

Integration of GPUs in WLCG

J. Flix

GDB

11/November/2023



Ciemat Centro de Investigaciones
Energéticas, Medioambientales
y Tecnológicas



Integration of GPUs in WLCG

The **WLCG MB** asked the GDB to bring back information from the LHC experiments on the **integration of GPUs in the WLCG computing infrastructure**

→ 2022 feedback was used for April 2022 C-RRB - **New feedback from 2023 has been collected:**

<https://twiki.cern.ch/twiki/bin/view/LCG/GPUsPlans>

QUESTIONS

1. Which is the current status of integration of GPUs for **offline** computing, at the software level?
2. Are you already using GPUs for **offline** computing activities? Please, comment on the experience.
3. If using these resources for offline, do you *account* for the GPU usage in some way?
4. Are there any future plans on GPU resource demands or plans for future utilization at sites?
Please, indicate mid-term and long-term plan (if available).
5. How many GPUs are needed for offline for the next two years?
6. Any other plans, e.g. on FPGAs?
7. Other comments or questions?

ALICE (main messages)

Uses a **common O2 software framework** for online and offline computing which is capable of offloading certain reconstruction steps to GPUs. Code can be run on different GPU backends

A substantial part of the **2022 and 2023 p-p data** sample was reconstructed on the EPN computing farm **using the farm's GPUs** (AMD MI50 and MI100)

Equivalence accounting between the GPU-equipped EPN nodes and CPU-only nodes with known HEPscore. This accounting method is possible and quite accurate, since **the same code is running on both and on the same data sample** → included in their **CRSG reports**

Currently **optimizing the use of various GPU models** through the resources made available to ALICE in various labs and computing centres. Once the optimization process is in sufficiently advanced stage, **ALICE may ask for a larger GPU deployment at WLCG centres** to run offline reconstruction on larger scale with GPUs. In addition, investigating the possibility **to use HPCs equipped with GPUs** for the same offline reconstruction tasks

ATLAS (main messages)

Several **R&D projects ongoing** investigating GPU usage for the HL-LHC era. The Athena offline software supports the integration of GPUs.

GPUs are already **integrated into offline computing** and are used for **limited analysis use cases**, e.g. ML training. Several sites have made small-scale GPU resources available to all ATLAS users through PanDA

GPU usage is measured **in the same way as CPU** usage, but considered **unpledged**. It is good to **start now developing a pledge framework** within WLCG for possible future use cases

Towards the **end of Run 3**, prototypes of event generation, detector simulation, and/or reconstruction software able to use GPUs may become available, so **sufficient resources should be available for testing and validation**. **Large-scale GPU deployment** at sites would be needed **by the start of HL-LHC** (~2028), assuming the offline software has been developed to use them

CMS (main messages)

All the components for the **usage of NVidia GPUs** are **present in the CMSSW** stack to exploit the GPU equipped HLT farm for Run 3 (the online and offline software release is the same software). Submission of **workloads on Grid worker nodes** equipped with NVidia GPUs is **fully supported**, both for user and production jobs

The transparent integration of GPU resources into the CMS Global Pool is being improved in **close collaboration with the HTCondor development team**

GPU usage is **opportunistic and not accounted**: HLT reconstruction code, machine learning studies, and user jobs. A strategic goal of CMS is to **efficiently use all of the resources made available** by WLCG sites

CMS will follow **WLCG guidance** for how to **include GPUs as part of our resource request** when CMS has a broader set of applications commissioned and sees that using GPUs would be cost effective. The goal is offloading ~10% of the offline reconstruction to GPUs by the end of 2023

LHCb (main messages)

The Allen framework is used in Run3 to execute on the **event building farm GPUs the first HLT stage (HLT1)**. Offline GPUs could be used to emulate the HLT1 in simulation through Allen

A few LHCb analysts use GPUs, but they use resources on institute clusters which are outside WLCG. **Actual usage is unknown. Fast simulations** use machine-learning techniques currently implemented on CPUs, but they might be ported on GPUs in future. LHCb might be using opportunistically the GPUs on the **online event builder farm when not in use for data taking**

No established plans on GPU resource demands or for future utilization at sites. Some discussion have started in the context of GPU usage for physics analysis, in the context of analysis facilities. **November 2023 workshop** to understand how to use GPUs/accelerators also in later processing stages, in the context of future Run4/Run5 LHCb Upgrades

LHCb is totally **dominated by simulation** (using above 90% of compute resources). **R&D towards porting (parts of) simulation on GPUs**, provided they can be advantageous in terms of performance and required resources

(main) conclusions derived from answers

- Outside HLT farms, **GPU usage is marginal**. No short-term plans to include GPUs at scale from WLCG sites (in general), but **sites offering GPUs help ongoing R&D activities**
- **FPGAs only for online**, no plans for offline yet
- No benchmarks for GPU atm, which affect the accounting → **Benchmark WG on it!**
- Good to start now developing a **GPU pledge framework** within WLCG (related to previous item)
- All of these resources treated as **opportunistic**, for the moment, waiting for guidance from WLCG, though the usage is not yet at scale
- The survey will be **conducted again in 1 year**, to know what's changed