

# A holistic study of the WLCG energy needs for the LHC scientific program

David Britton<sup>1</sup>, Simone Campana<sup>2\*</sup>, and Bernd Panzer-Stradel<sup>2</sup>

<sup>1</sup> University of Glasgow, Glasgow G12 8QQ, United Kingdom

<sup>2</sup> CERN, Esplanade des Particules 1, 1211 Geneva 23, Switzerland

**Abstract.** The WLCG infrastructure provides the compute power and storage capacity needed by the Large Hadron Collider (LHC) experiments at CERN. The infrastructure is distributed across over 170 data centres in more than 40 countries. The amount of energy consumed by the WLCG to support the scientific program of the LHC experiments, and its evolution, depends on different factors: the luminosity of the LHC and its operating conditions; the data volume and the data complexity; the evolving computing models and the offline software of the experiments; the ongoing R&D program in preparation for the next LHC phase (HL-LHC); the evolution of computing hardware technology towards better energy efficiency; and the modernization of the facilities hosting the data centres to improve Power Usage Effectiveness. This contribution presents a study of the WLCG energy needs and their potential evolution during the future LHC program based on the factors mentioned above. Some of the information is obtained from the CERN experience but then extrapolated to the whole of WLCG. The study provides, therefore, a holistic view for the infrastructure rather than a detailed prediction at the level of the individual facilities. It presents a clear view of the trends and offers a model for more refined studies.

## 1 Introduction

The WLCG infrastructure provides the compute power and the storage capacity needed by the LHC experiments. The infrastructure is distributed across over 170 data centres in more than 40 countries. CERN, the Tier-0, provides roughly 20% of the WLCG resources. The annual CERN energy consumption during the LHC runs is approximately 1.25 TWh [1] and the IT infrastructure contributes up to 5% of that consumption. Therefore computing is a non-negligible contribution to the energy needs of the LHC program, which motivates this study.

A precise estimate of the energy consumption in the WLCG infrastructure is outside the scope of this document. Each WLCG site is autonomous in the procurement of hardware; in handling the hardware lifecycle; and in defining the way the services are deployed and operated. In addition, each data centre is unique in its Power Usage Effectiveness (PUE)

---

\* Corresponding author: Simone.Campana@cern.ch

and the way it varies with external conditions. We can therefore only perform a holistic study rather than specific measurements and predictions. Finally, the resource needs of the experiments for the High Luminosity phase of the LHC (HL-LHC) have large uncertainties at the moment, as does the evolution of the hardware technologies. We should not expect, therefore, to obtain precise predictions but rather general trends in different scenarios.

One must also clearly distinguish between power usage measured in kWh and the corresponding emission of CO<sub>2</sub>. The conversion factor between the two values is strongly country dependent [2] and can vary by more than a factor 10. We therefore do not intend to discuss here the impact on the CO<sub>2</sub> emission from the energy consumed by WLCG. However, there exist interesting studies about the impact on CO<sub>2</sub> in WLCG federations [3].

## 2 Factors impacting the WLCG energy needs

The amount of energy consumed in WLCG to support the scientific program of the LHC experiments and its evolution depends on different factors:

- In the future, the experiments will require an increasing amount of computing resources to fulfil their scientific program. The HL-LHC program, in particular, expects to produce a factor five more luminosity (and therefore, data) than the previous runs combined, and more complex events. This implies more compute resources and more storage.
- The progress made with different computing R&D activities and particularly with offline software efficiency and performance, plays a key role defining the computing resource needs of the WLCG experiments. The evolution of the computing models and the software therefore impacts the energy needs.
- The hardware technology, particularly for compute, evolves in the direction of improving energy efficiency for the same processing power. At the same time, the lifetime of computing hardware is increasing and less modern hardware tends to remain online for longer than before, still providing reliable compute and storage capacity
- The facilities hosting the WLCG hardware are progressively being modernized to improve the Power Usage Effectiveness (PUE) as part of their renovation.

