# A holistic study of the WLCG energy needs for HL-LHC

S. Campana, B. Panzer (CERN)

# Introduction

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

- **The LHC luminosity:** more luminosity (and therefore, data) and more complex event data => increased need in more compute resources and more storage.

- **The evolution of the computing models and the software:** the progress made with different computing R&D activities and particularly software efficiency and performance will play a key role defining the resource needs

- **The hardware technology**, modern hardware improves energy efficiency and also reliability

- **The facilities** hosting the WLCG hardware are progressively being modernized to improve the Power Usage Efficiency (PUE) as part of their renovation.

# LHC luminosity and parameters

https://espace.cern.ch/HiLumi/WP2/Wiki/HL-LHC%20Parameters.aspx

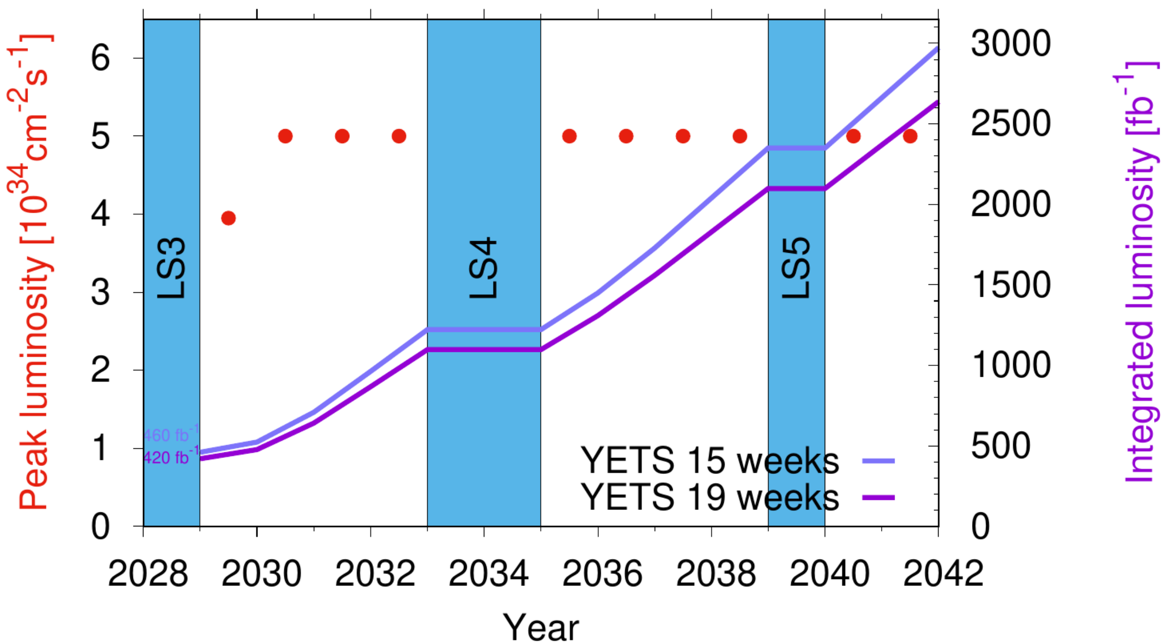


Those are the current expectations for the HL-LHC schedule and parameters

Those are not the conditions we used in this study. We used the conditions agreed with the LHC Physics Committee in early 2022

# LHC luminosity and parameters

The conditions we used for this study are in the table below

2024 and 2025 are expected ~ as 2023: < 100/fb luminosity, < 50 pileup, 6x10^6 seconds of pp time

| | Run-4 (2029-2032) | Run-5 (2035-2038) |
|---|---|---|
| ATLAS/CMS luminosity | <270/fb (<135/fb in 2029) | <340/fb (<170/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 |

# WLCG data centers power consumption

The pie chart shows the breakdown of the power consumption at the CERN data center

Most of the power is consumed for data processing (CPUs). Large part of the "services" are in fact CPUs

