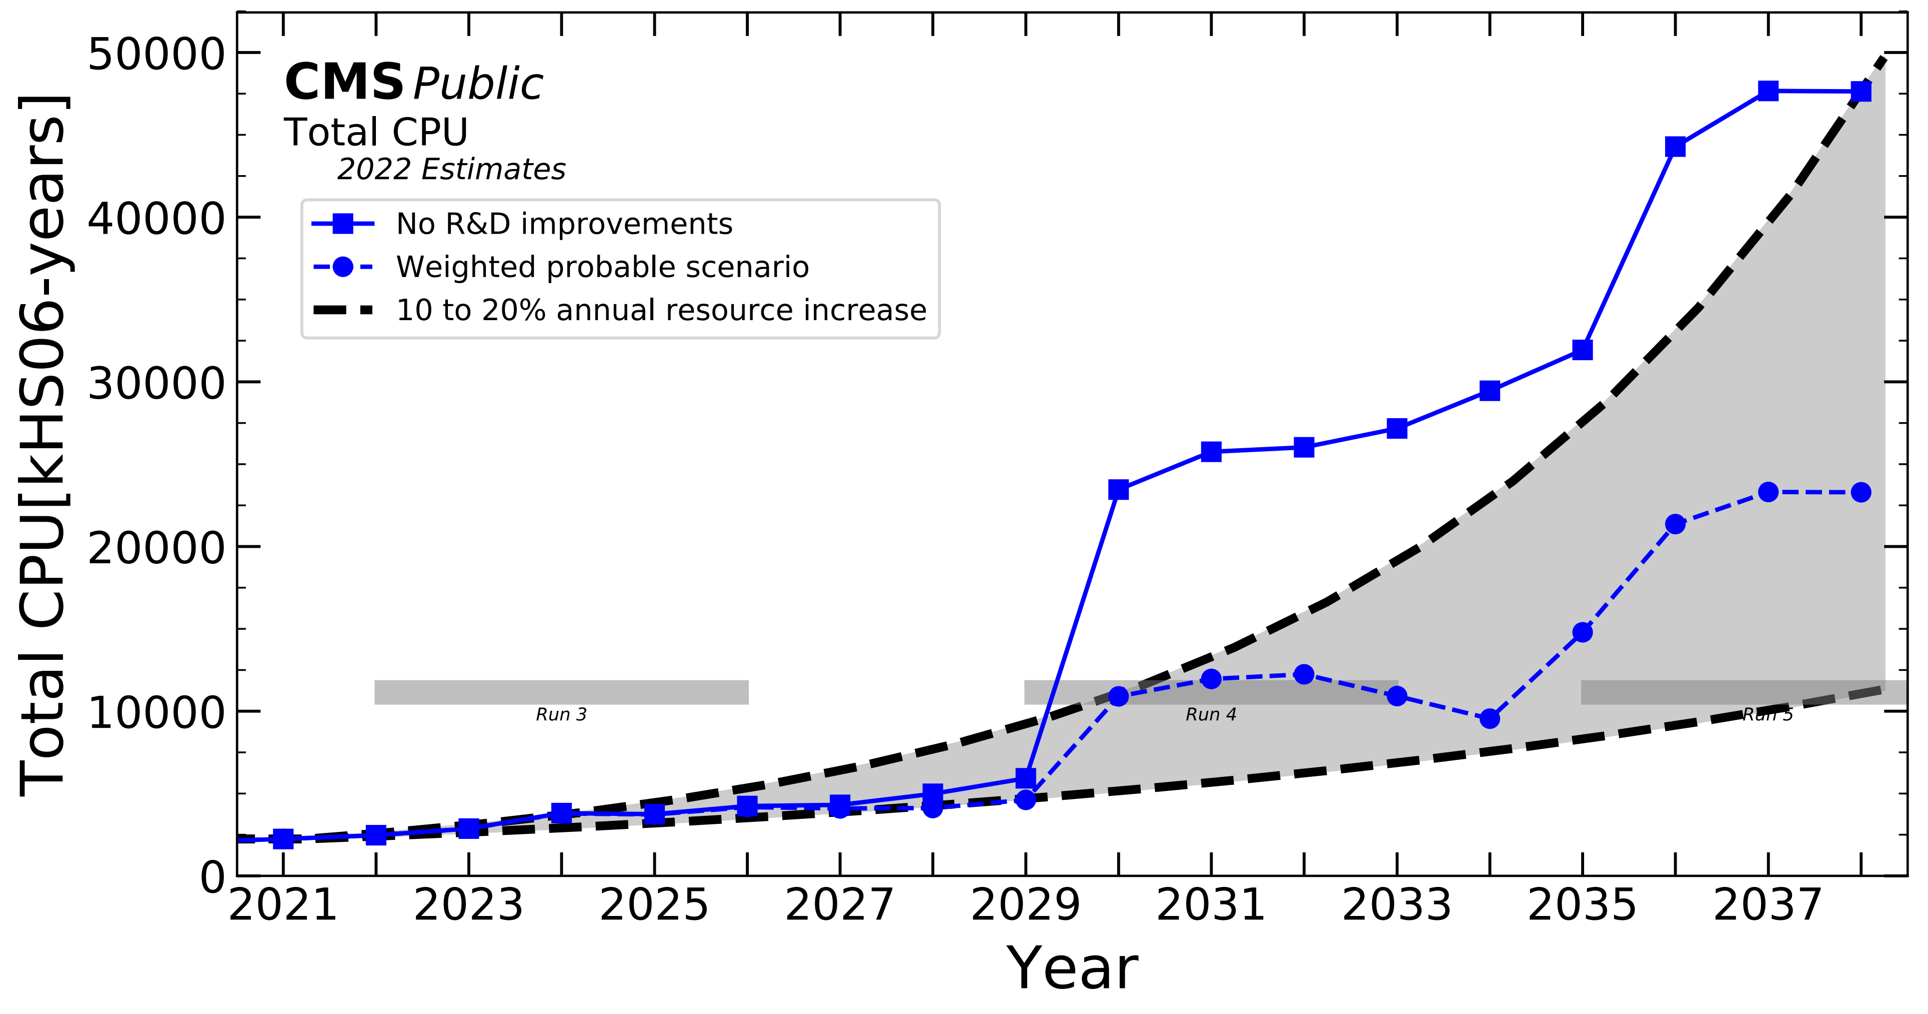

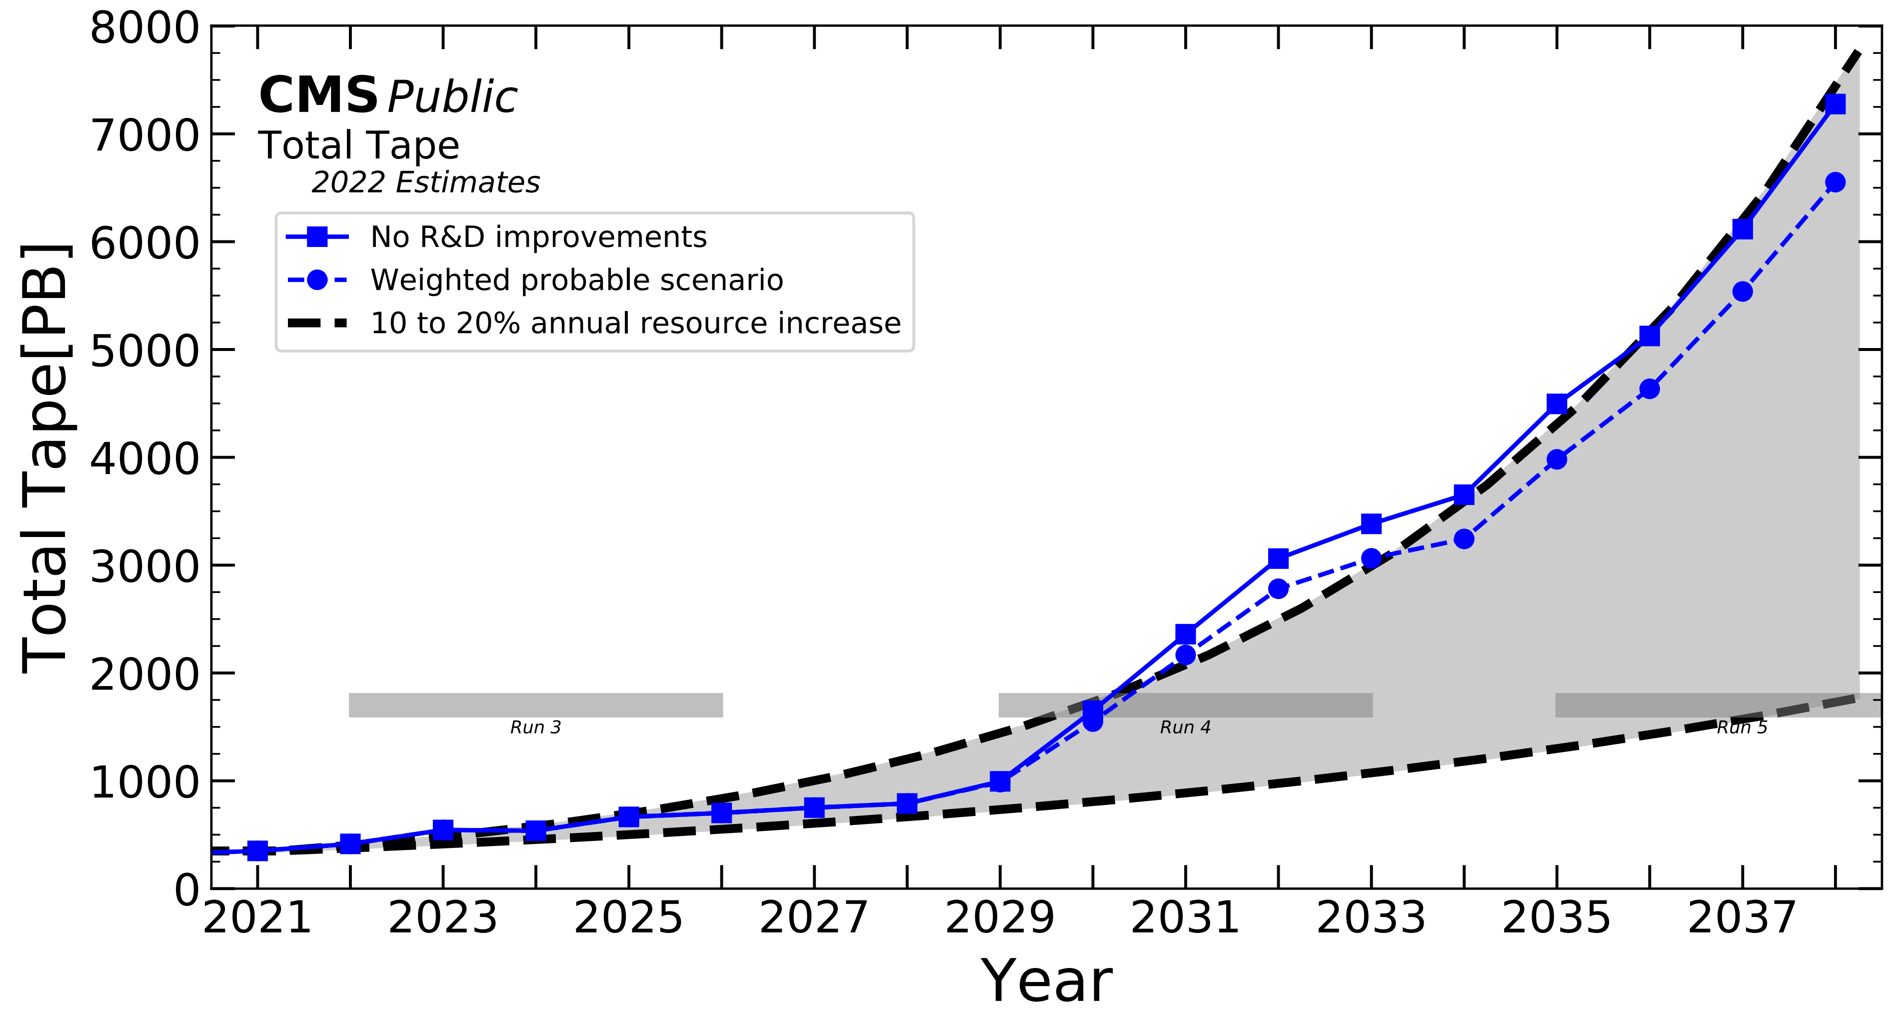

Changing hardware landscape

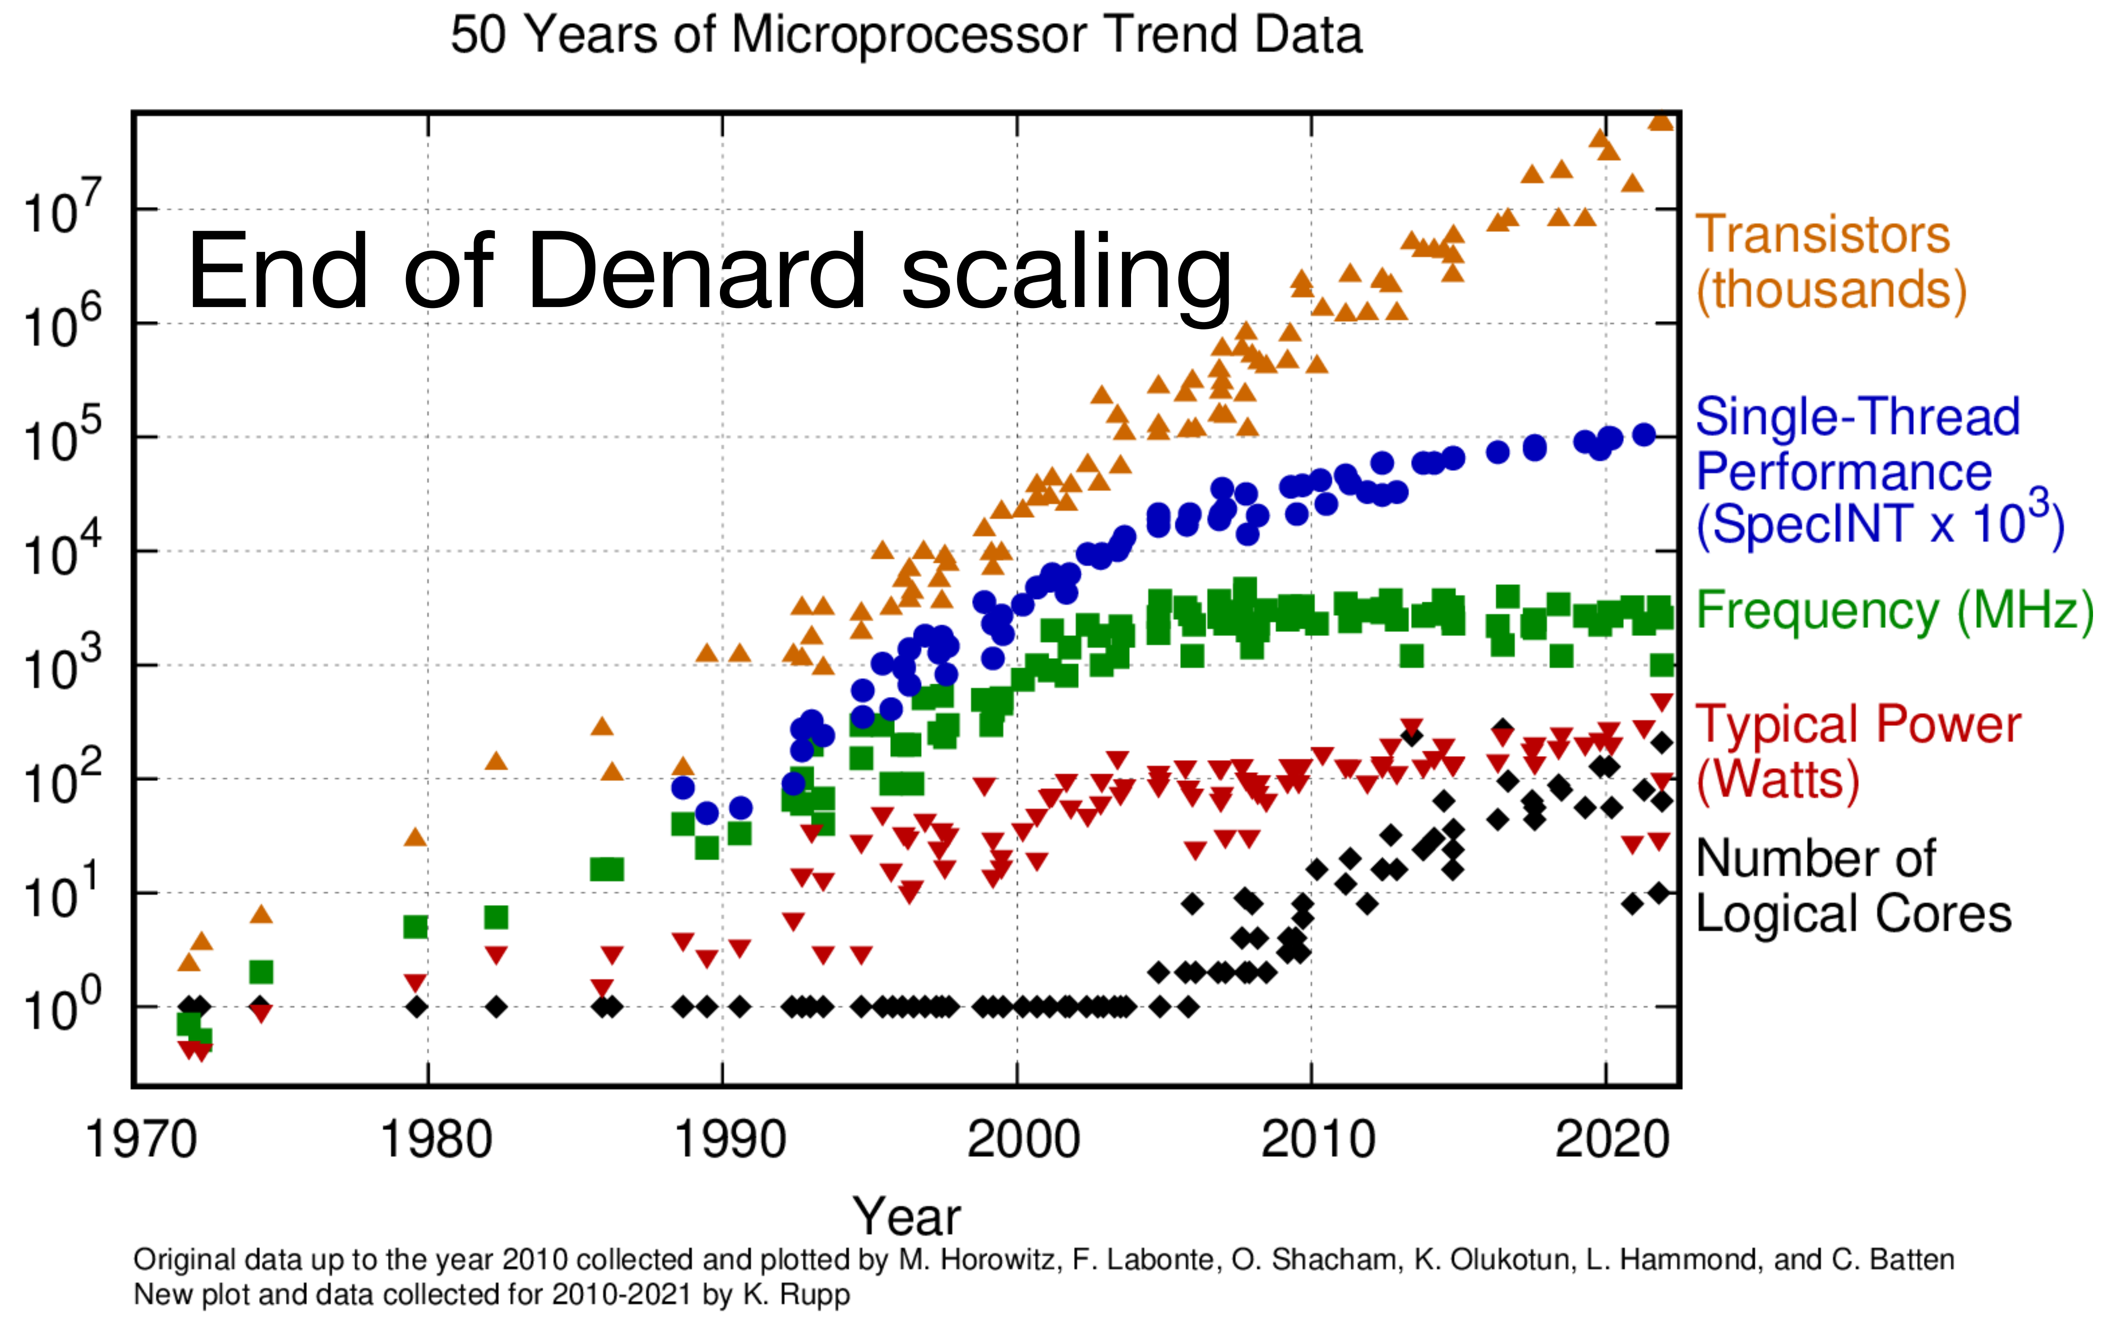

Proactively adapting to new hardware trends: adopted multi-threaded computing

- Crucial to the success of CMS computing

Accelerated heterogeneous compute: Opportunities and challenges

- GPU running at the HLT for CMS Run 3 data-taking.

Highlight R&D activities towards HL-LHC, projects led by USCMS/FNAL.

