Environmental Sustainability at LHCb

WLCG Environmental Sustainability Workshop
12 December 2024

Henryk Giemza
On behalf of the LHCb collaboration

Goal

Reduction of the environmental impact of the LHCb collaboration

Easy to say but how can we achieve it?

What can LHCb do? Where can it help?

- Identify **key areas**, assess their environmental impact and set focus on the major issues
- Measure, monitor and visualize
- Give and receive **feedback**
- Apply standards and good practices, define policies, request pro-environment changes
- Redesign and **implement** new eco-friendly software and services
- Balance CO_2 emission versus job execution rate and/or speed
- Budget CO_2 emission
- Train and educate
- Join forces with new and existing initiatives

Key areas

- DIRAC
  - Services and Agents
  - Monitoring and visualization
  - Job distribution and execution
- Run-time and simulation software
- Hardware
- Facility
- People

Services and Agents (DIRAC)

The future of DIRAC is DiracX [1]. DiracX is going to be a complete rewrite of DIRAC, and so it is a chance to incorporate solutions and technologies which will benefit the environment.

- Include battle-tested off-the-shelf technologies, as usually they are optimized and efficient solutions with community which updates and supports it
- Agents can be run as **event-driven tasks**, what decreases the hardware resources as we scale them depending on the actual needs (no hard allocations, space for over-provisioning).

Monitoring and Visualization (DIRAC)

DIRAC is well known of its **accounting**, **monitoring** and **visualization** capabilities. It relays on MySQL and Elastic/OpenSearch databases and provides native (legacy) plotting services, as well as integration with Grafana.

Right now we have a robust solution in place that collects data essential to calculate the environmental impact (e.g. job execution time, job final status, execution site, worker node performance, etc.).

The missing elements (in the basic scenario), like CO_2 emissions and energy usage (or its estimation) over time, can **already be integrated** from existing sources.

Our current aim is to visualize and monitor environmental impact by calculating and/or estimating CO_2 emissions generated by our grid activities.

We will be happy to share this knowledge, as not every site is able to prepare such measurements and estimations itself.

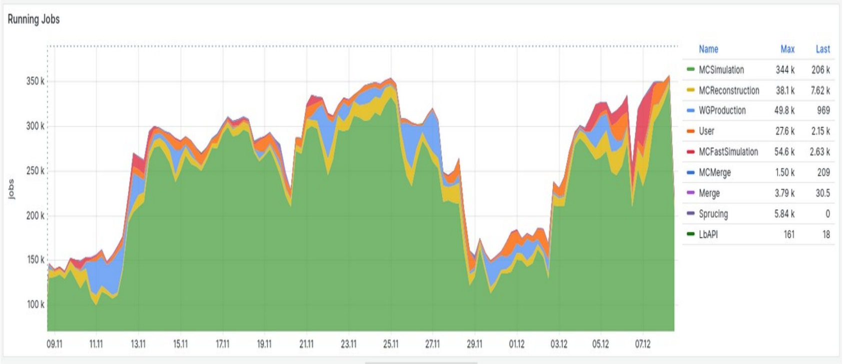

Job distribution and execution (DIRAC)

In LHCb, the majority of GHG emissions are generated during the execution of our jobs.

The jobs are running at **many different sites**, each of which has its own local **energy mix**, optimizations (e.g. waste heat recovery, downclocking), policies, etc.

Job failure rate and downtimes are also important factors when selecting an execution site for new jobs.

Any job execution above the pledged capacity at sites with a high carbon footprint should be managed carefully

- Q: Should we stop submitting jobs to allow the servers to be **turned off**?
- Q: Should we provide the sites with “a button” to quickly stop our jobs in an orderly manner?

Right now we do not take the aforementioned aspects into account when submitting jobs, but this is **going to change**.