In this study we will focus on the energy needs for CPUs

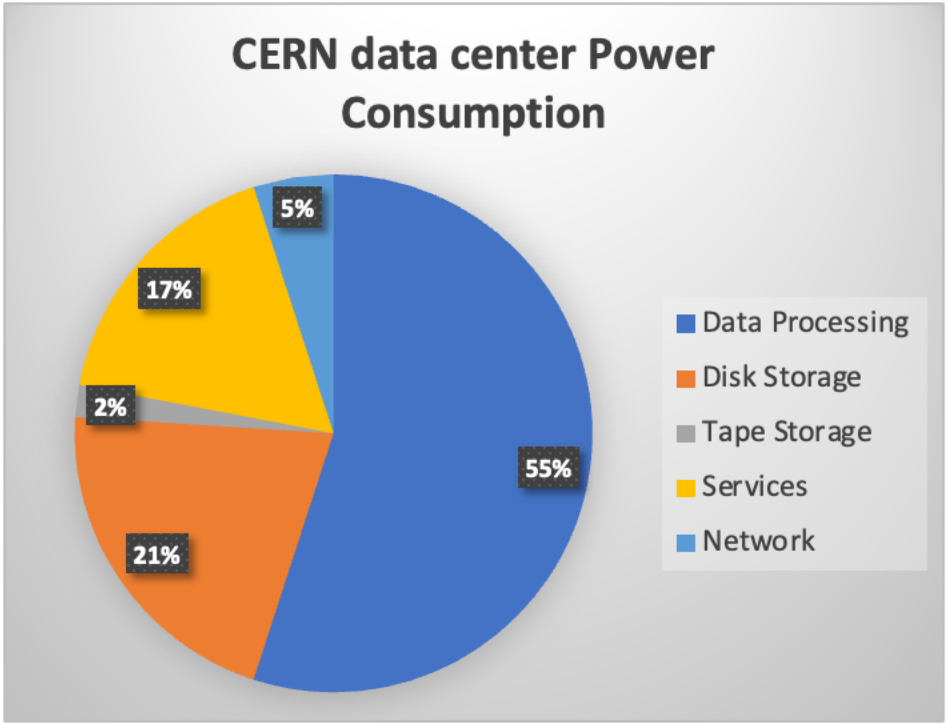

# ATLAS and CMS CPU needs for HL-LHC

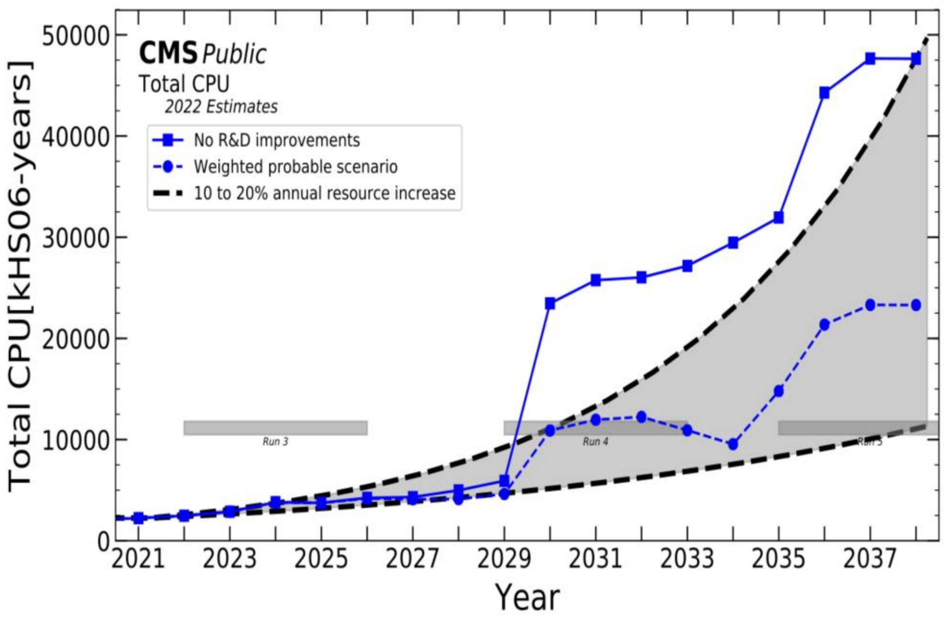


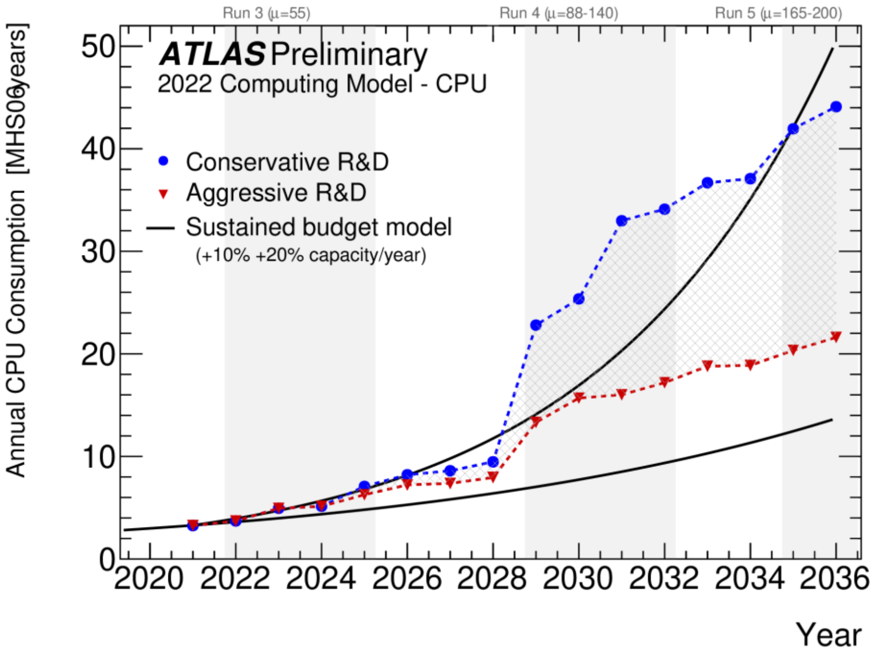


The projected CPU needs of ATLAS and CMS at HL-LHC. Estimates produced for the 2021 LHCC review and updated in 2022 to reflect the changes in the LHC schedule

ALICE and LHCb will require considerably less in Run-4, while no firm estimates are available for Run-5. We will consider ATLAS and CMS only in this study.