Efficient reconstruction with increased parallelism: mkFit

Track reconstruction is a large part of the CMS reconstruction CPU budget

Increasingly so at higher pileup of HL-LHC

CMS uses a Kalman Filter algorithm for track reconstruction

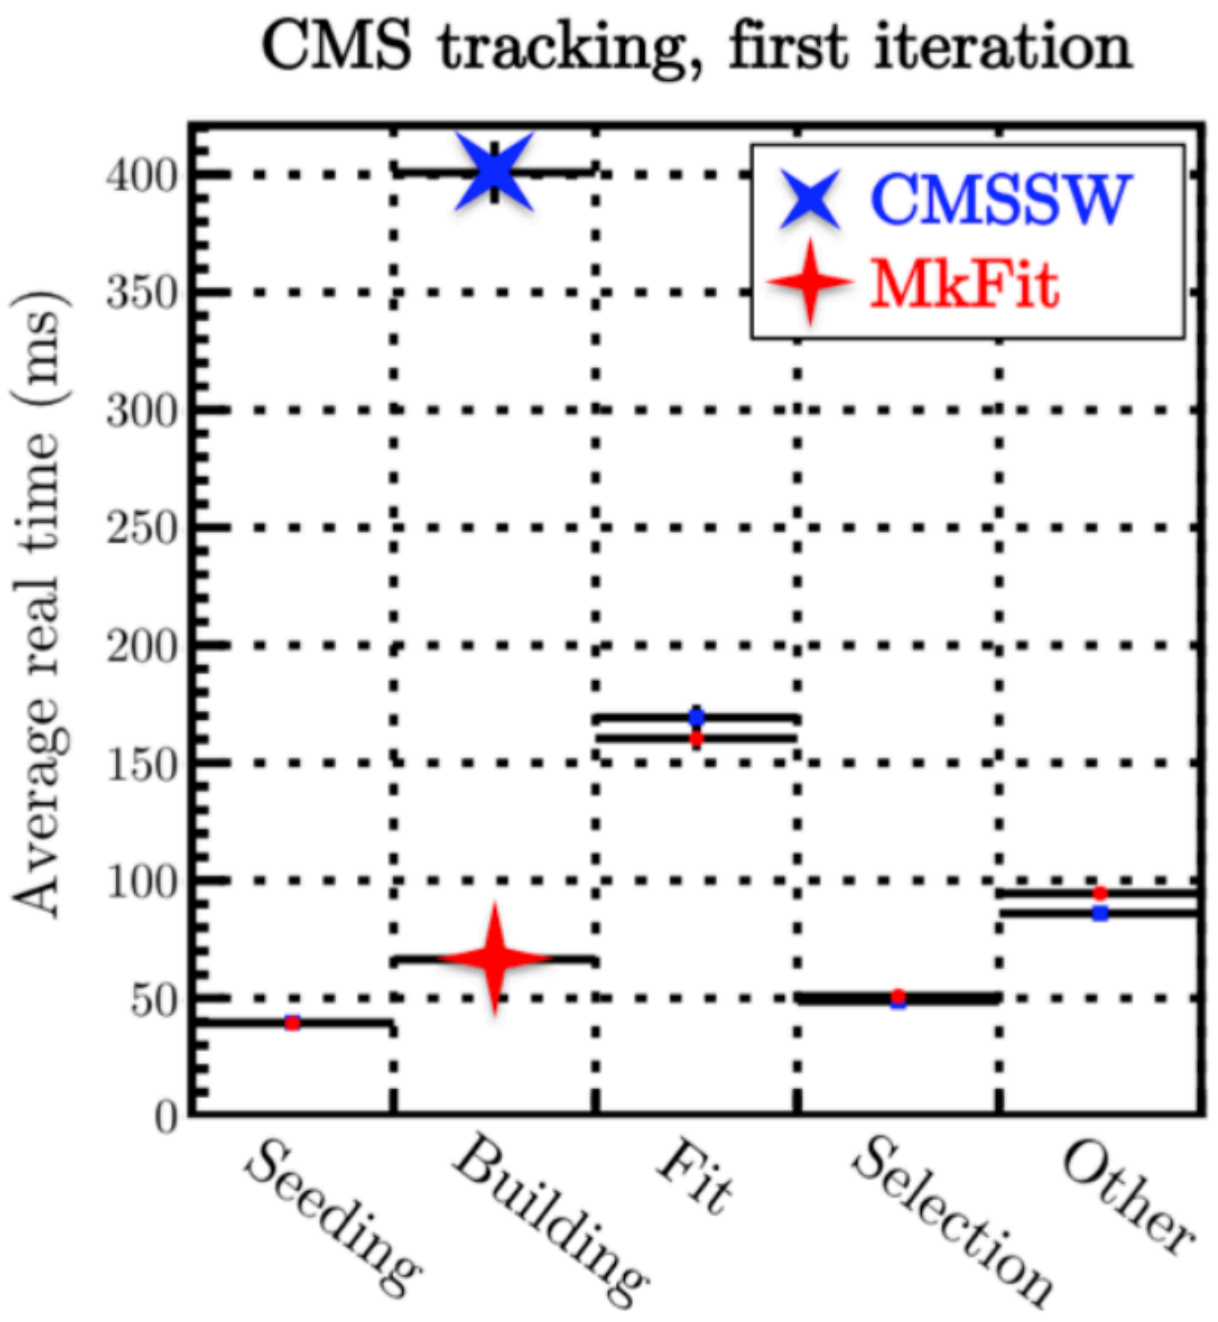

https://iris-hep.org/projects/mkfit.html

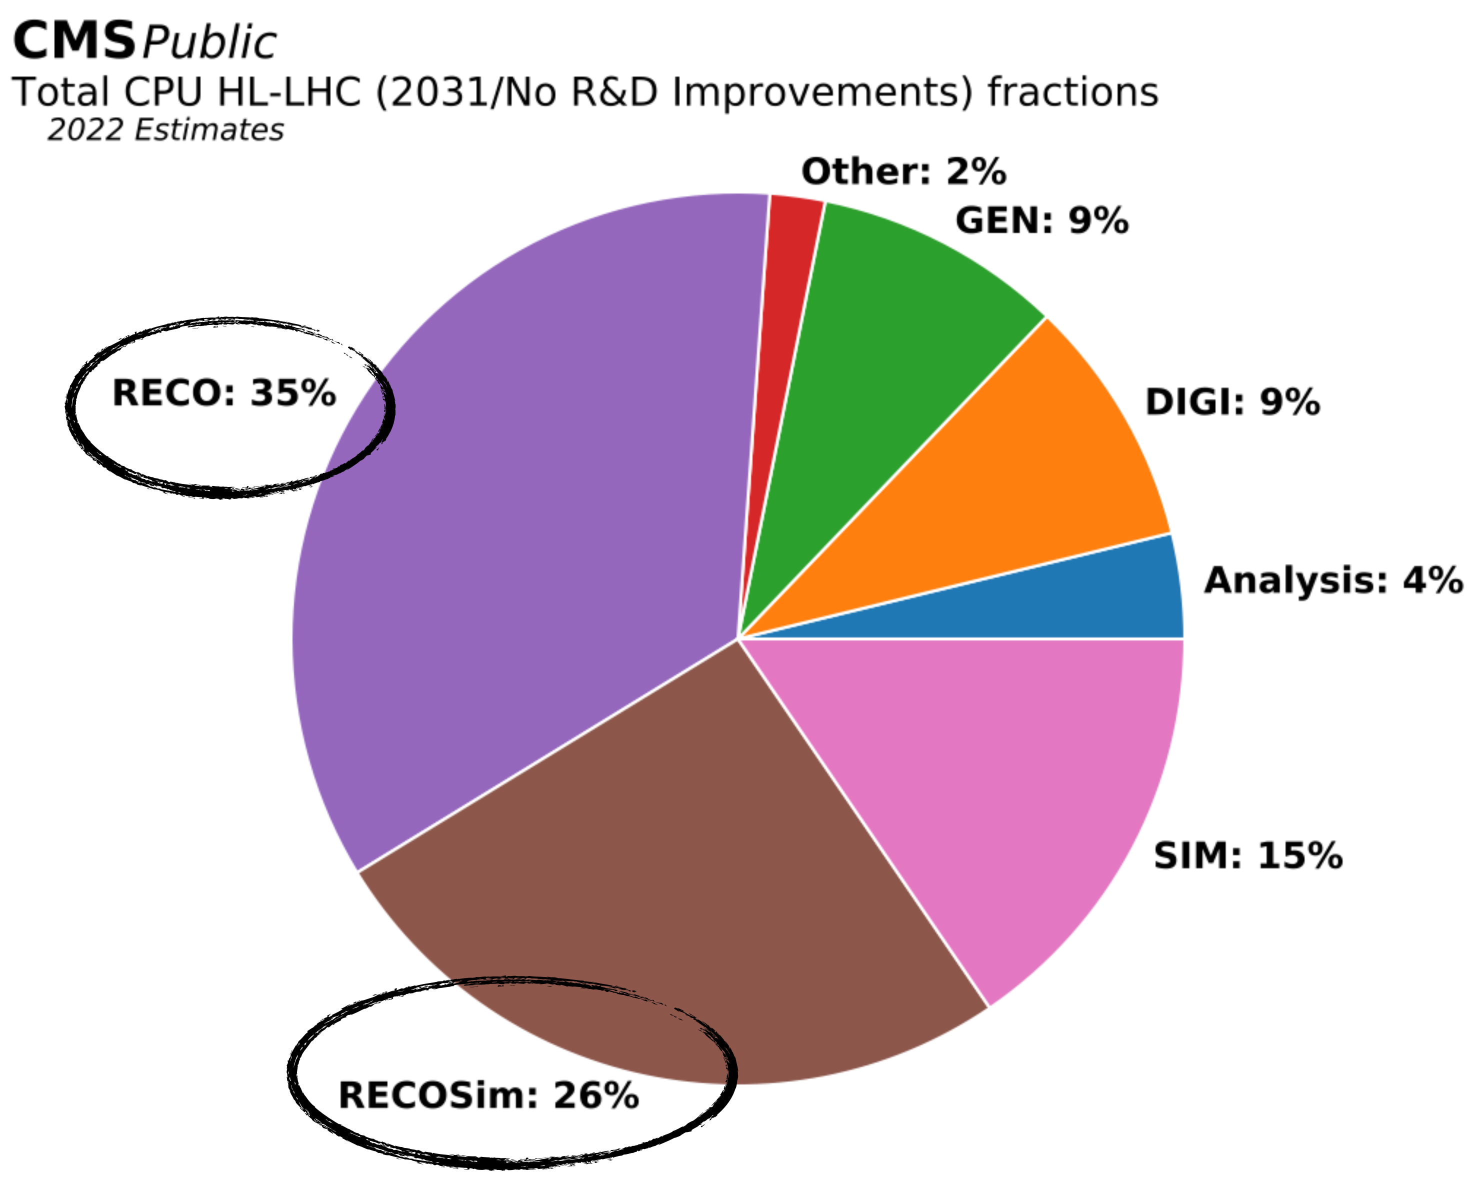

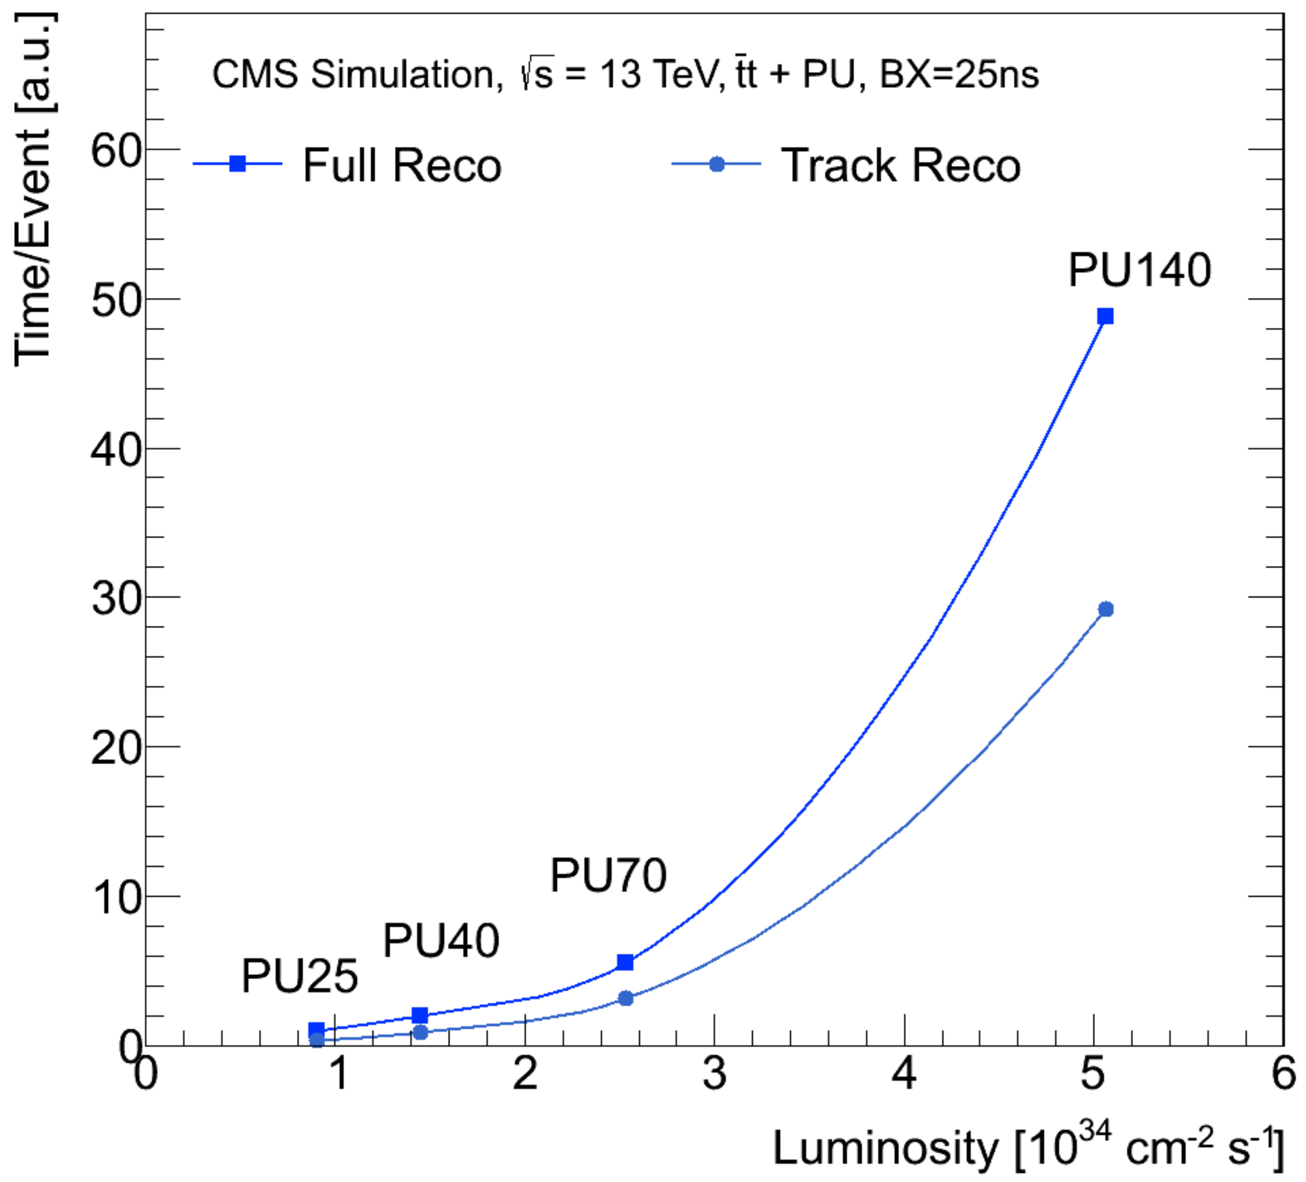

Increases parallelization of algorithm and uses vector units in modern CPUs

Track building up to **6x faster** compared to standard algorithm

Algorithm now deployed with standard CMS software for Run 3, further development includes optimization for GPUs and other heterogeneous computing.

Heterogenous computing: software portability

Available computing resources (including at HPCs) increasingly non-CPU (e.g. GPUs), each flavor of GPU/accelerator may require difference software

e.g CUDA only works on NVIDIA GPUs.