Further details are in the next presentation:
Development and implementation of an environmentally friendly Workload Management System in DiracX

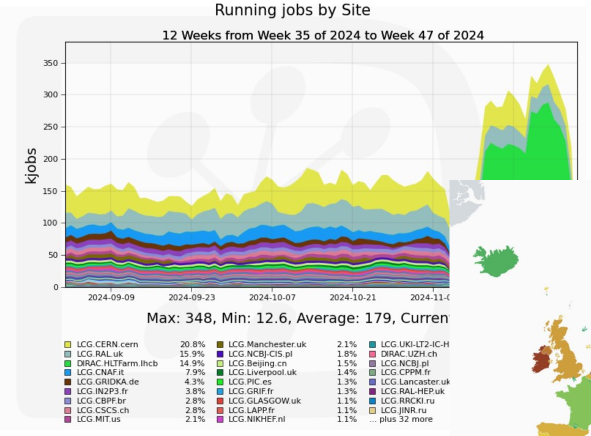

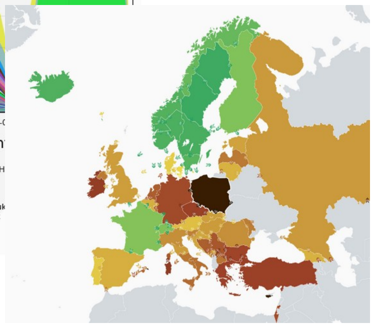

Run-time and simulation software

Updates in the real-time processing software done in 2021 have **decreased** *energy per trigger* [2].

For the simulation software, we have implemented a mechanism that selects the most optimal version from a wide range of validated options, based on the features available at given system.

Currently, we do not have measurements or estimates of the carbon footprint of our simulation software. However, updates to the monitoring system will provide feedback on the environmental impact of **different software versions** and may assist developers in incorporating more eco-friendly solutions.

Architecture	Energy per trigger (mJ)	Gain	Total gain
E5-2630-v4 Xeon			
Before SW optimization	39.9	1.0x	
w/Physics optimizations	21.0	1.9x	1.9x
w/SIMD optimizations	8.4	2.5x	4.8x
7502 EPYC			
w/SIMD optimizations	3.2	2.6x	12.5x
Event Building Node, NR			
1 GPU	3.1	1.03x	12.9x
2 GPUs	2.4	1.29x	16.6x
3 GPUs	2.1	1.15x	19.0x
Dedicated GPU machine			
4 x 2080 Ti + 2 Network Cards	2.8	1.14x	14.3x
5 x 2080 Ti + 3 Network Cards	2.5	1.12x	16.0x
Pure GPU machine			
8 x 2080 Ti + Onboard Network	2.1	1.15x	19.0x

An estimation of CO_2 emission, displayed in the results obtained from test pipelines during development, could also be beneficial.

The majority of CPU usage at LHCb (> 80%) is generated by the **MC simulation software**. While integrating new, relevant optimizations into the simulation algorithms may be challenging, recompiling the software to **new architectures** appears to be a reasonable goal, so we have already initiated works in this area.

Optimization of the algorithms does not always provide a long-term solution for the environment, as there is always a temptation to scale up the simulations and utilize all available resources. The optimization have to be a continuous effort.

Hardware

- Research in the area shows that using **non-x86_64**, energy efficient architectures are beneficial in context of the environment sustainability.
- DIRAC has been successfully tested in job submission to **ARM and GPU** based Computing Elements
  - Right now we are waiting for physics validation on ARM
  - RISCV processors will come next
- In case of GPUs, we have a successful story of integrating this type of resources in **real-time processing** at HLT1 [2].
  - Implemented solution has **eliminated** dedicated servers, as GPUs were incorporated into the Event Building (EB) machines
  - Further works will focus on upgrading HLT2 farm in a similar way (important target before the HL-LHC era)
- While LHCb does not have direct control over the hardware provided by the sites, adapting our computations to the new architectures will enable grid sites to select more **eco-friendly solutions**.

Facility (Online)

- LHCb targets **PUE < 1.1** within A2 envelope (10-35 degrees C) [3]
- A lot of effort was put to achieve this goal
  - Winning solution: **Modular Data Center**
  - Total power (design): **2.3 MW**
  - Indirect Free Air Cooling System
  - Regulation: outside and inside fan speeds, water pump speed