In the next sections we will analyse these different factors in some details.

### 2.1 LHC luminosity and run conditions

The LHC physics program is characterised by data taking periods (runs) interleaved with multi-year shutdowns of the accelerator complex to allow upgrades of the accelerator and the experiments. The integrated luminosity delivered by the LHC increases after every upgrade. In Run-1 (2009-2012) the LHC delivered 29 fb<sup>-1</sup> of integrated luminosity, in Run-2 (2015-2018) it delivered 156 fb<sup>-1</sup>. The current baseline for Run-3 (2022-2025), Run-4 (2029-2032) and Run-5 (2035-2038) is roughly to produce respectively 350, 900

and  $1200 \text{ fb}^{-1}$  of physics from proton-proton collisions at  $\sim 14 \text{ TeV}$ . Physics at HL-LHC (Run-4 and Run-5) is not only characterised by more data to store and process but also by an increased complexity of the physics events. In Run-3 the average pile-up (number of particle interactions per event) is about 60 at the ATLAS and CMS detectors; this number will increase to 140 in Run-4 and 200 in Run-5. More event complexity implies a larger data volume and longer processing time per event. The conditions used for this study are summarised in Table 1, which were the most up-to-date numbers at the time of the analysis. Note that those values represent a scenario with optimistic performance of the LHC machine and therefore are a conservative scenario in terms of computing planning. Official figures concerning the expected physics delivery of the LHC might differ and be generally lower particularly in terms of the integrated luminosity. The High Level Trigger (HLT) rate and the amount of simulated data to be produced are specific parameters to each experiment, but are partially driven by the values in the common parameters of the table.


**Table 1.** The LHC input parameters used in this study.

	Run-4 (2029-2032)	Run-5 (2035-2038)
ATLAS/CMS luminosity	$<270/\text{fb}$ ( $<135/\text{fb}$ in 2029)	$<340/\text{fb}$ ( $<170/\text{fb}$ in 2035)
ATLAS/CMS average pile -up	$<140$ ( $<70$ in 2029)	$<200$ ( $<100$ in 2035)
LHCb luminosity	15/fb	50/fb
Alice luminosity (pp)	100/pb	100/pb
Running time (pp)	6 M seconds	6 M seconds
Running time (ions)	1.2 M seconds	1.2 M seconds

## 2.2 ATLAS and CMS CPU need for HL-LHC

The WLCG data centres power consumption is generally driven by the CPU needs of the experiments. For example, at CERN, 70% of the data centre power is consumed for data processing (CPUs); 25% is consumed for storage (disk and tape), and only 5% by network equipment. In this holistic study we therefore focus on the energy needs for CPUs. The projected CPU needs of the ATLAS and CMS experiments at HL-LHC are documented respectively in [4] and [5]. Those estimates were initially produced in the scope of the 2021 HL-LHC computing review driven by the LHC scientific Committee (LHCC). They were reviewed in 2022 to reflect the changes in the HL-LHC schedule. There are no publicly available estimates from the LHCb and ALICE experiments at this point in time, and in any case the impact on computing resources for HL-LHC, particularly in Run-4, is considerably lower than ATLAS and CMS. For the purpose of this study, we will use the ATLAS and CMS information. The evolution of the CPU needs for the two experiments as a function of time is shown in Figure 1. Both ATLAS and CMS present two scenarios:

A more conservative software R&D scenario estimates the CPU needs assuming limited computing R&D for ATLAS and no R&D at all for CMS in the next years. In our study we will label this as the “Pessimistic Scenario”. A more aggressive R&D scenario estimates the CPU needs assuming a larger amount of R&D succeeds in the coming years. The amount differs for the two experiments and in some cases resources to perform this R&D have not been fully identified. We will label this as the “Optimistic Scenario”. The total volume of CPU capacity needed is the sum of the ATLAS and CMS requirements in each scenario.