# ATLAS and CMS CPU needs for HL-LHC

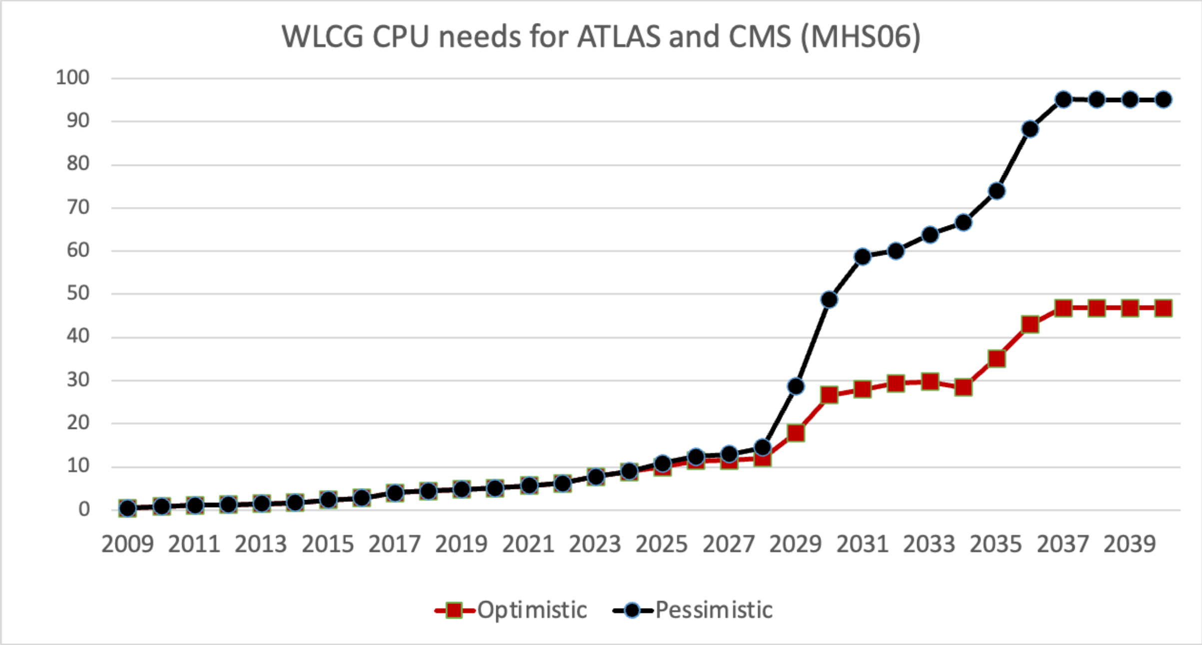


ATLAS (conservative R&D) + CMS (no R&D improvement) = **PESSIMISTIC** scenario

ATLAS (aggressive R&D) + CMS (weighted probable scenario) = **OPTIMISTIC** scenario

# HS06 into Watts

Consider a dual AMD 7302 processors, 4 TB SSD, 256 GB memory, 10 Gbit NIC

Incorporate 4 separate serves into a common chassis with common redundant power supply (minimize the infrastructure over-head)

The mentioned 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

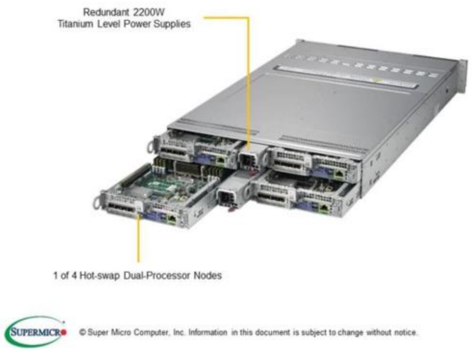


Consider 80% CPU efficiency (as an average in WLCG) => **350 W/kHS06**

**(this is the number we use for a processor of "today" in this study)**

# HS06 into Watts

The underlying semiconductor manufacturing technology for processors is continuously improving and the feature size is shrinking

- increase of performance or the reduction in energy usage or a compromise between the two

**Advertised PPA Improvements of New Process Technologies**
Data announced during conference calls, events, press briefings and press releases

| | TSMC | | | | | | |
|---|---|---|---|---|---|---|---|
| | N7 vs 16FF+ | N7 vs N10 | N7P vs N7 | N7+ vs N7 | N5 vs N7 | N5P vs N5 | N3 vs N5 |
| Power | -60% | <-40% | -10% | -15% | -30% | -10% | -25-30% |
| Performance | +30% | ? | +7% | +10% | +15% | +5% | +10-15% |
| Logic Area Reduction % (Density) | 70% | >37% | - | ~17% | 0.55x -45% (1.8x) | - | 0.58x -42% (1.7x) |
| Volume Manufacturing | | | | Q2 2019 | Q2 2020 | 2021 | H2 2022 |

CPU performance per energy usage (HS06/W) reduced by ~50% over the last 5 years. We assume this trend will continue in the next years

The hardware replacement strategy also plays a role: we assume a 5 years lifecycle

# Power Usage Efficiency

The overall power usage has to include the PUE (Power Usage Efficiency) factor

- extra power is needed to provide the cooling of equipment
- The PUE is temperature dependent: day-night and seasonal modulation