Portability solutions allow offloading to different GPUs with the same code base

Compiler pragma-based approach/Libraries/Language extension

CMS investigated code portability solutions, Alpaka was chosen

New features/compiler supports/new backends

Software →

Hardware ↓

	CUDA	Kokkos	SYCL	HIP	OpenMP	alpaka	std::par
NVIDIA GPU			*intel/llvm compute-cpp*	*hipcc*	*nvc++ LLVM, Cray GCC, XL*		*nvc++*
AMD GPU			*openSYCL intel/llvm*	*hipcc*	*AOMP LLVM Cray*		
Intel GPU			*oneAPI intel/llvm*	*CHIP-SPV: early prototype*	*Intel OneAPI compiler*	*prototype*	*oneapi::dpl*
x86 CPU			*oneAPI intel/llvm computecpp*	*via HIP-CPU Runtime*	*nvc++ LLVM, CCE, GCC, XL*		
FPGA				*via Xilinx Runtime*	*prototype compilers (OpenArc, Intel, etc.)*	*protytype via SYCL*	

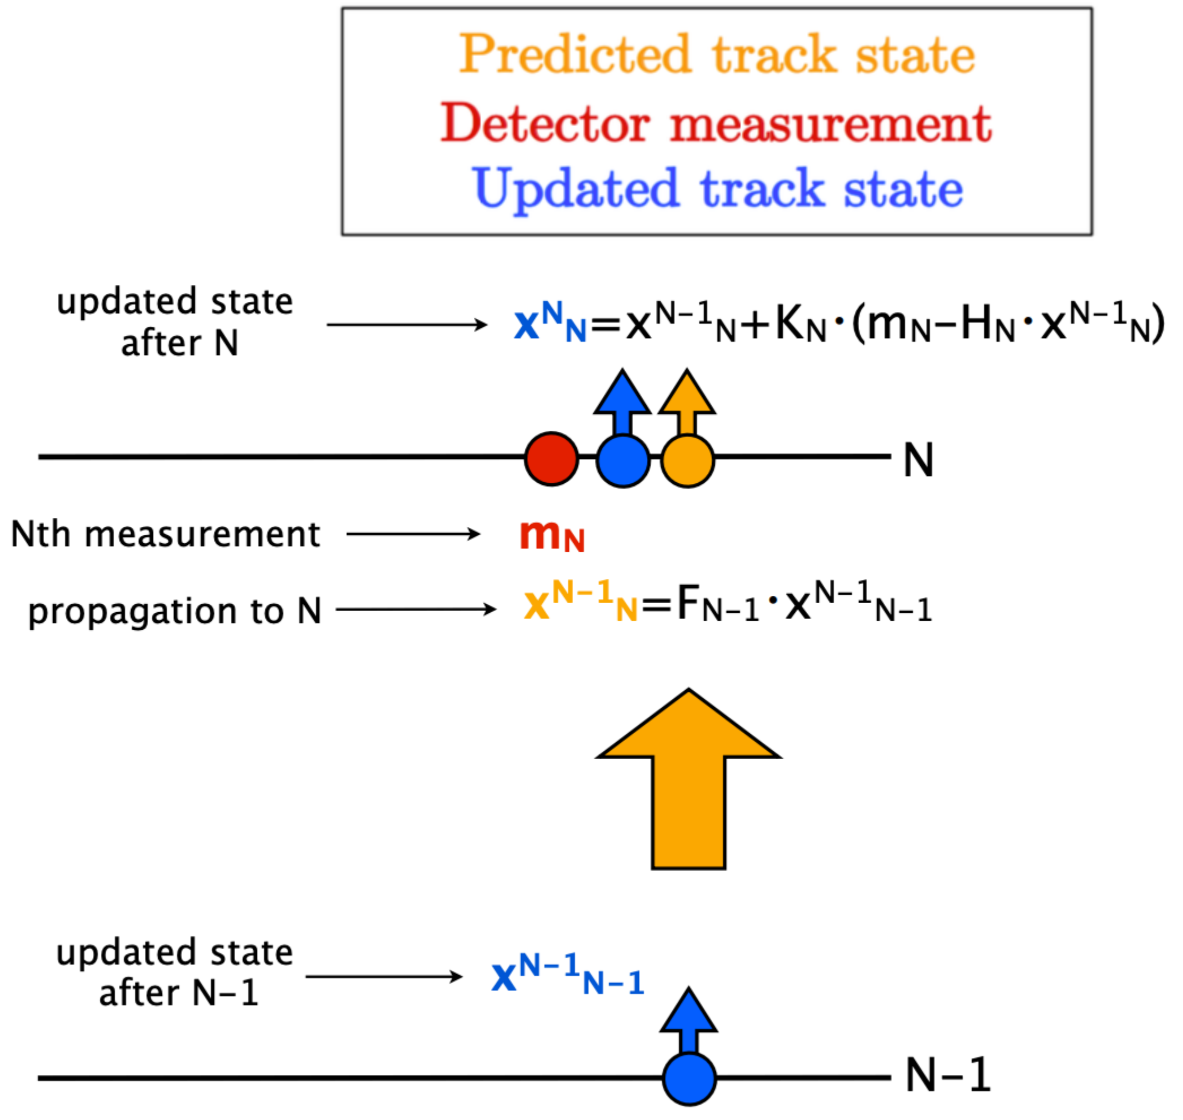

Portability talk at CHEP 2023

Steps of the mkFit core math of parallelized track reconstruction tested on various hardware software combinations.

Heterogenous resources ‘as-a-service’: SONIC

‘Services for Optimized Network Inference on Co-processors’

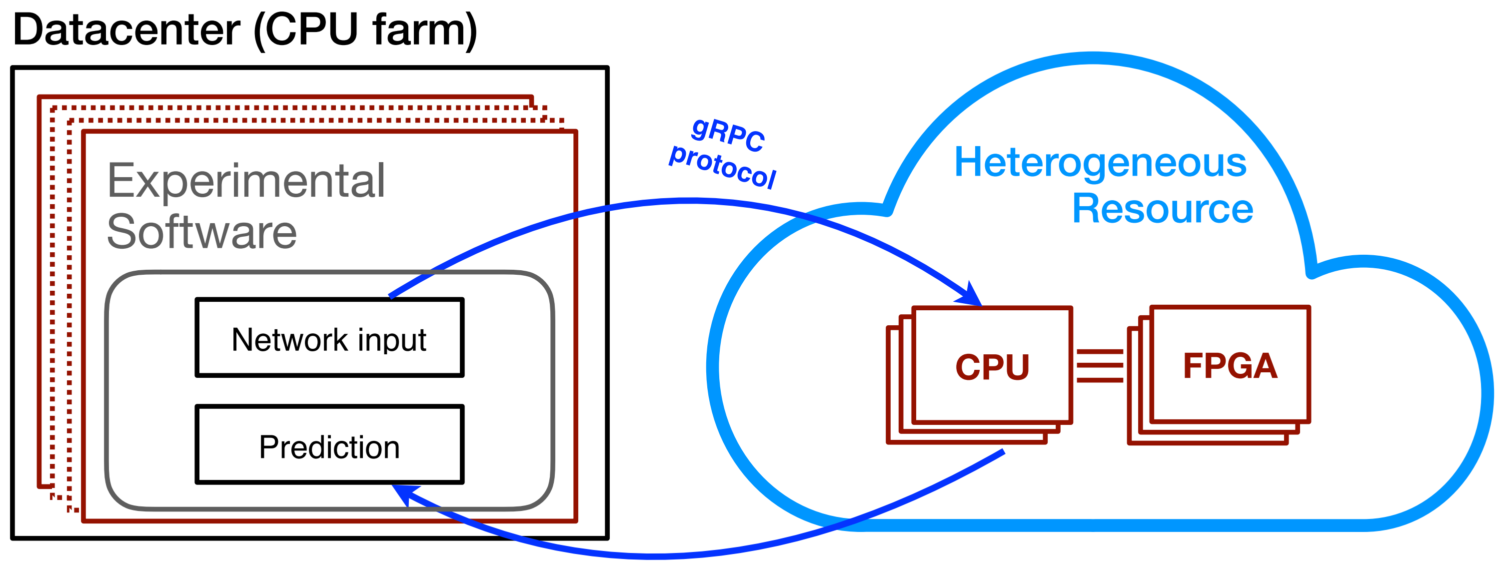

Acceleration ‘as-a-service’

NVIDIA Triton: ML inference server

Flexible CPU to GPU ratio

Maximize acceleration with heterogenous resources

Flexible fallback solutions with homogenous client software

Server handles heterogeneous hardware types, acceleration kernels (CUDA & alpaka).

Recent result presented at CHEP 2023.