**Figure 1.** The estimated WLCG CPU requirements in Millions of HS06 [6] of ATLAS and CMS, as a function of time. The two scenarios, “Optimistic” and “Pessimistic” highlight the potential impact of the current R&D programs.

### 2.3 Watts per HS06

To estimate the future energy needs for ATLAS and CMS computing we need to know the amount of power needed per HS06 and how that evolves with time. We considered a dual AMD 7302 processor with 4TB of SSD, 256 GB of memory and a 10Gbps NIC. In the hardware configuration at CERN at the time of this study we deploy four separate servers into a common chassis with common redundant power supply, to minimise the infrastructure overhead. This configuration has, per server, a performance value of 1040 HS06, an idle power value of 120 W, and a full-load value of up to 420 W. Factoring in an average CPU efficiency of the applications running in the CERN data centre, the power consumed per unit of compute is roughly 350W/kHS06. This is the value we use in this study for a processor of the current generation.

The underlying semiconductor manufacturing technology for processors is continuously improving and the feature size is shrinking. This implies an increase in performance of

the processor, or a reduction of the energy consumed or a compromise between the two. The energy consumption per unit of computation decreased by approximately 50% over the last five years. We assume this trend will continue in the next years. The hardware replacement strategy and policy of the data centres also plays a role, as a more frequent rotation allows us to benefit from more energy efficient technologies, but it obviously also has an implication on the cost. The WLCG data centres normally commit resources not older than five years, and this is the number we will use in this study.

## 2.4 Hardware technologies and trends

The 50% decrease every five years in energy consumption per unit of computation is not a gradual process. Table 2 shows the hardware deployed at the University of Glasgow, a relatively large WLCG data centre, between 2015 and 2022 and the main specifications in terms of performance and power consumption

**Table 2.** The specifications in terms of performance and power consumption for the hardware at the University of Glasgow data centre

Year	Glasgow Hardware	Loaded Power Watts	HS06	Watts per kHS06	KWh per kHS06-year
2015	2*Intel(R) Xeon(R) CPU E5-2640 v3 @ 2.60GHz	160	371	431	3776
2016	2*Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40GHz	210	342	613	5373
2017	2*Intel(R) Xeon(R) CPU E5-2630 v4 @ 2.20GHz	210	416	505	4422
2020	2*AMD EPYC 7452 32-Core Processor @ 2.30GHz	390	1766	221	1934
2021	2*AMD EPYC 7452 32-Core Processor @ 2.40GHz	420	1766	238	2083
2022	2*AMD EPYC 7513 32-Core Processor @ 2.60GHz	480	2112	227	1991

A step-change happened when the data centre started procuring a different technology while there was a more moderate improvement in other years. The question is then, what is going to be the next technology bringing considerable benefits in terms of energy efficiency? Advanced Rick Machine (ARM) chips have low power consumption and heat generation. They are used extensively in portable, battery-powered devices such as smartphones, laptops, and tablets. The LHC experiments have been looking into this technology since at least 2013 [7]. At the time of those early studies, whilst the power consumption of ARM chips was considerably lower than that of the X86 architecture, so was the performance. As the process of porting High Energy Physics (HEP) applications to non-X86 architectures is not trivial, the use of ARM remained at the proof-of-concept level for many years. However, over the last 10 years we have seen a very large performance improvement of mobile devices, and this has prompted the LHC experiments to reconsider ARM as a possible architecture. Software releases currently exist for many of the experiment workflows, and most have now been validated for use in production. GridPP compared [8] the performance and the power consumption of two machines, roughly similar in terms of specifications and prices:


- x86\_64: Single AMD EPYC 7003 series (Milan). CPU: AMD EPYC 7643 48C/96T @ 2.3GHz (TDP 300W). RAM: 256GB (16 x 16GB) DDR4 3200MHz. HDD: 3.84TB Samsung PM9A3 M.2 (2280)
- arm64: Single socket Ampere Altra Processor. CPU: ARM Q80-30 80 core 210W TDP processor. RAM: 256GB (16 x 16GB) DDR4 3200MHz HDD: 3.84TB Samsung PM9A3 M.2 (2280)

