Using GPUs at HPC centers for CMS

Lightning talk @ tCSC25

Nikita Shadskiy
October 7, 2025

Computing resource needs for HL-LHC

- **Significant increase** in **data** readout, processing and storage in upcoming years

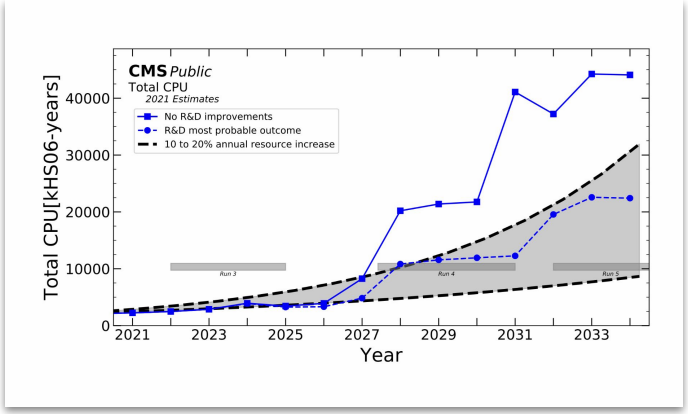

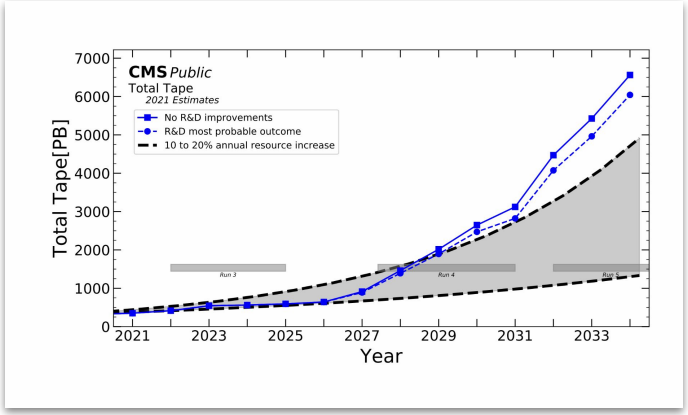

CMS-NOTE-2022-008

Computing resource needs for HL-LHC

- **Significant increase** in **data** readout, processing and storage in upcoming years
- Solutions for WLCG
  - Invest more money

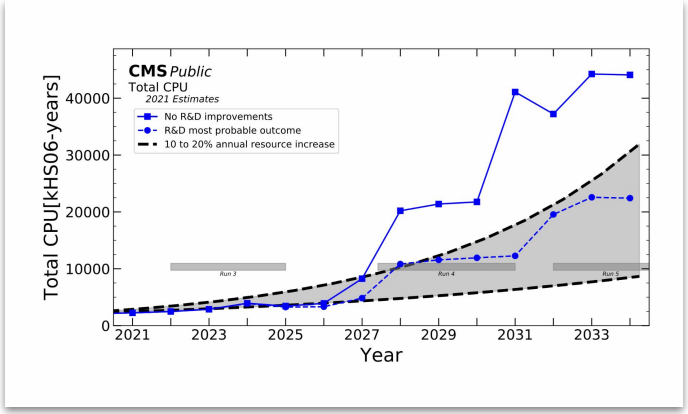

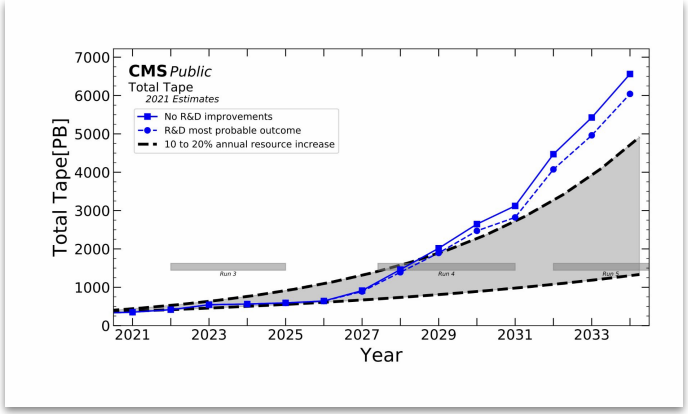

CMS-NOTE-2022-008

Computing resource needs for HL-LHC

- **Significant increase** in **data** readout, processing and storage in upcoming years
- Solutions for WLCG
  - Invest more money
  - Optimization of resource usage