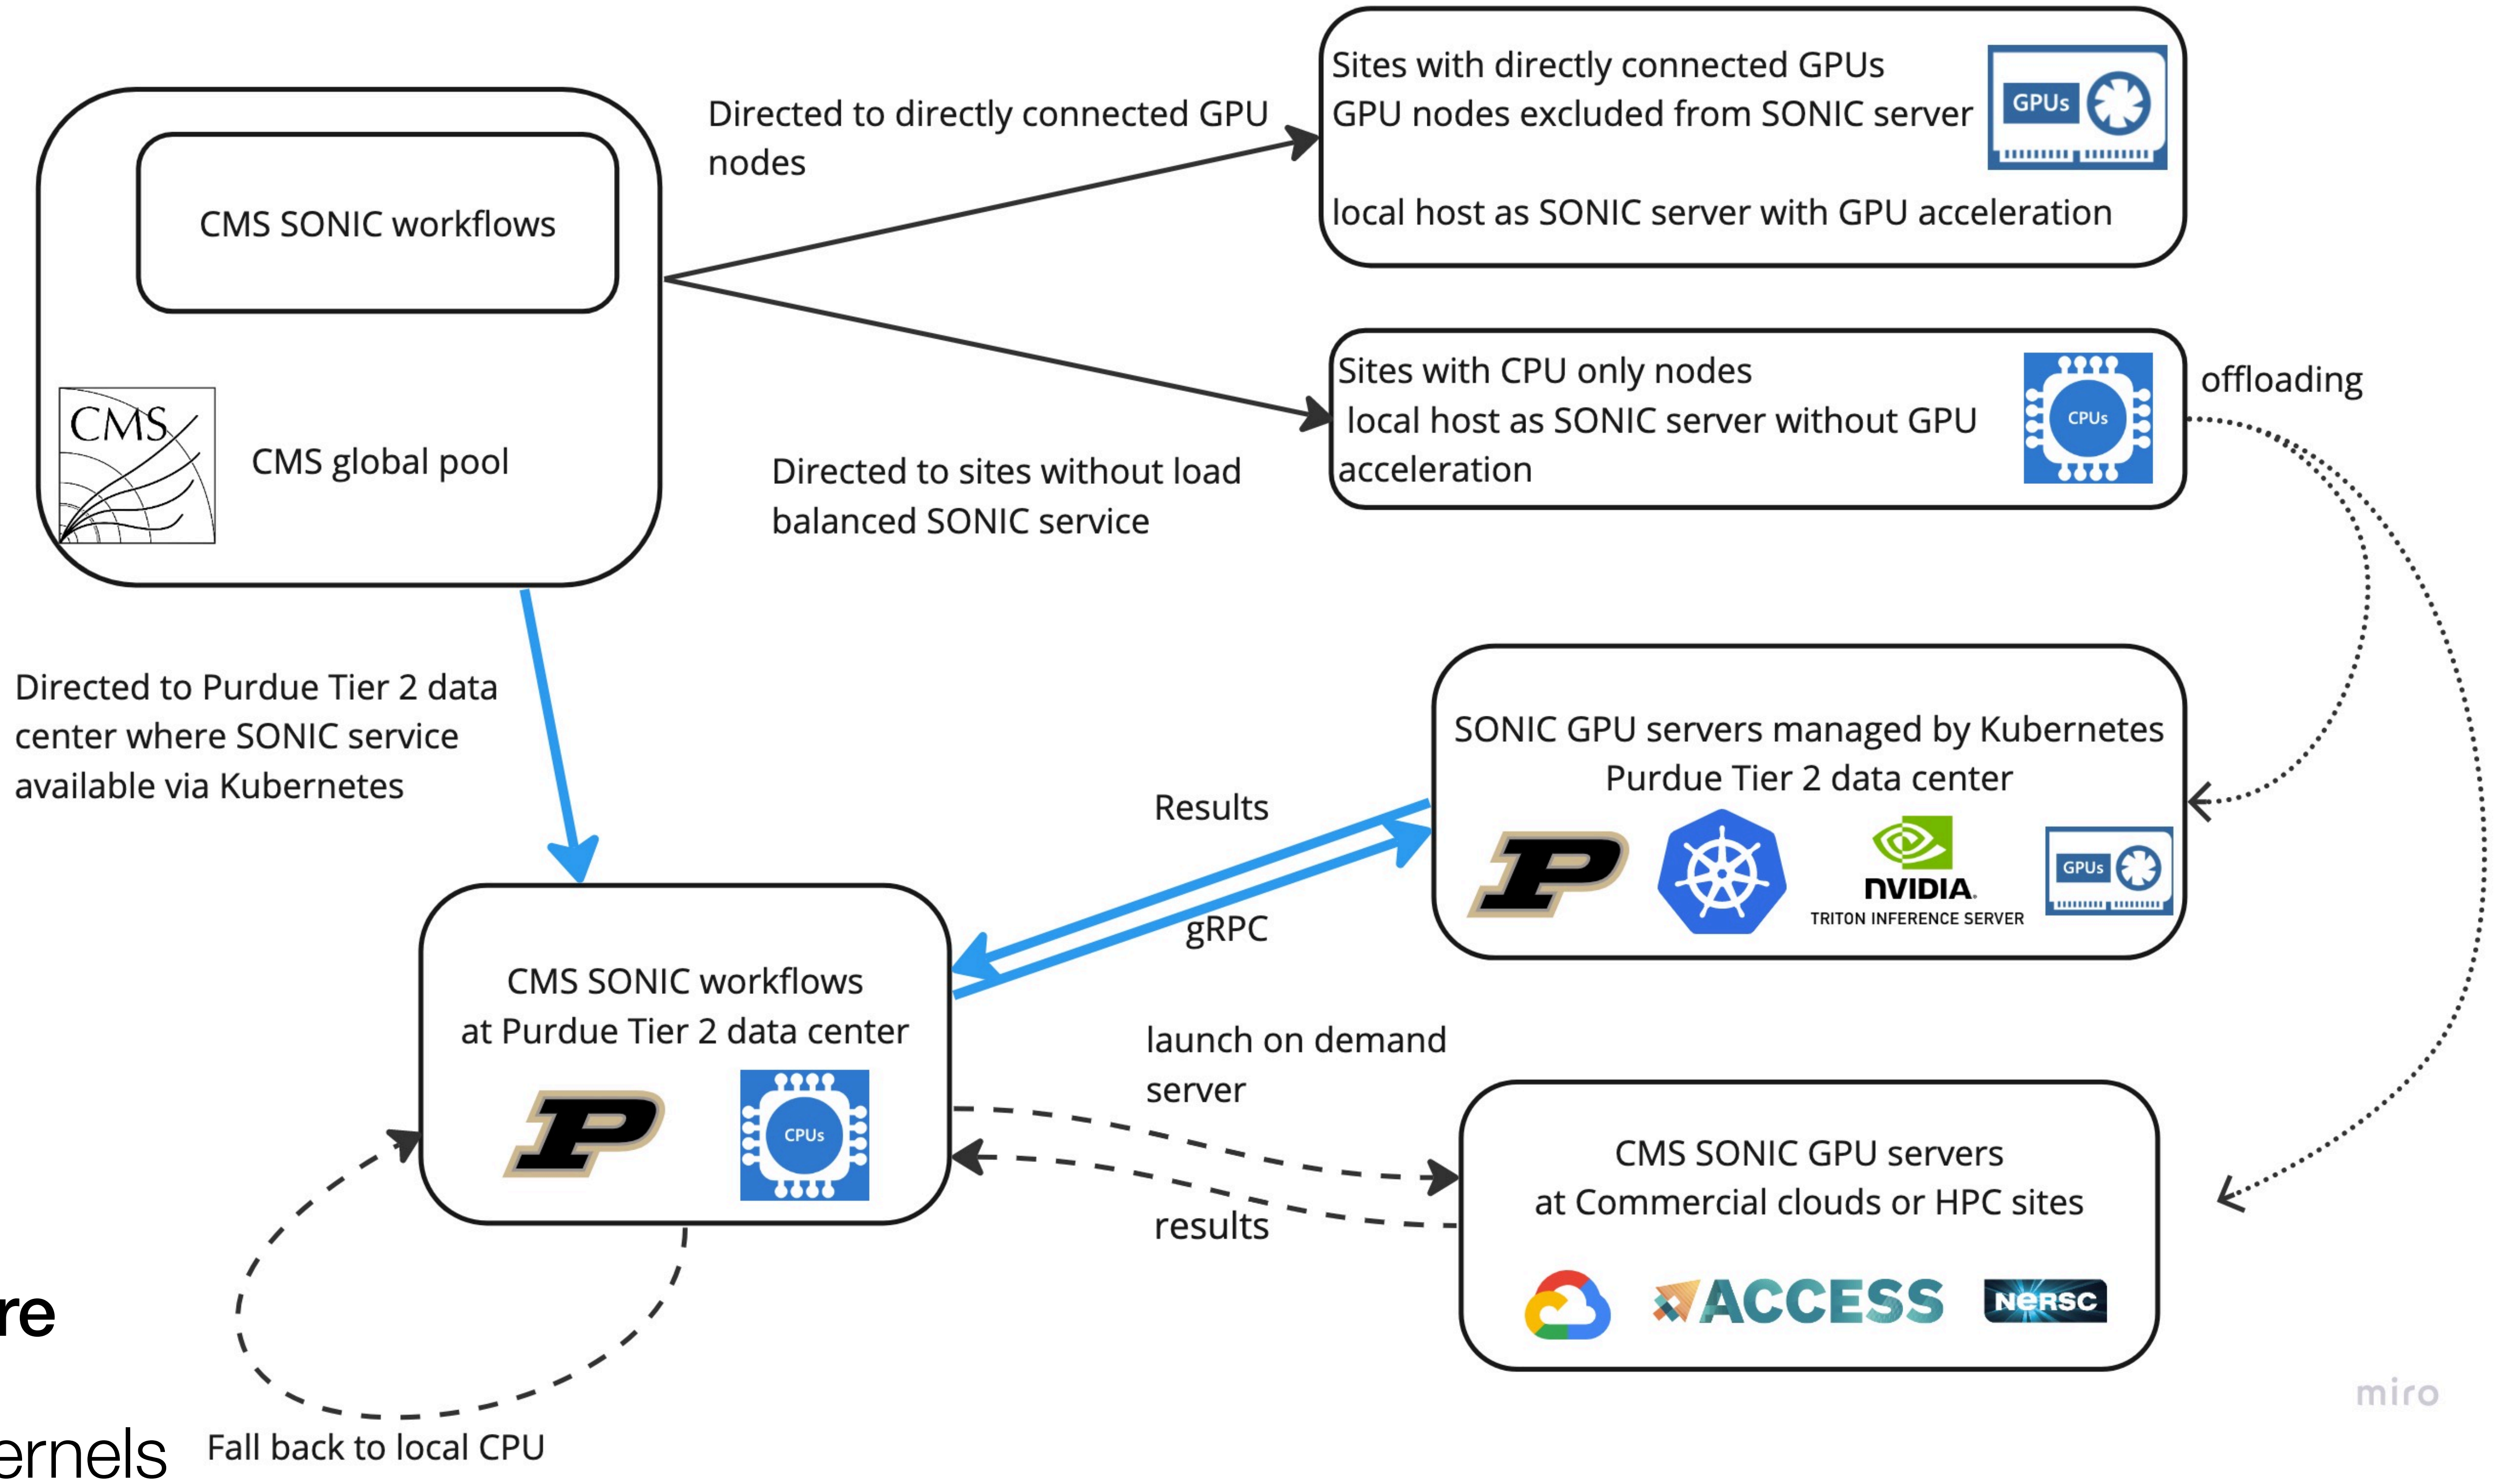

SONIC workflow in CMS production system

Analysis Software

Pythonic scientific computing ecosystem

Seamless interface with data science and ML tools

Columnar analysis (Coffea)

Leverage many many cores!

Investing in facility support at CMS sites (next slide).