The HEP community implemented the HEPscore [9] benchmark suite to replace HS06. The advantage of HEPscore is the use of the High Energy Physics experiments workflows as applications in the benchmark, giving a more realistic estimate of the performance of the processors for the HEP needs. The workflows available at the time of this study in HEPscore were compiled and executed in the two architectures above. Detailed results are available in [8] but in summary: using hyperthreading on the x86 AMD processor gives gains both in terms of speed and energy efficiency. However, the gains are highly workload dependent. The ARM processor (that does not support hyperthreading) was significantly more energy efficient and generally a little faster than the hyperthreaded AMD. It also showed a strong dependency on the type of the workload. In summary ARM looks like a potential step-changing technology. It motivates the efforts in porting and validating the experiment software for many workloads. Several WLCG data centres plan to start procuring ARM processor in parallel to X86 already in 2023.

Up to 2023, the use of GPUs in WLCG has been very limited and focused on special applications such as Machine Learning. At the time of writing, there is no committed GPU capacity in WLCG, though non-negligible resources are available opportunistically at WLCG data centres. Several of the LHC experiments are, however, relying on a hybrid CPU-GPU system for data acquisition. This allows the potential gains in compute performance and energy efficiency to be studied when using GPUs for processing. CMS, for example, is deploying a hybrid CPU-GPU solution for the HLT farm. In this solution, by offloading 40% of the processing to the GPUs, CMS measured [10] a 70% gain in throughput (number of events processed per second) with respect to a pure CPU based HLT farm, and a 50% gain in terms of events processed per kWh. Porting HEP software to GPUs is a challenging task, particularly for the most CPU time consuming offline use-cases (e.g. event generation, simulation, offline reconstruction). CMS is already working to profit from accelerator-ready code for offline processing in LHC Run-3. The aim is to be able to offload 10% of the computations in offline reconstruction by the end of 2023. In addition, the ALICE experiment is using a hybrid CPU-GPU system for data processing. The system was designed for synchronous reconstruction of the data, but it has recently been used also for asynchronous reconstruction, demonstrating the potential for ALICE to use GPUs at Grid sites. In summary, GPUs offer another opportunity to reduce the energy needs of the LHC experiments. Porting the LHC offline software to GPUs however requires specialized expertise.

### 3 ATLAS and CMS CPU energy needs at HL-LHC

Combining the information from the sections above, we can estimate the power needs of ATLAS and CMS for CPUs at HL-LHC. The results are shown in Figure 2. The energy requirement peaks in 2036 (start of Run-5) when a factor of five higher than 2022 is needed in the pessimistic scenario, and 50% higher in the optimistic scenario. These estimates assume 50% improvement in hardware efficiency every 5 years; 5 years hardware lifecycle; and an average PUE of 1.45 for the facilities (see next section).