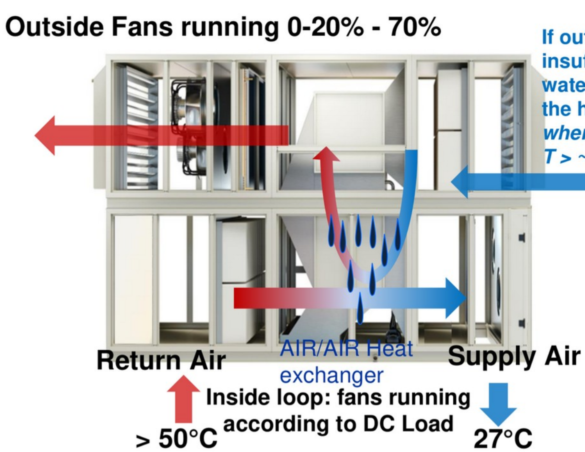

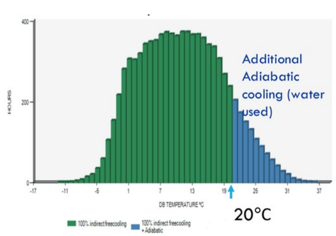

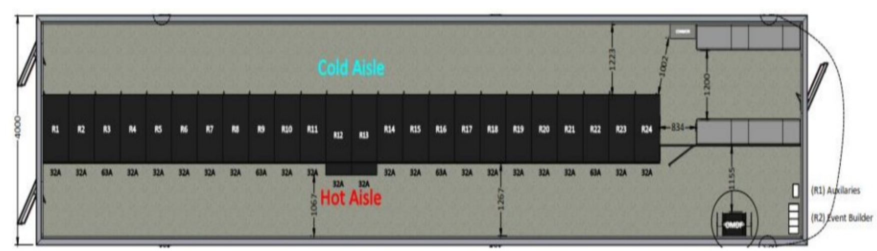

20kW (85%) and 40kW (15%) racks

Facility (Online)

- The goal has been **achieved**
  - PUE continuously **below** the target
- What next?
  - Reduction of **water-usage**
  - Optimization of cluster configuration
  - Preparation for the use of **DLC**
  - Investigation of energy-reuse possibilities
- Lessons learned will influence final choices for LHCb upgrade II for Run5 and Run6

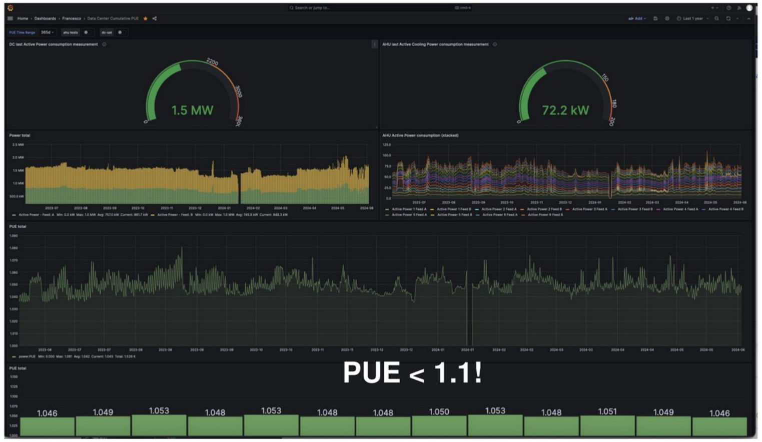

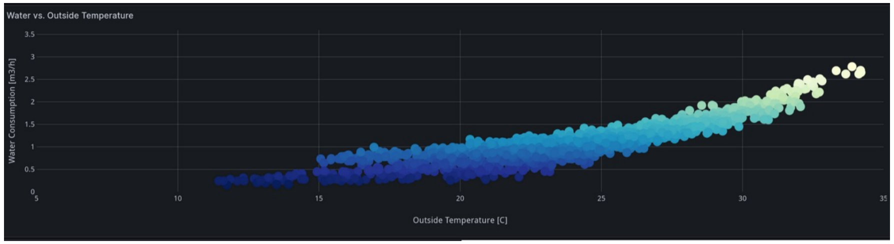

People

Human behavior is the key factor in achieving the **carbon neutrality**.

As a collaboration we can proactively support environmentally friendly practices by:

- Showing our members their individual, and the collaboration as a whole, influence on the CO_2 emissions
- Training developers and engineers to implement eco-friendly solutions
- Educating Shifters, Production Managers, Data Managers and other administrators about the importance of their work in the overall collaboration's impact on the environment
- Discussing the best hardware solutions and optimal scheduling policy with resource centers representatives

Thank you!

References

[1] DIRAC current, upcoming and planned capabilities and technologies, F. Stagni et al., https://doi.org/10.1051/epjconf/202429504018

[2] Evolution of the energy efficiency of LHCb's real-time processing, R. Aaij et al., https://cds.cern.ch/record/2773126/files/document.pdf

[3] Sustainable computing solutions: a case-study of the LHCb data-centre, N. Neufeld et al., ICHEP 2024