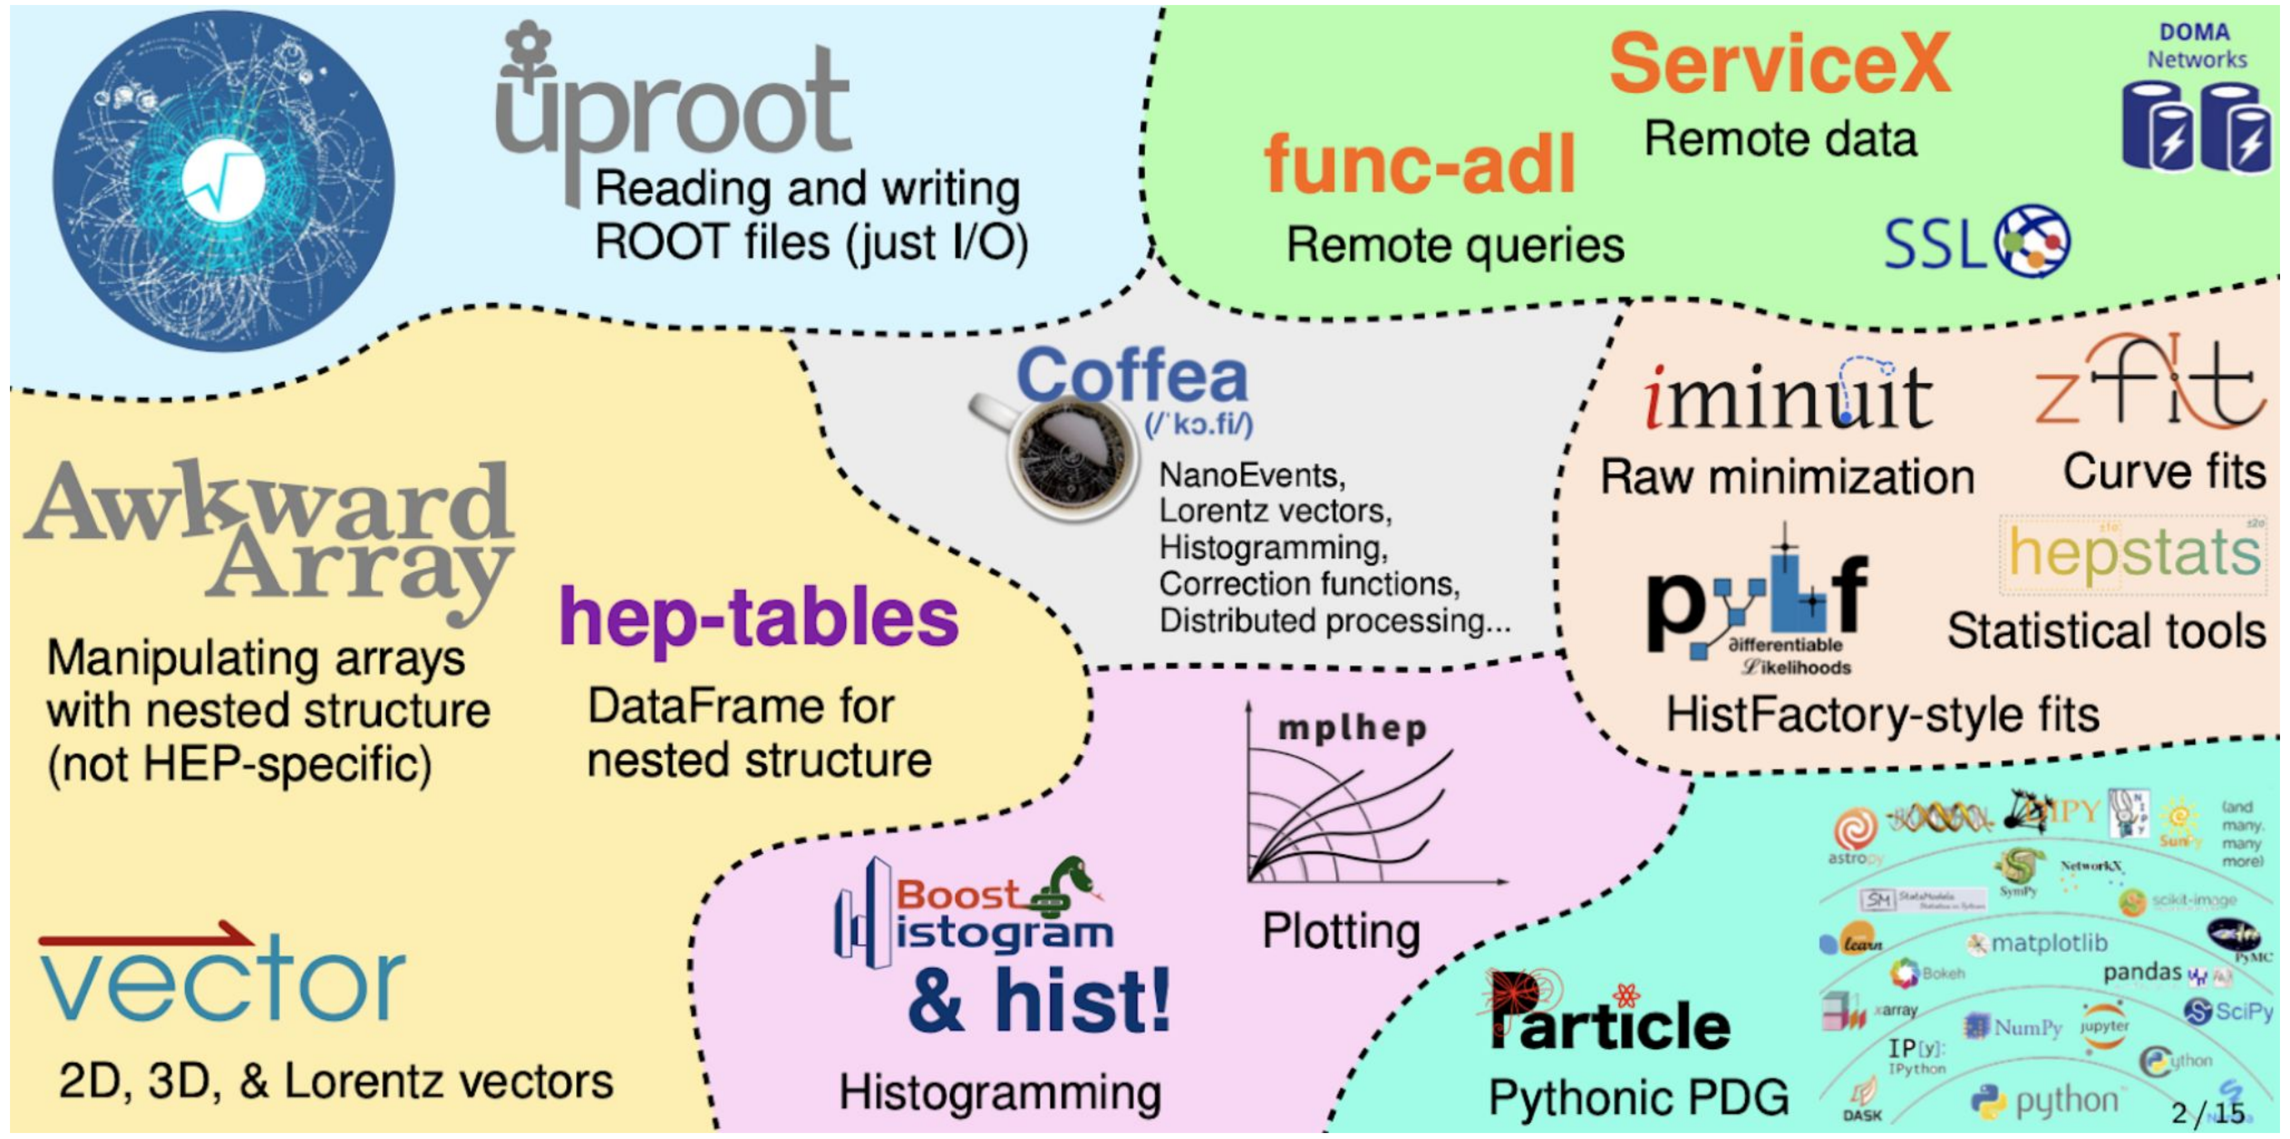

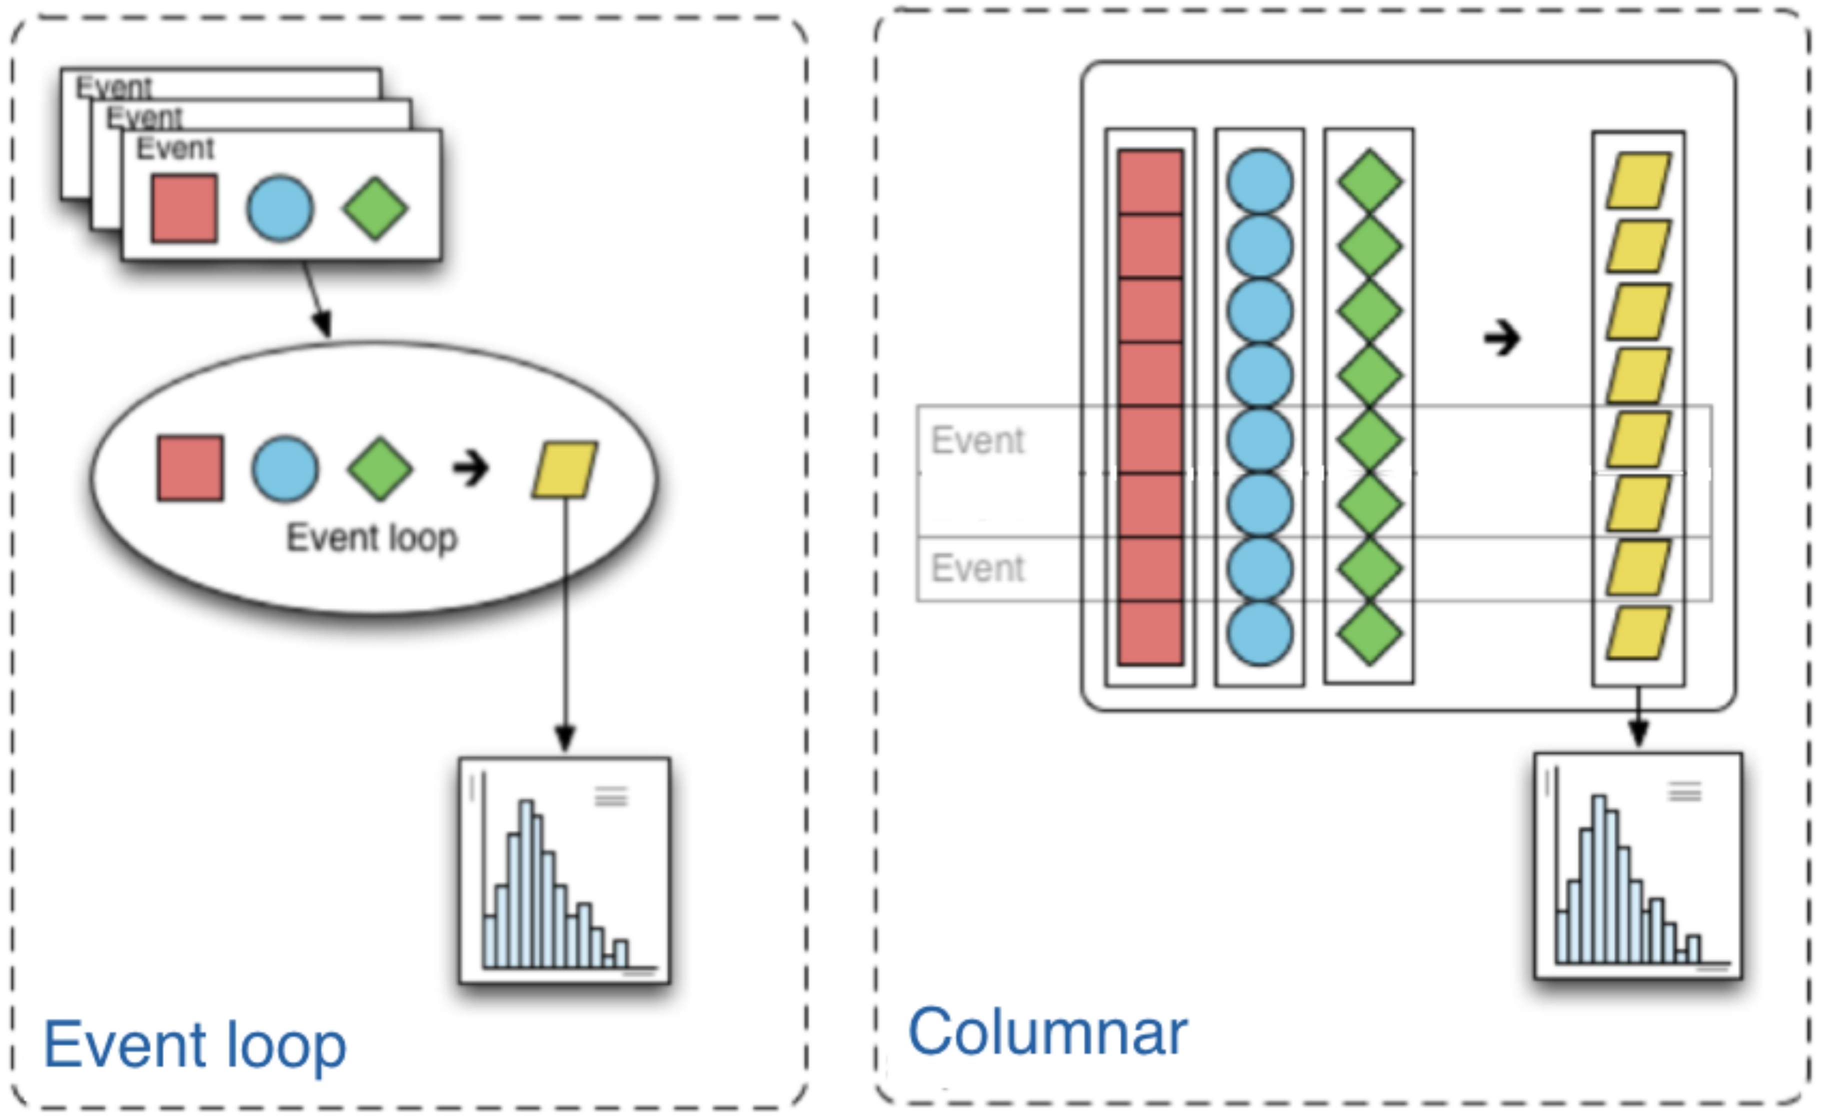

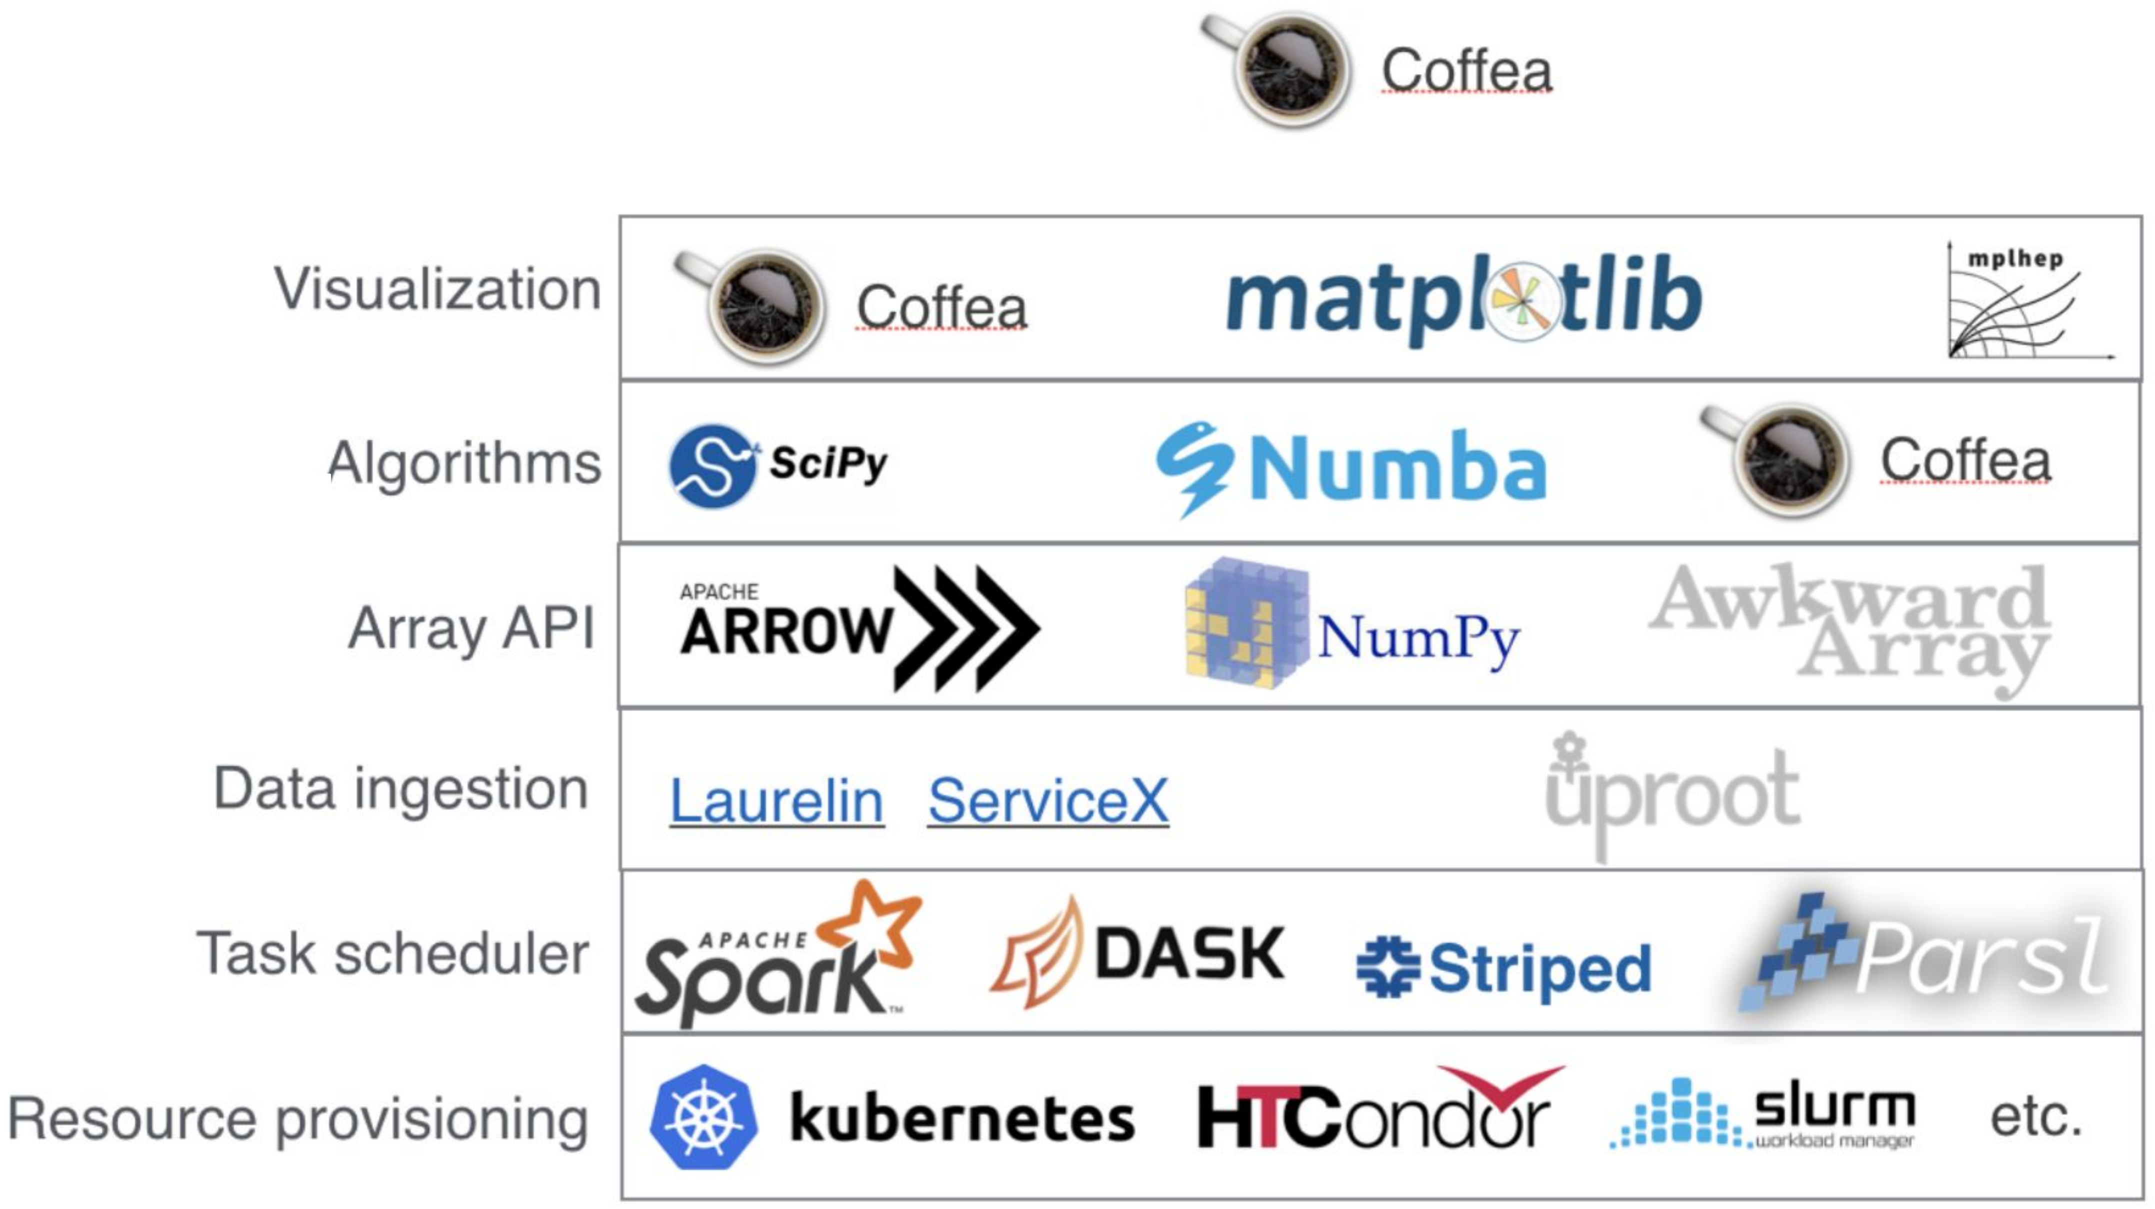

Analysis Facilities

Analysis computing presently run on facilities designed for reconstruction

Investing in analysis facilities at various USCMS sites

Build infrastructure support for columnar analysis with modern softwares

Traditional style workflows can also run on these facilities, or possible with new languages, e.g Julia