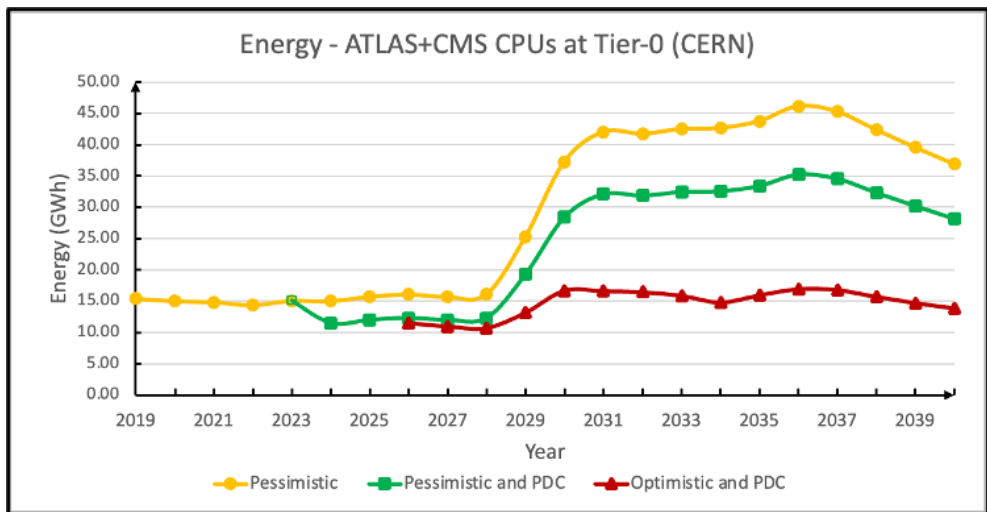
**Figure 2.** The estimated WLCG CPU energy needs of ATLAS and CMS in GWh as a function of the year. The two scenarios, “Optimistic” and “Pessimistic” highlight the potential impact of the current R&D program.

#### 3.1 Impact of the Power Usage Effectiveness

To estimate the power that will be needed by the LHC experiments in future we need to include the PUE factor. This accounts for the extra power that is needed by the data centre infrastructure on top of what is consumed by the processors and the servers; this is mainly for cooling. The PUE of the CERN data centre hosting the computing capacity for the LHC experiments is currently 1.45. This is an average over the year and refers to the data centre in Meyrin. Two extra containers are currently used at CERN and these have a better PUE. Worldwide, the average data centre PUE was 1.55 in 2022 [11]. It is difficult to quantify the PUE value for all WLCG facilities. In this study we used the value 1.45 as an average over WLCG.

In [11] we can also see how the modernisation of the facilities generally reduces the PUE. Modern data centres have a PUE below 1.4. CERN for example is building a new data centre at its Preveessin site and that is expected to be in use at some point in 2024. The new data centre will be able to provide up to 12 MW of power to the equipment, though

only 4 MW will be commissioned for 2024. The Prevešin Data Centre PUE is expected to be around 1.1. The new facility will also provide the possibility of heat recovery, of up to 3 MW out of the first 4 MW. The upgrade of WLCG facilities around the world is a continuous process and many plans are in place. It is difficult to project the impact of these plans in our estimates as some of them are very preliminary. For this study we consider the average WLCG PUE value 1.45 constant over the years. This is certainly a pessimistic assumption and, therefore, the estimates in Figure 2 should be seen as a conservative scenario. We can, however, estimate the effect of the Prevešin Data Centre on the CERN energy needed to provide processing capacity to ATLAS and CMS. The results are in Figure 3.



**Figure 3.** The estimated CERN CPU energy needs of ATLAS and CMS in GWh as a function of the year. The terms “Optimistic” and “Pessimistic” highlight the potential impact of the current R&D program. In addition, the “Pessimistic” scenario is presented both with and without the effect of the Prevešin Data Centre (PDC) becoming available in 2024. The “optimistic” scenario is presented only in the case of PDC being available in 2024.


For CERN CPUs, the introduction of the Prevešin Data Centre reduces the energy needs by 30% at the time of HL-LHC. One must note that the successful completion of the R&D program reduces the needs by an additional 50%.

### 3.2 Key Performance Indicators

In the 2019-2020 CERN environment report [12], the  $\text{GWh}/\text{fb}^{-1}$  Key Performance Indicator (KPI) was introduced to measure the improvements in energy consumption with respect to the amount of luminosity delivered by the LHC. It represents the amount of energy needed to **produce** a given amount of data. In our study we consider a similar indicator, but for us  $\text{GWh}/\text{fb}^{-1}$  represents the amount of energy needed to **analyse** the data.