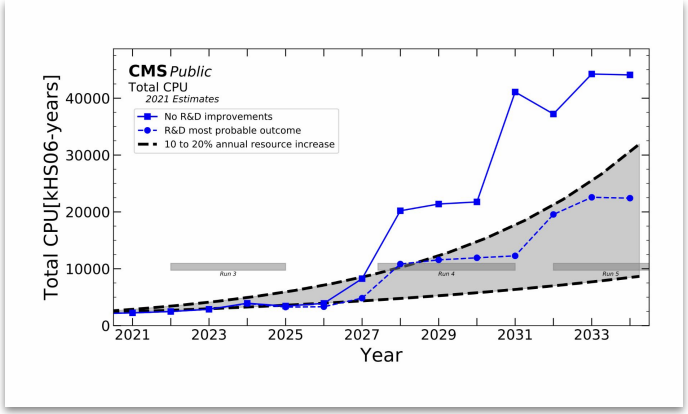

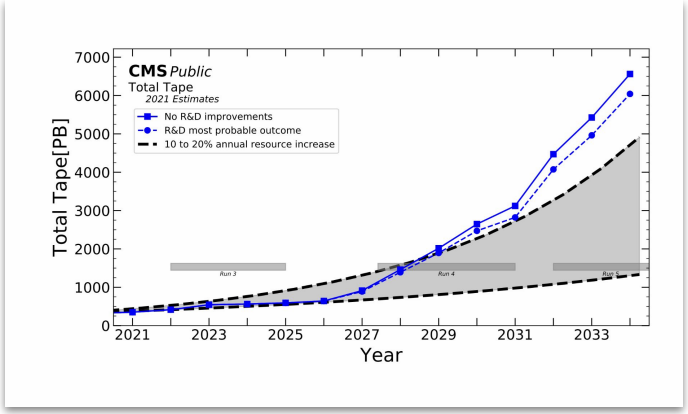

CMS-NOTE-2022-008

Computing resource needs for HL-LHC

- **Significant increase** in **data** readout, processing and storage in upcoming years
- Solutions for WLCG
  - Invest more money
  - Optimization of resource usage
  - Look for already existing computing infrastructure → **HPC**