Quick analysis turn around time demonstrated: ~1 hour for a Run 2 analysis demonstrated (instead of a day or more)

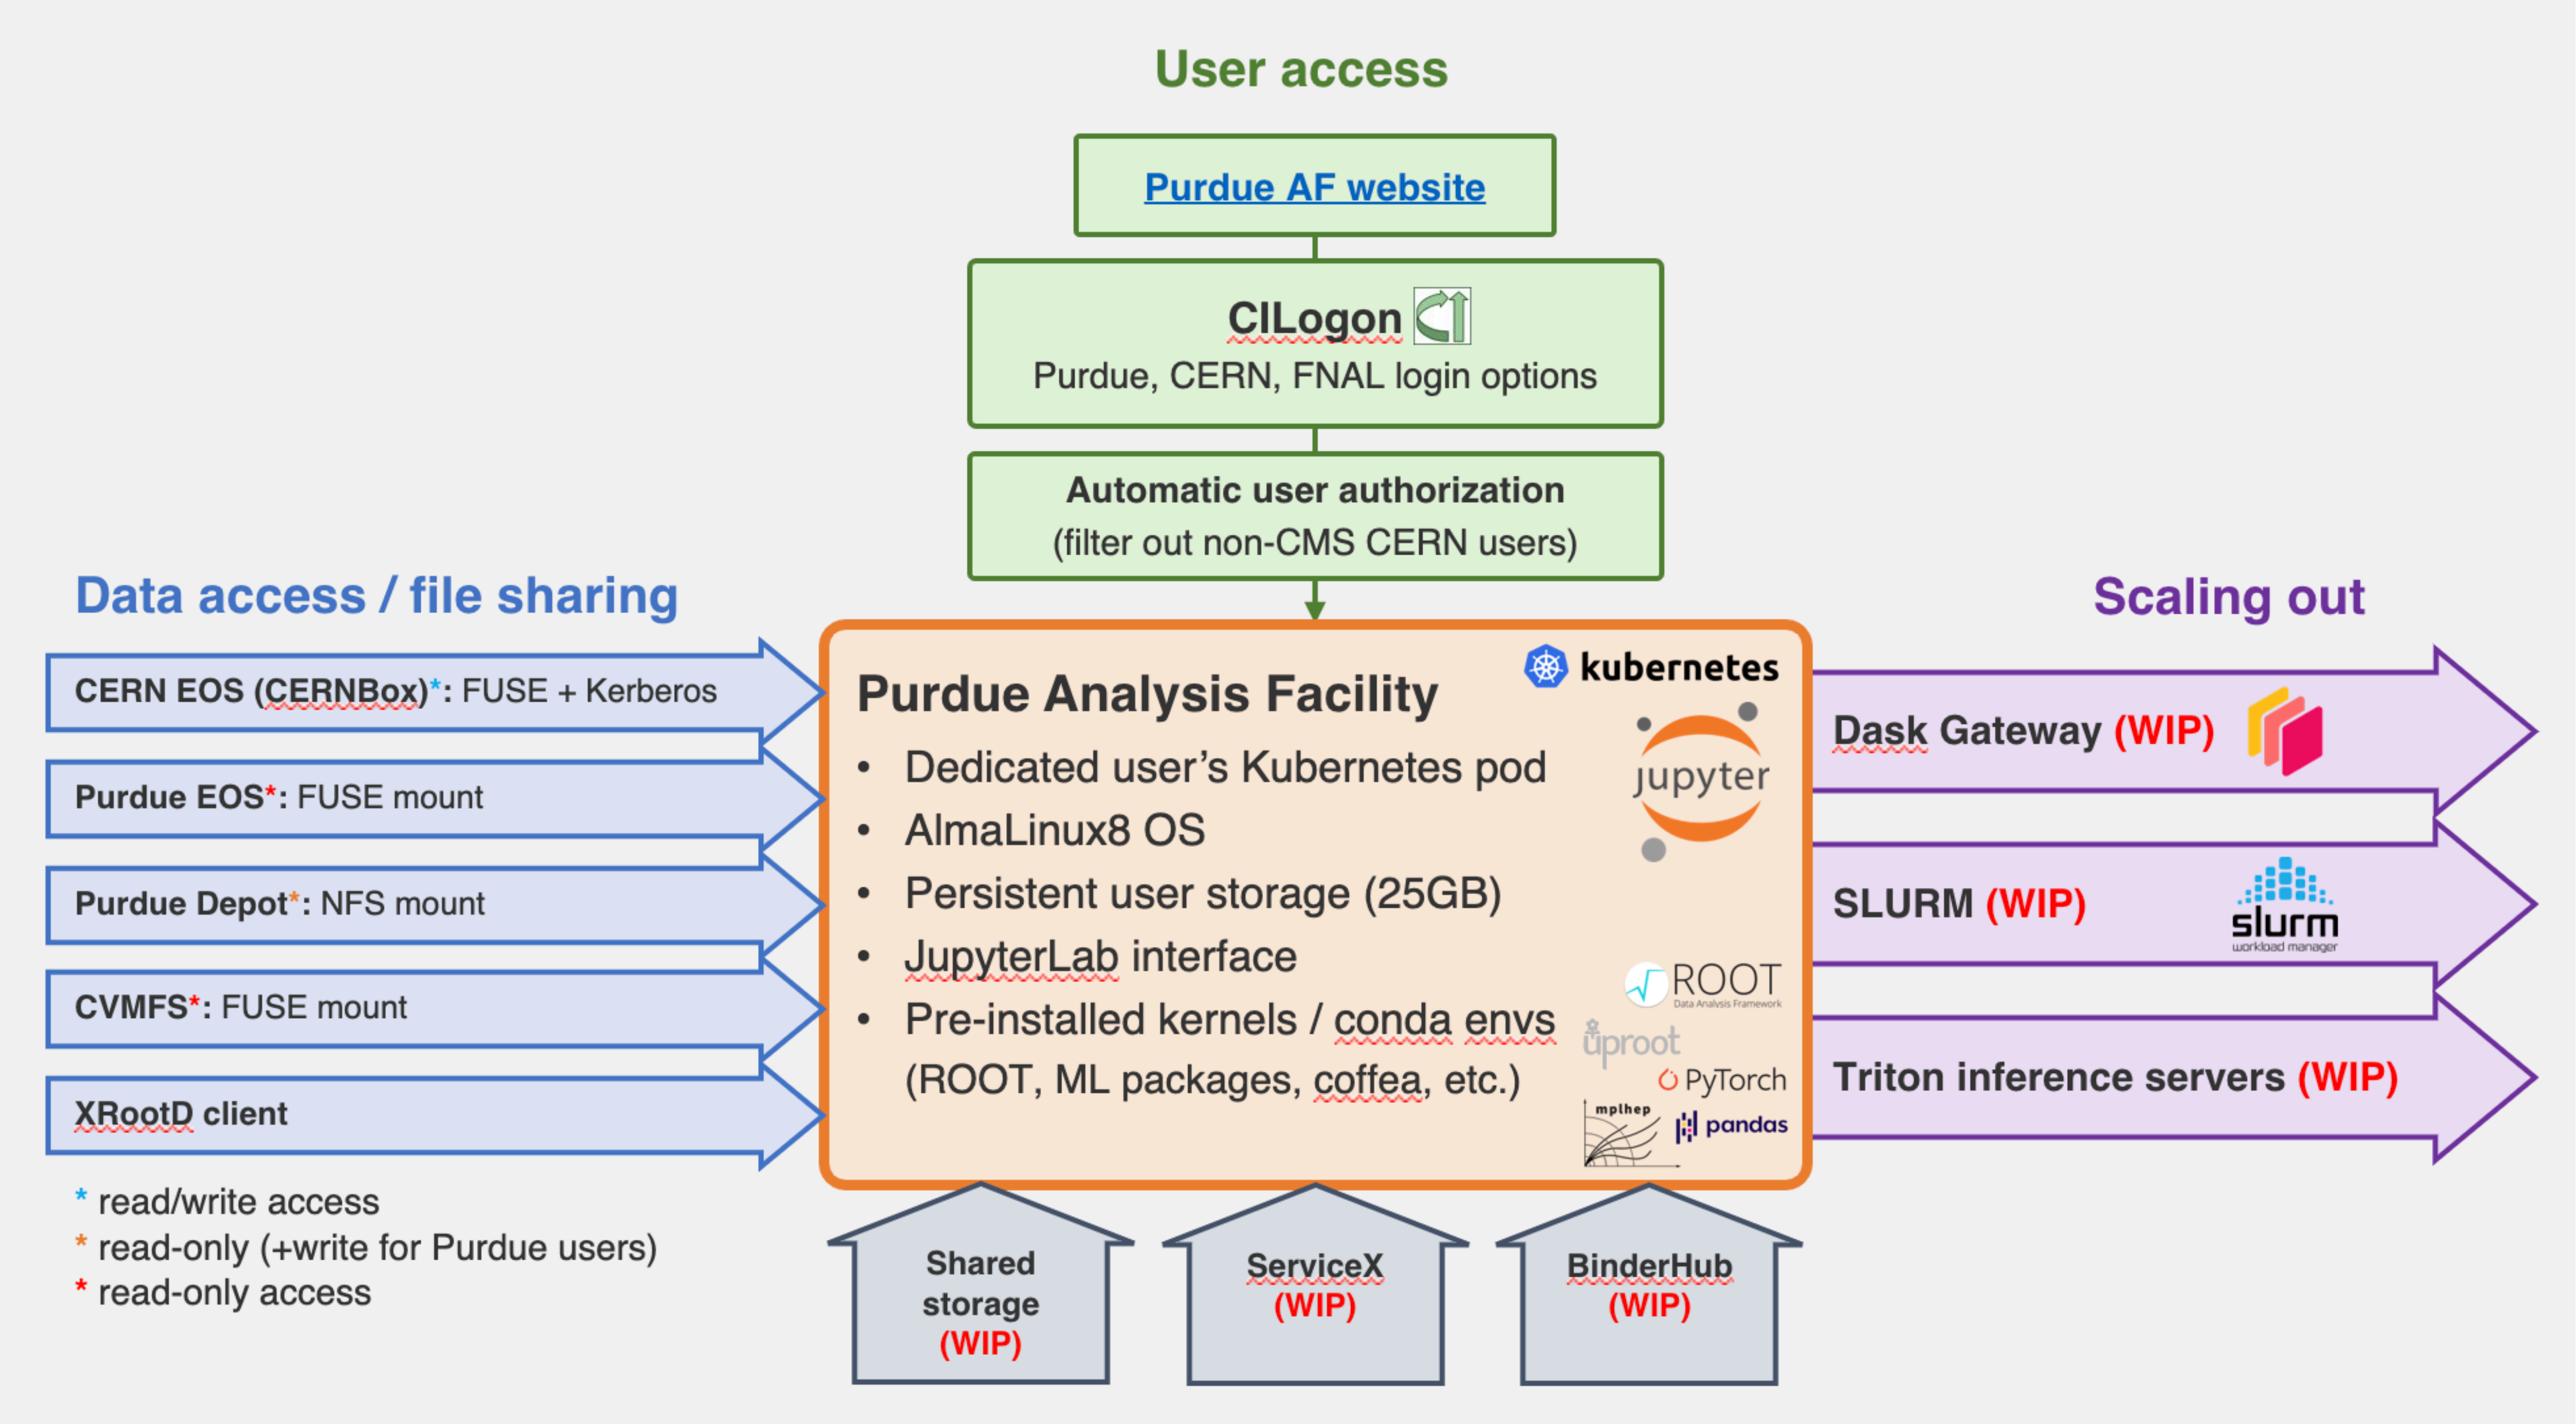

Example: Analysis facility prototype at Purdue

Storage R&D

Object data formats provide novel data management capabilities

- Compared to **current tier-based file model**
- Reduce disk storage requirements, obviate the need to define data tiers

