WLCG Sustainability - Next Steps

https://indico.cern.ch/event/1450885/timetable/

Sites

Centrally collect and expose information about how much power a site is using and from where the power is being taken. Ideally we should collect the HS23s of jobs and total power consumed in a job. But this is rather complicated to obtain from the job, while we could ask this information from a site. There is agreement that we ask the sites to report the power consumption per core, but we also have a mechanism to fix a default value in case the site does not report. This is an incentive for sites to update numbers every time the numbers improve. In the meanwhile we should investigate how a job-based measurement can be done or at least how the numbers above can be validated. For example, we know,or can collect, the hardware inventory at the sites and, based on this, we can estimate the average power consumption per HS23. Alternatively, we could look into the batch system capabilities to report back energy consumption and report it upstreams in the accounting tools. **This topic will require a dedicated discussion, but there is consensus that we would like to measure these two metrics and understand how the ratio improves.**

The same approach as above should be used for PUE. We should have a dual approach where a default value is assumed and the sites can update it to demonstrate improvement. WLCG wants to reward sites that make an effort to improve the infrastructure to be more environmentally friendly, while not penalising the ones that have no immediate means to do that. This is not just for PUE but also and more generally for carbon emissions. The more important aspect is to measure improvements, more than absolute values that might be difficult to interpret. We need to be careful about variations obtained by adding elements to the metrics as that could give the impression of degradation from the previous years.

The metrics above refer to the power consumed by a single core. We agree that **we need an additional study looking more holistically at the power consumption of the full workflow** (storage access, network, …) besides the single core. Particularly the impact of storage on energy consumption is about 50% of the CPU impact, and therefore not negligible.

At the end of the process we would like to move from power consumption to carbon footprint. The community is aware that the CO2 embedded in the equipment is significant. Given the complexity of assessing this contribution we suggest, as a first step, to collect the reliable data that sites have and make this available to all partners. This will allow, over time, to move towards a more uniform understanding of the carbon footprint.

These efforts will require setting up a team of experts looking at the aspects described above. **We agree to set up this team with a mandate along the lines of what was discussed**. The WLCG management board will define that mandate, noting that this area fits well the **mandate of HEPIX** and there are two working groups (TechWatch and Benchmarking) that cover most of the areas (see below)

Experiments

There is continuous progress in software performance, but it has been difficult to track how that impacts the energy and carbon footprint. CPU/Walltime efficiency is another aspect of the discussion.

The experiments occasionally release the status of the progress in individual components (e.g. reconstruction time / event) but folding that into a metric for general progress over a period of time is more complicated and can be done only periodically (every few years).

To make progress, we agree it will be important to **identify a limited number of relatively stable workflows and measure the progress against them**. In addition, we should **measure the amount of "waste" (e.g. failed jobs and others) and understand how that improves**. This information is produced regularly by the experiments and can be exposed, but it should be done coherently by the four experiments. The focus should be on production-like (organised) activities as today the impact of chaotic activities, like end user analysis, is moderate. In addition, we need to understand how the improvements in a given workflow impacts the total energy consumption and CO_2 footprint. **In consequence we need an estimate of how many events are processed/simulated with a given workflow.**

We also need a metric that looks at the amount of "physics" per consumed volume of energy. This item should be discussed further as there is not a clear mechanism to define "physics".

There is consensus that **we need a discussion among experiments to agree on how to implement the areas related to these topics. The WLCG management will discuss with the experiments how to best follow this up.**

Software performance and portability

It is very important for experiments to continue making progress on utilising heterogeneous architectures. The sites can support this process by providing the platforms that are needed for the development, integration and validation. To be able to use ARM resources for production, a set of ingredients was necessary: a benchmark to measure the performance, a platform for development and integration and enough resources for physics validation. The work on GPUs is at a less mature stage as the benefits are not currently fully understood for offline and the GPU market trends go in a direction that is not very suitable for e.g. MC generation, simulation and reconstruction. At the same time, there are examples where the GPU porting work really paid off, such as the ALICE reconstruction. There are GPU resources available for offline use cases at various WLCG sites and these should be sufficient for most of the stages above. Already at the level of benchmarking there are however very few workflows that today can be used. In addition, how to benchmark a machine with CPUs and GPUs and what to optimise is also not currently clear. We agree that **progress should be made in benchmarking for hybrid systems and that this will need input from the experiments by providing a representative number of workloads**.

In addition, **WLCG should identify the means to support the experiments adopting new technologies and facilitate the process**. For example WLCG could provide special access to resources with profiling tools enabled, noting that a level of expertise will be required to use them.

The aspects of training are very important, given the complexity in programming modern architectures, but also given the improvements that can be achieved in software running on standard CPU architectures. The role of the HSF in this area has been and is very important and **WLCG should continue supporting it.** There might be future funding calls from e.g. the EU for software efficiency and environmental sustainability that could be leveraged.

The HSF started a process to affiliate software projects with different levels depending on their status of maturity. At the moment it is a lightweight process looking mostly at the management of the project. It is not straightforward to enhance the affiliation process to take into account sustainability and software performance but this will be discussed in the HSF.

There are several categories of users with different levels of expertise. The scientists at the sites can provide bottom-up effort based on first hand experience accumulated by using analysis software. Explaining the impact of software improvements on environmental sustainability to these users, will serve as a motivation.

We also note that a very large fraction of the compute cycles are consumed by a small number of large packages that are used by several experiments. Improvements of these packages have the biggest impact. We encourage the community to support efforts to further improve these critical building blocks.

Miscellanea

We discussed mostly about data processing but less about storage. That also includes data replication models. This is an interesting area to look at, to compare how different replication models and storage solutions impact the carbon footprint

Next Steps

The following actions have been identified to advance WLCG's sustainability efforts:

1) Benchmarking Compute, Storage, and Networking:

WLCG MB will discuss with HEPIX the possibility of reviewing the mandate of the Benchmarking WG to expand its scope beyond CPU benchmarking to include storage and potentially networking, which are significant contributors to energy consumption. Collaboration with experiments will ensure realistic workflows and metrics are integrated into this effort.

2) Metrics: Physics per Unit of Energy:

A focused discussion will be initiated among the experiments to define a metric for "physics per consumed unit of energy" and its implications for the evolution of computing models, ensuring alignment with sustainability goals.

3) Facilitating the Adoption of Heterogeneous Architectures:

WLCG will identify ways to support experiments in adopting new technologies, including GPUs and other heterogeneous architectures. This includes providing access to resources for development, profiling, and validation, as well as offering tools and expertise to optimise their usage.