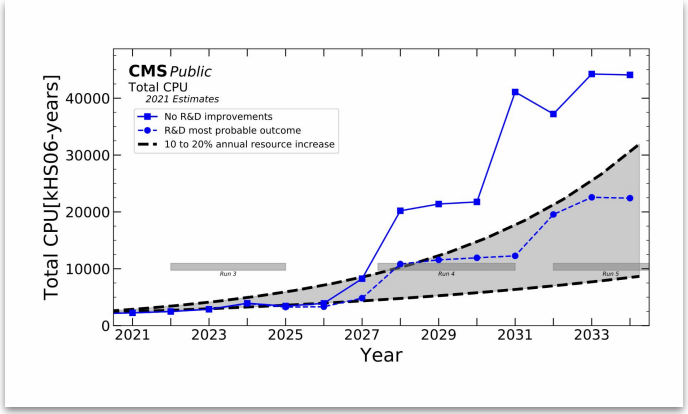

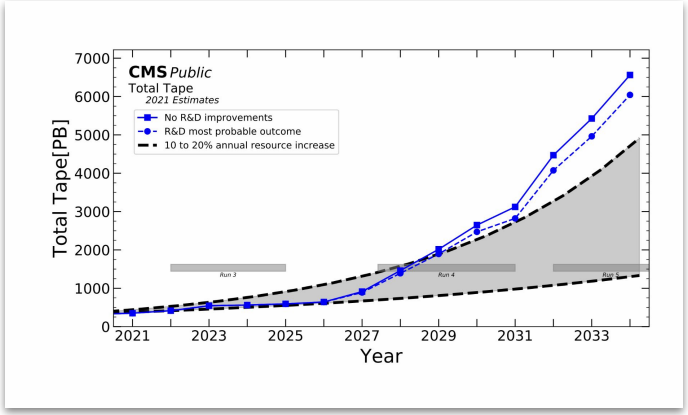

CMS-NOTE-2022-008

Computing resource needs for HL-LHC

- **Significant increase** in **data** readout, processing and storage in upcoming years
- Solutions for WLCG
  - Invest more money
  - Optimization of resource usage
  - Look for already existing computing infrastructure → **HPC**
  - Use other hardware → **GPUs**

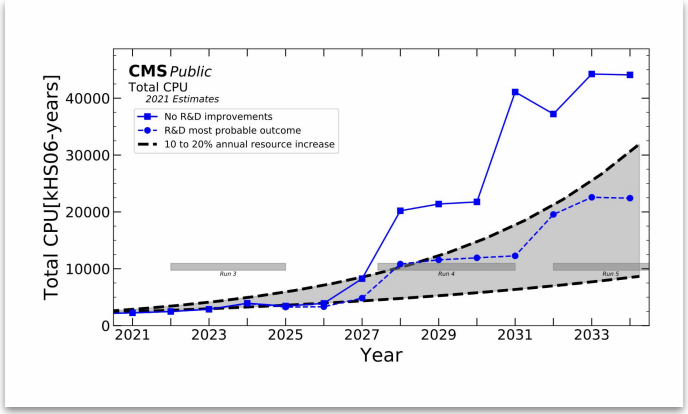

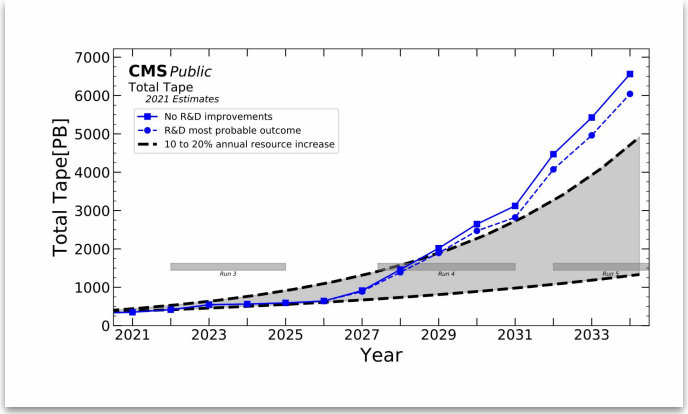

CMS-NOTE-2022-008

HoreKa @ KIT

- **HoreKa** is part of 9 **HPC center**s in Germany

  → National High Performance Computing Alliance (NHR)