Developing a prototype object store service

- Using Ceph S3 protocol
- In collaboration with FNAL CSAID storage group
- I/O scales better than ROOT files

Tier-based scheme

MiniAOD Data product	KB per event	
	v1	v2
packed+pruned genParticles	5.7	5.7
slimmedElectrons	1.3	1.3
Others	48.7	48.7
Total	55.7	55.7

Object store scheme

MiniAOD Data product	KB per event	
	v1	v2
packed+pruned genParticles	5.7	-
slimmedElectrons	1.3	-
Others	48.7	-
Updated slimmedElectrons	-	1.3
Total	55.7	1.3

Talk at CHEP 2023

Big data challenges in CMS and beyond

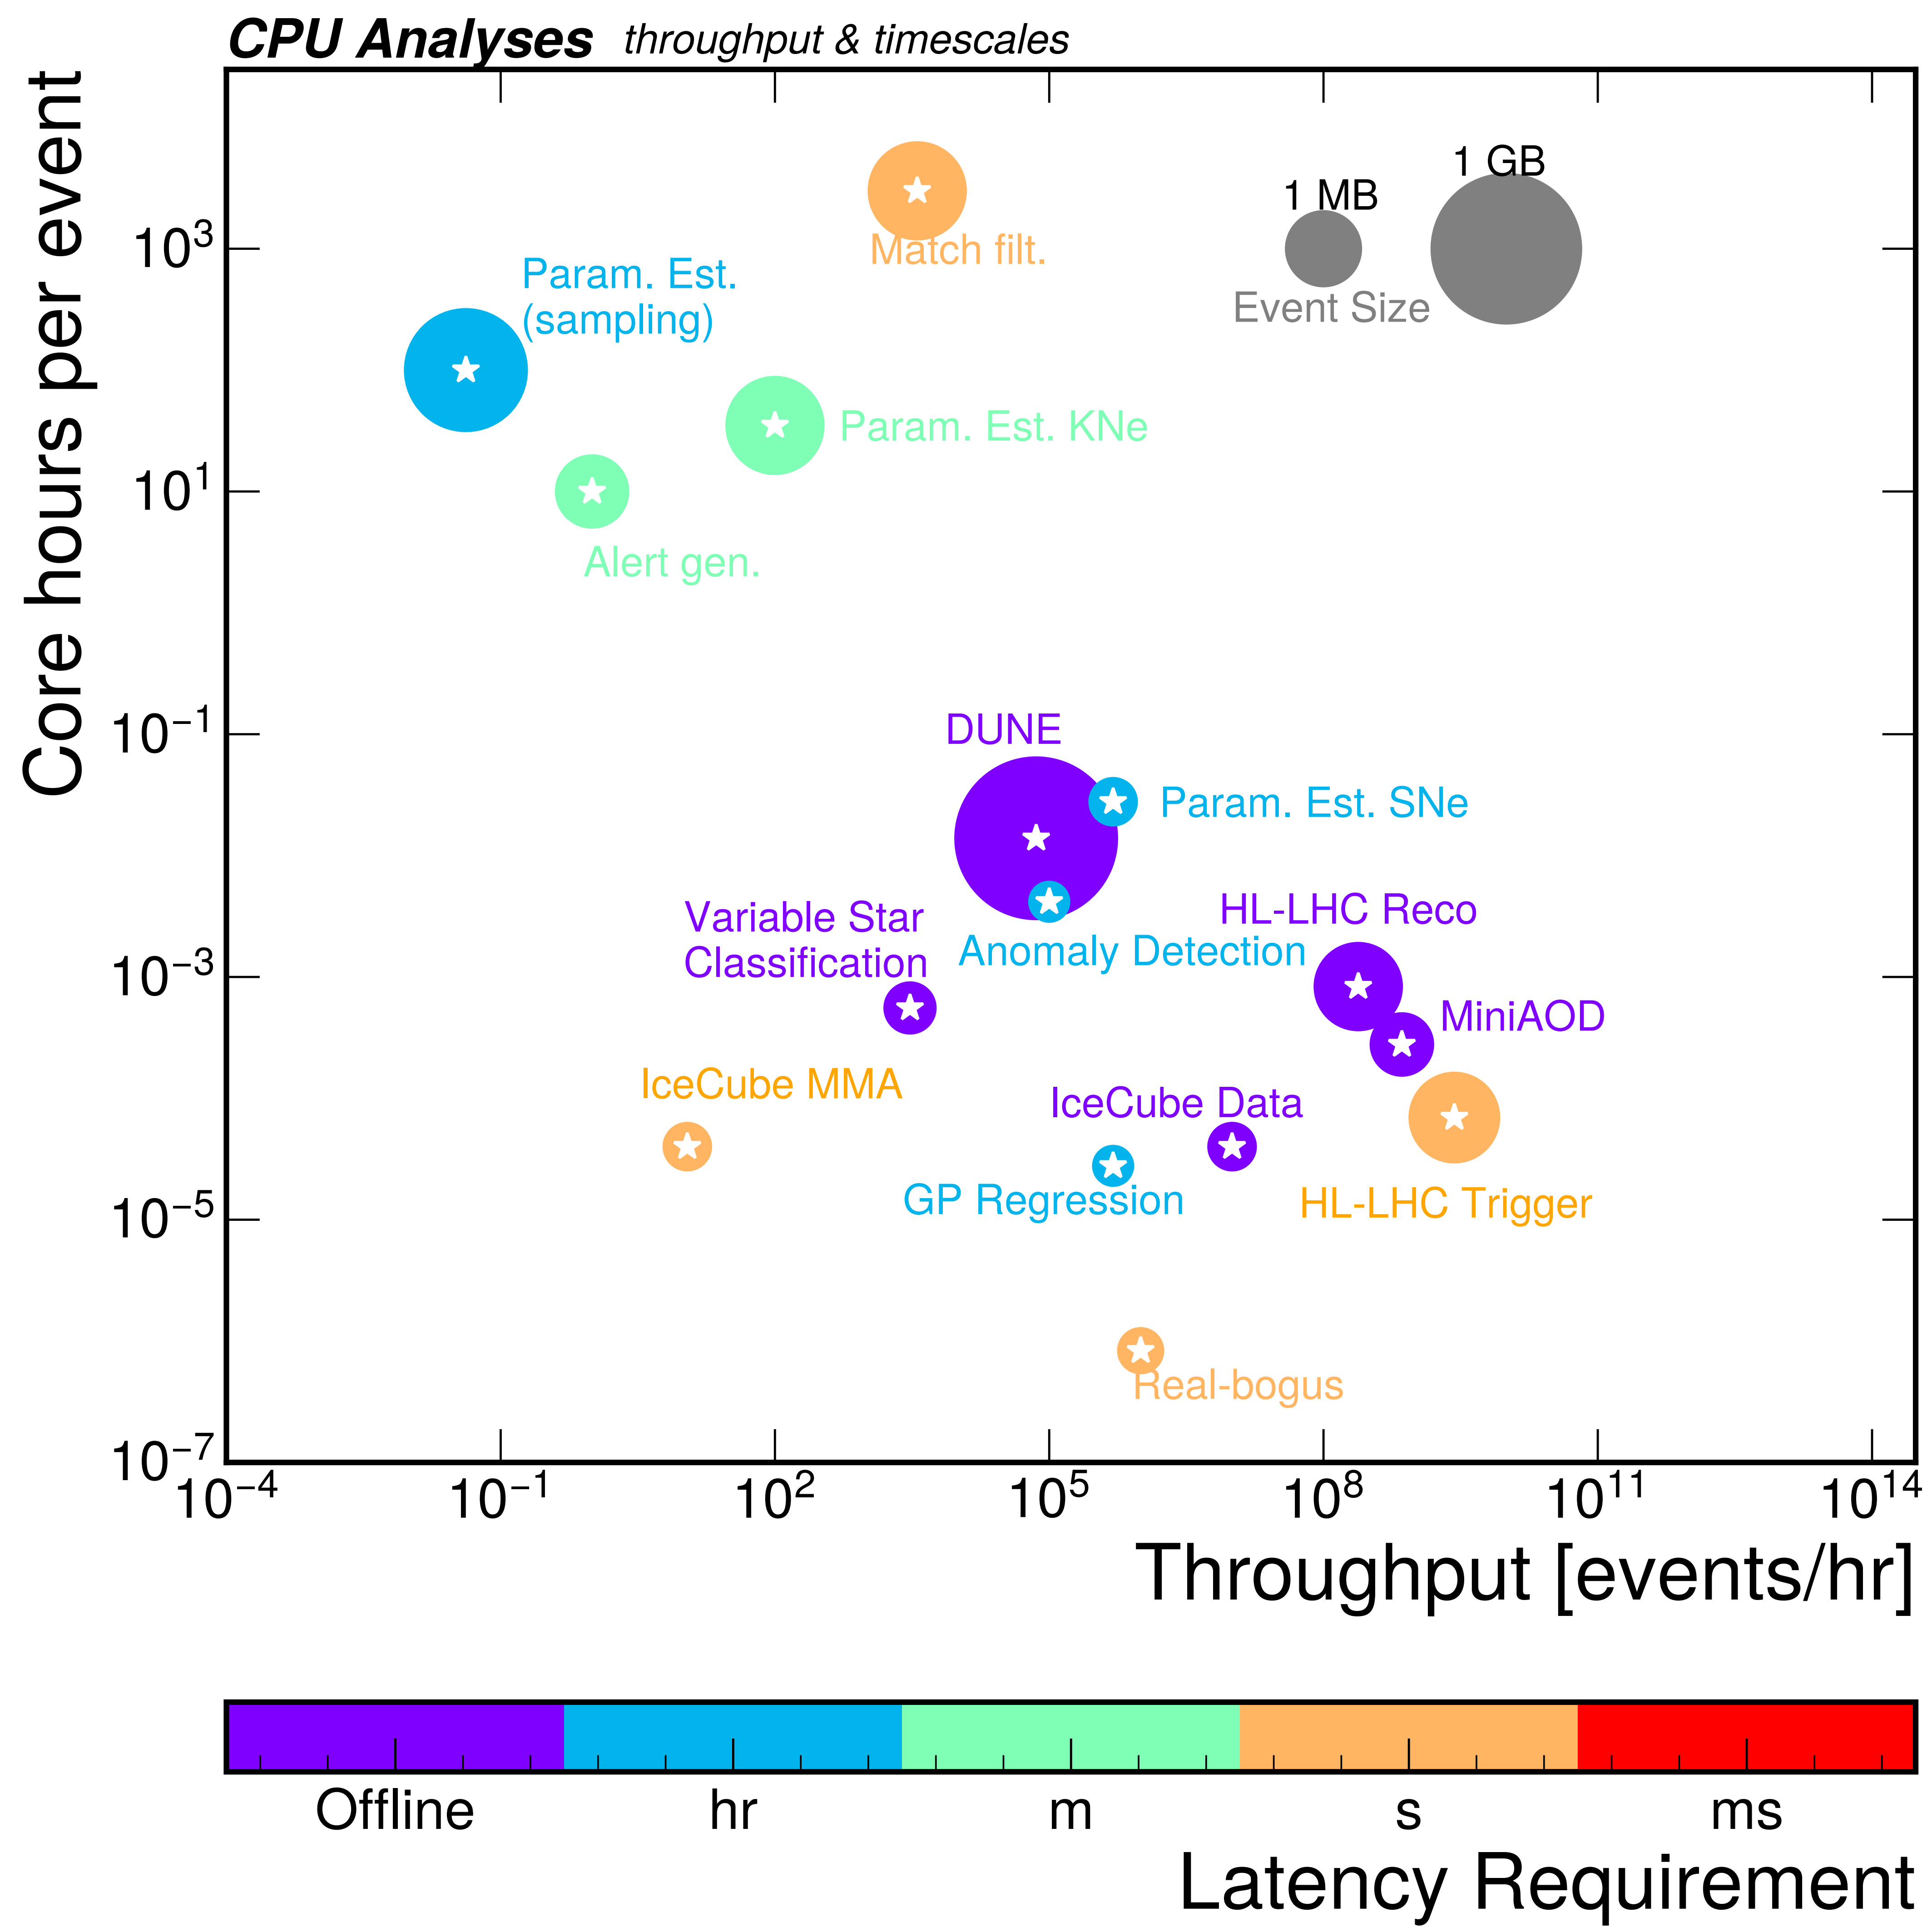

Computing challenges present in current and next generation HEP, astrophysics experiments and beyond.

White paper: "Applications of Deep Learning to physics workflows "

CMS members leading developments in general solutions, tools, computing paradigms.

Shared resources, computing infrastructure, network challenges

Summary

Upcoming High-Luminosity LHC run present many physics opportunities

Enhanced statistics with 10 times larger dataset, new ideas in looking for BSM deviations

Computing challenges at the LHC/High-Luminosity LHC (Run 4) starting in 2029

Collider operation conditions, physics needs

Changing computing hardware landscape

Various R&D activities

Big data science era. Shared challenges, tools and solutions that enable new computing paradigms and opportunities