To predict the trend of  $\text{GWh}/\text{fb}^{-1}$  for HL-LHC we use the results in Figure 2 and estimate the WLCG energy needs for CPUs averaged for each LHC run. The results are shown in Figure 4. The markers labelled “Energy” show the results in the pessimistic and optimistic R&D scenarios. The Y-axis scale for those markers is the one on the right-hand side of the plot. The average energy needs in Run-4 and Run-5 are roughly a factor two higher than in Run-2 in the pessimistic scenario. They are only 10% higher than in Run-2 in the optimistic scenario.



**Figure 4.** The “Energy” markers show the estimated average CPU energy needs in WLCG for the ATLAS and CMS experiments. The two scenarios, “optimistic” and “pessimistic” highlight the potential impact of the current R&D program. The scale those markers refer to is the one on the right Y-axis of the plot. The “Energy/Lumi” markers show the ratio between the energy needs as above and the integrated luminosity expected for a given run. The scale those markers refer to is the one on the left Y-axis.

The energy estimates per each run have also been divided by the integrated luminosity expected in that run (or collected, in the case of Run-1 and Run-2) to obtain  $\text{GWh}/\text{fb}^{-1}$ . We assume that in the shutdown period the CPUs are mostly used to continue processing the data from the previous run, which matches the experience we have so far after Run-1 and Run-2. The results are also shown in Figure 4 with the markers labelled “Energy/Lumi”. The Y-axis scale for those markers is the one on the left-hand side of the plot. Note that this is a logarithmic scale. The energy needs per  $\text{fb}^{-1}$  decrease by a factor ten between Run-1 and Run-5. In Run-5,  $\text{GWh}/\text{fb}^{-1}$  in the optimistic scenario is half compared to the pessimistic scenario.

## 4 Conclusions

In this study we estimated the energy needs of the WLCG facilities to support the CPU needed by the ATLAS and CMS experiments at CERN, focusing on HL-LHC. The study is holistic in the sense that does not claim to make accurate predictions, but rather look at the contributing factors and estimate the trends. The study shows that, generally, the energy needs in HEP computing can be kept under control, if not reduced, with a multi-prong approach. The modernization of the facilities is one pillar, as it goes in the direction of more energy efficiency. It is normally a major capital investment that requires multi-year planning. Improving the software and computing model of the experiments is a second pillar. It is a gradual process bringing early benefits. It allows potentially all members of our community to contribute in different ways depending on the level of their expertise. The improvement in the hardware technologies is a third pillar, associated with an appropriate hardware lifecycle strategy. It requires effort in software portability and, like the second pillar, it allows many members of the community to contribute. Each of these pillars is important, but the improvements in software and computing models are an area where the largest gains should be expected according to our study. Finally, in all scenarios, the amount of energy needed to analyse a “unit” of physics decrease over time. Or, in other words, the HEP community will generate more knowledge per kWh over time.

## References

1. <https://hse.cern/content/energy-management> (April 2023)
2. <https://ourworldindata.org/grapher/carbon-intensity-electricity> (April 2023)
3. D. Britton et al, *Reducing the carbon footprint of computing*, in Proceedings of the International Symposium of Grid and Cloud, ISGC, 19-31 March 2023, Taipei, Taiwan (to be published)
4. <https://cds.cern.ch/record/2802918>
5. <https://cds.cern.ch/record/2815292?ln=en>
6. <https://w3.hepiv.org/benchmarking.html> (April 2023)
7. D. Abdurachmanov et al 2015 *J. Phys.: Conf. Ser.* **664** 092007
8. R.J. Sobie et al, *HEPScore : a new CPU benchmark for the WLCG*, in Proceedings of the Computing in HEP Conference 2023, CHEP, 8-12 May 2023, Norfolk, United States (to be published)
9. A. Bocci et al 2023 *J. Phys.: Conf. Ser.* 2438 012058
10. <https://www.statista.com/statistics/1229367/data-center-average-annual-pue-worldwide/> (April 2023)
11. <https://hse.cern/environment-report-2019-2020/energy> (April 2023)