- Purpose: **computing** infrastructure for **scientific research**

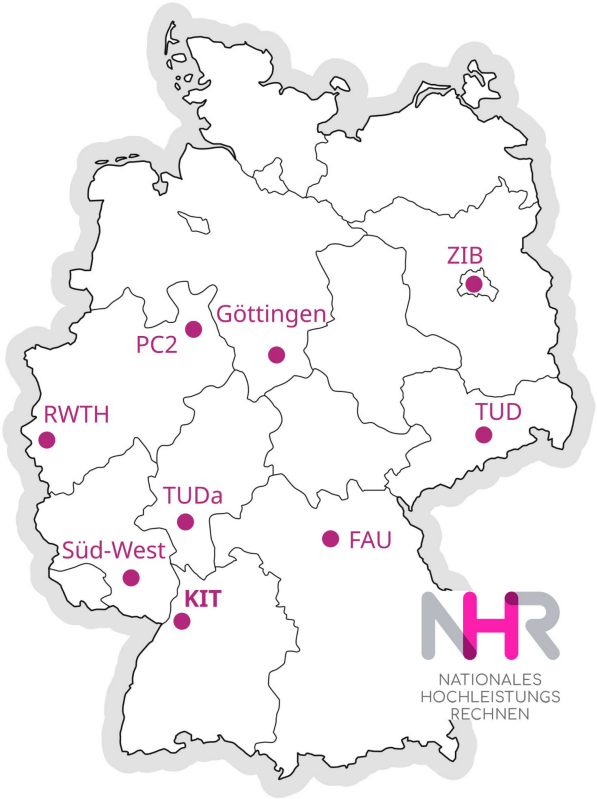

HoreKa @ KIT

- **HoreKa** is part of 9 **HPC center**s in Germany

  → National High Performance Computing Alliance (NHR)

- Purpose: **computing** infrastructure for **scientific research**
- HoreKa has multiple partitions:
  - **HoreKa Blue**: 610x CPU nodes (76x Cores)
  - **HoreKa Green**: 167x GPU nodes (4x A100)
  - **HoreKa Teal**: 22x GPU nodes (4x H100)
  - New since August 2025:

    HoreKa Ruby: 13x GPU nodes (4x H200)

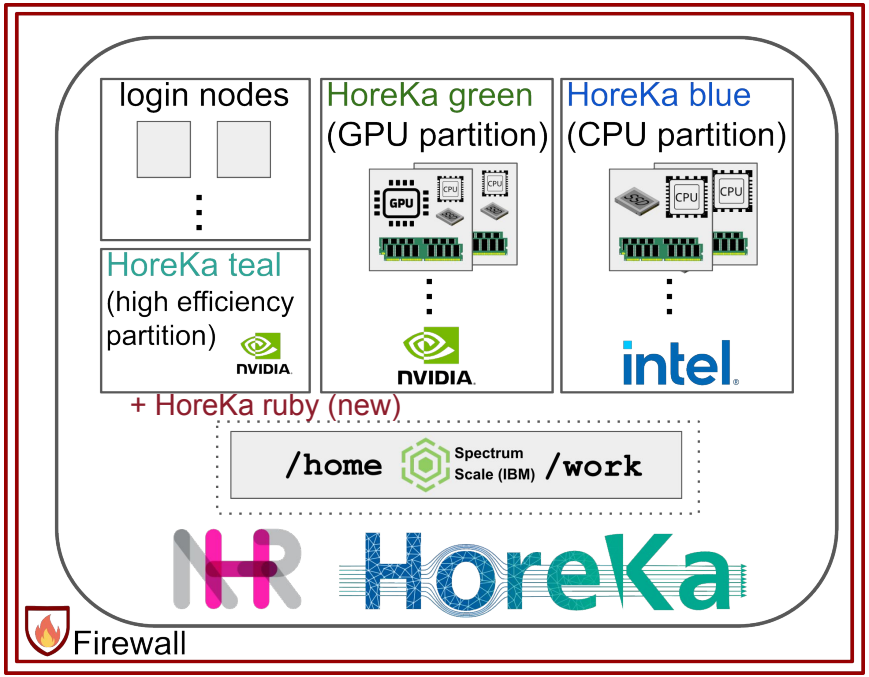

Source: Dr. Robin Hofsäss

HPC as opportunistic resource

- **GridKa** (Tier1) orchestrates **opportunistic resources** (CPUs) using
  - **COBalD**: balancing **resource allocation** and **utilization**
  - **TARDIS**: manager for **integration** of **heterogeneous resources**

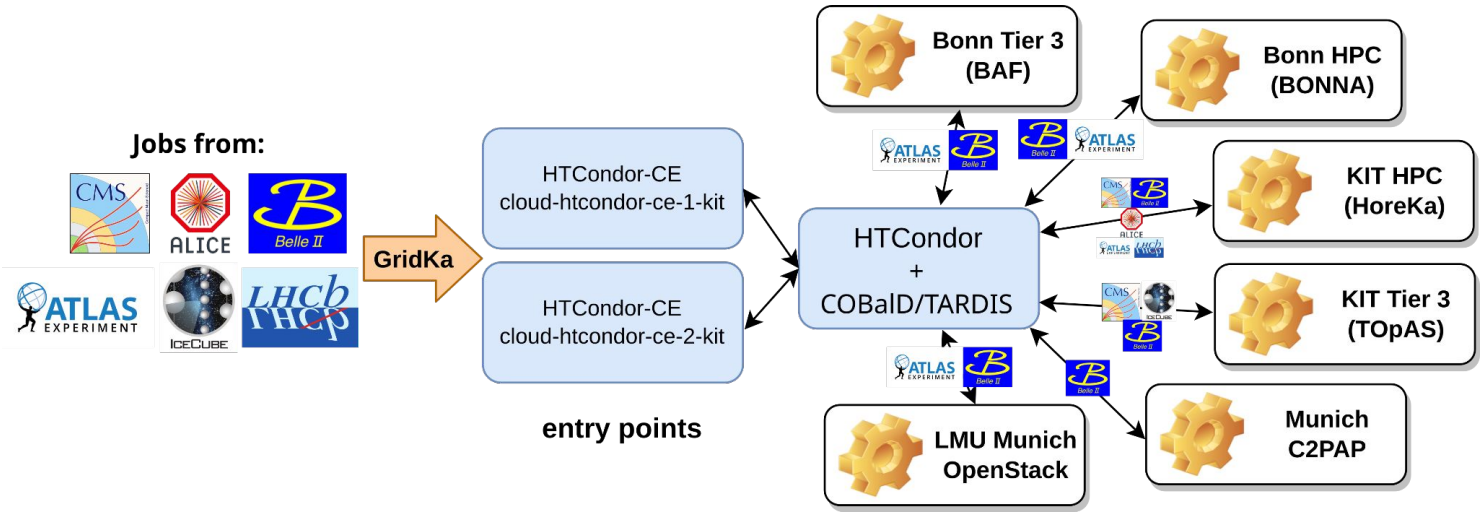

Source: Dr. Manuel Giffels

HPC as opportunistic resource

- **GridKa** (Tier1) orchestrates **opportunistic resources** (CPUs) using
  - **COBalD**: balancing **resource allocation** and **utilization**
  - **TARDIS**: manager for **integration** of **heterogeneous resources**
- Now also **GPUs** integrated!

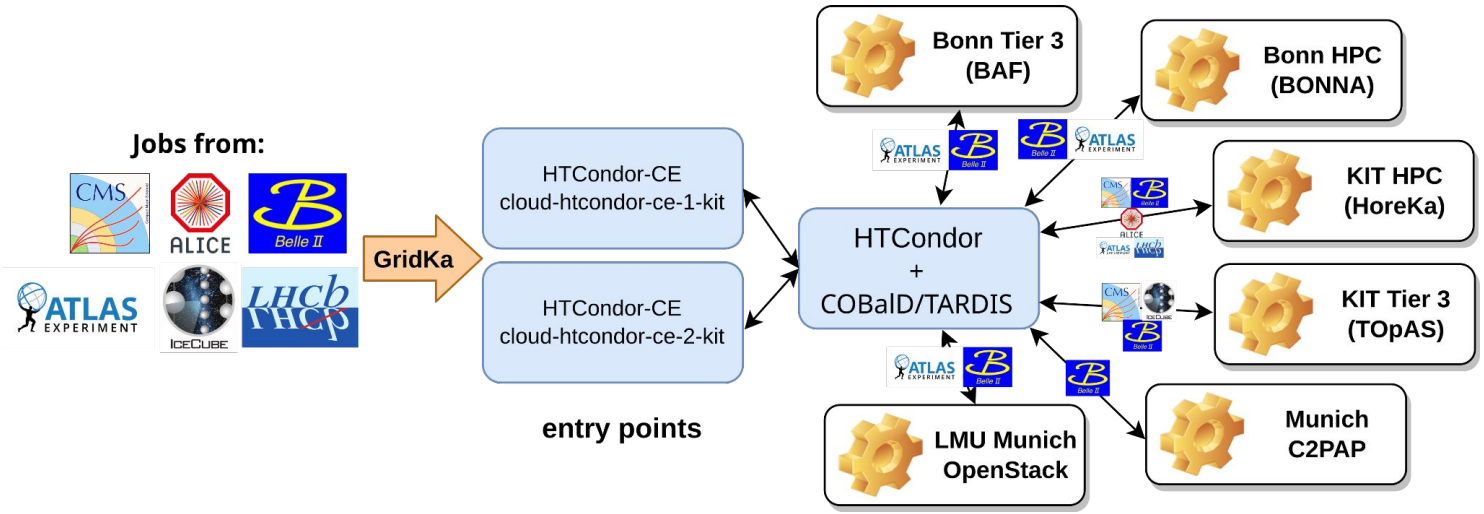

Source: Dr. Manuel Giffels

Testing the system (GPUs)

- **Tests** with part of available **HoreKa** resources → up to 8 GPUs
- **Centrally** submitted **CMS** and **ATLAS workflows** ran on HoreKa
- **User analysis** jobs ran on HoreKa
  - e.g. submitting via CMS Connect

One small CMS workflow for Simulation partially using GPUs

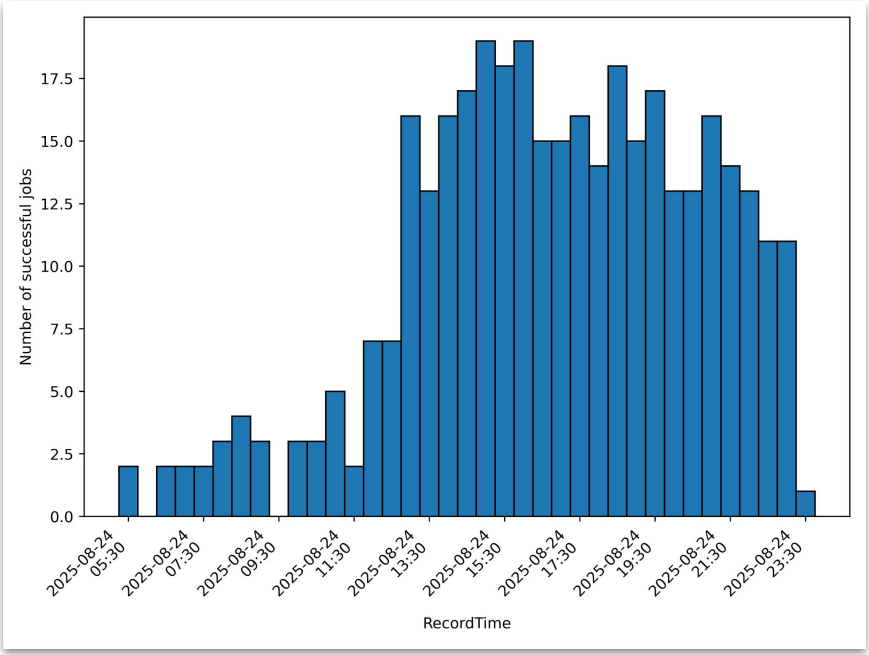

Testing the system (GPUs)

- **Tests** with part of available **HoreKa** resources → up to 8 GPUs
- **Centrally** submitted **CMS** and **ATLAS workflows** ran on HoreKa
- **User analysis** jobs ran on HoreKa
  - e.g. submitting via CMS Connect

- **But**: None of the GPU workflows are officially in production, yet (at least in CMS)
- **Developments** are **happening**!

  e.g. Daniele's talk MadGraph for GPUs

One small CMS workflow for Simulation partially using GPUs

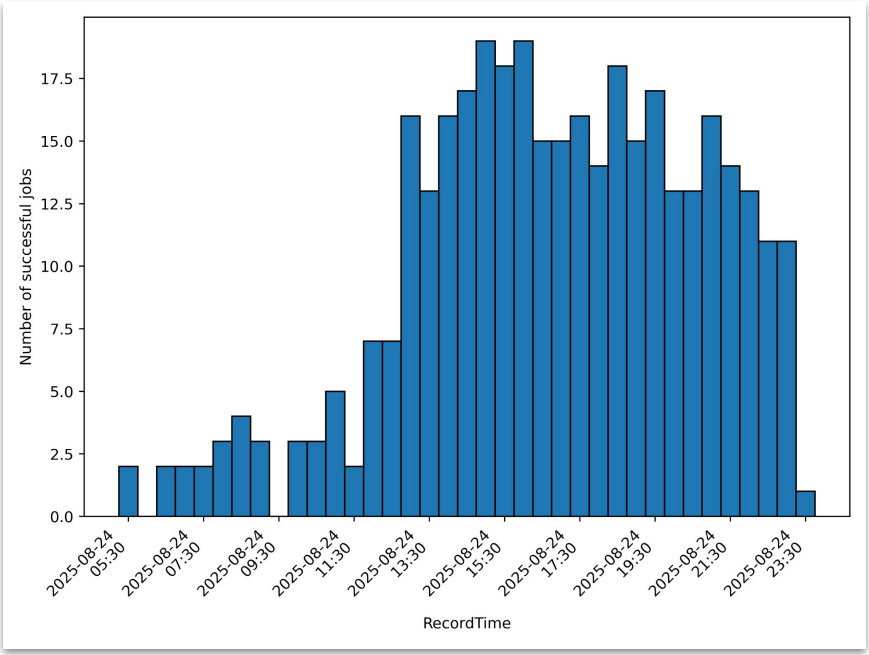

Testing the system (GPUs)

- **Tests** with part of available **HoreKa** resources → up to 8 GPUs
- **Centrally** submitted **CMS workflows** ran on HoreKa
- **User analysis** jobs ran on
  - e.g. submitting via C
- **But**: None of the GPU wor officially in production, yet (at least in CMS)
- **Developments** are **happening**!

  e.g. Daniele's talk MadGraph for GPUs

We should push to bring the developments to production.

The computing resources are (kind of) available.
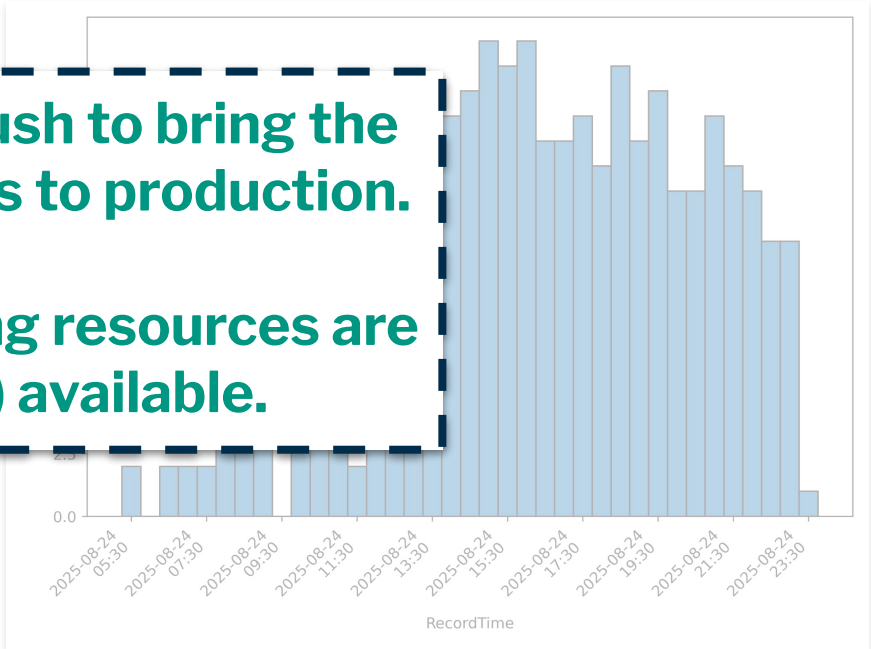
One small CMS workflow for Simulation partially using GPUs