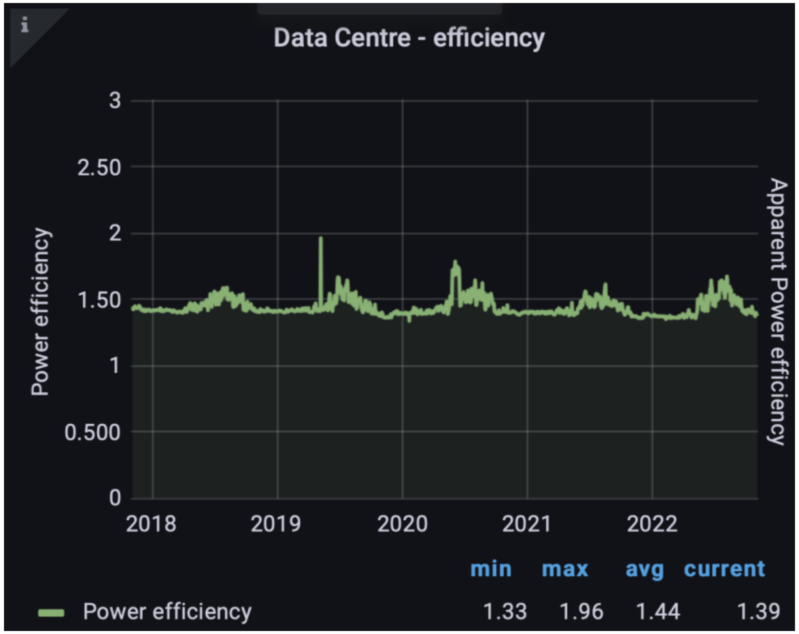


The average PUE for the CERN Meyrin (plot above) centre is about 1.45 over the year (range is 1.35 – 1.65)

It is hard to quantify this value for all WLCG facilities. We will use 1.45 for this study, constant over the years

# ATLAS and CMS CPU power needs for HL-LHC

This is power now
(GWh/year)

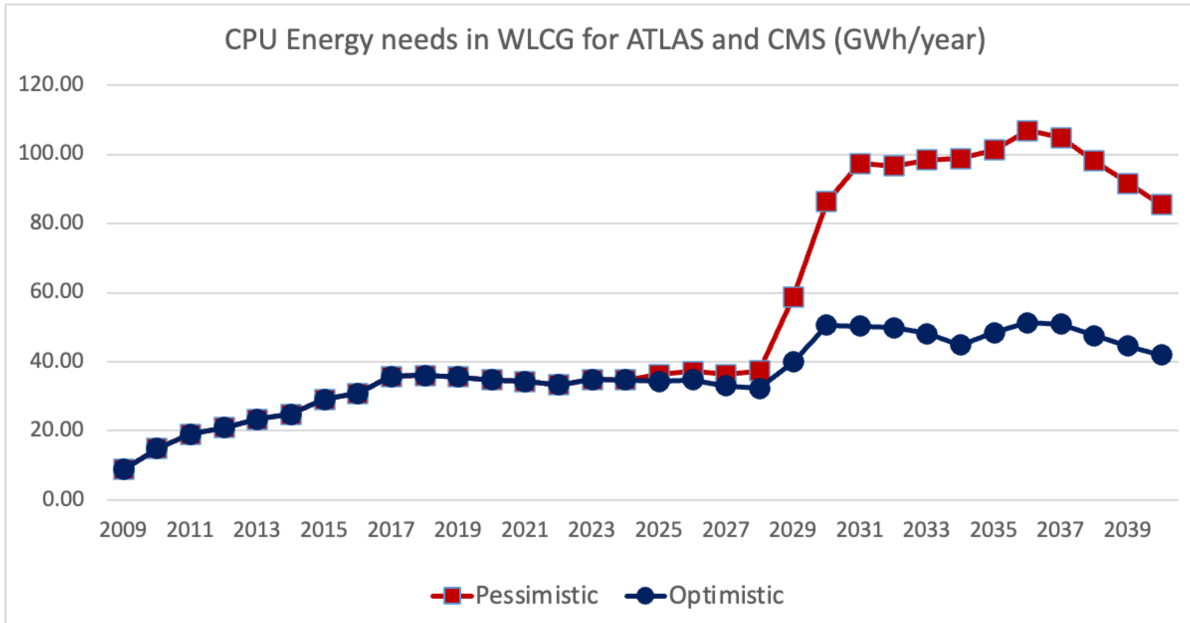


The peak of energy need happens in 2036 (start of Run-5): 400% higher than 2022 in the pessimistic scenario and 50% higher in the optimistic scenario

This assumes no improvement in the PUE of any of the facilities

# Key Performance Indicator for LHC

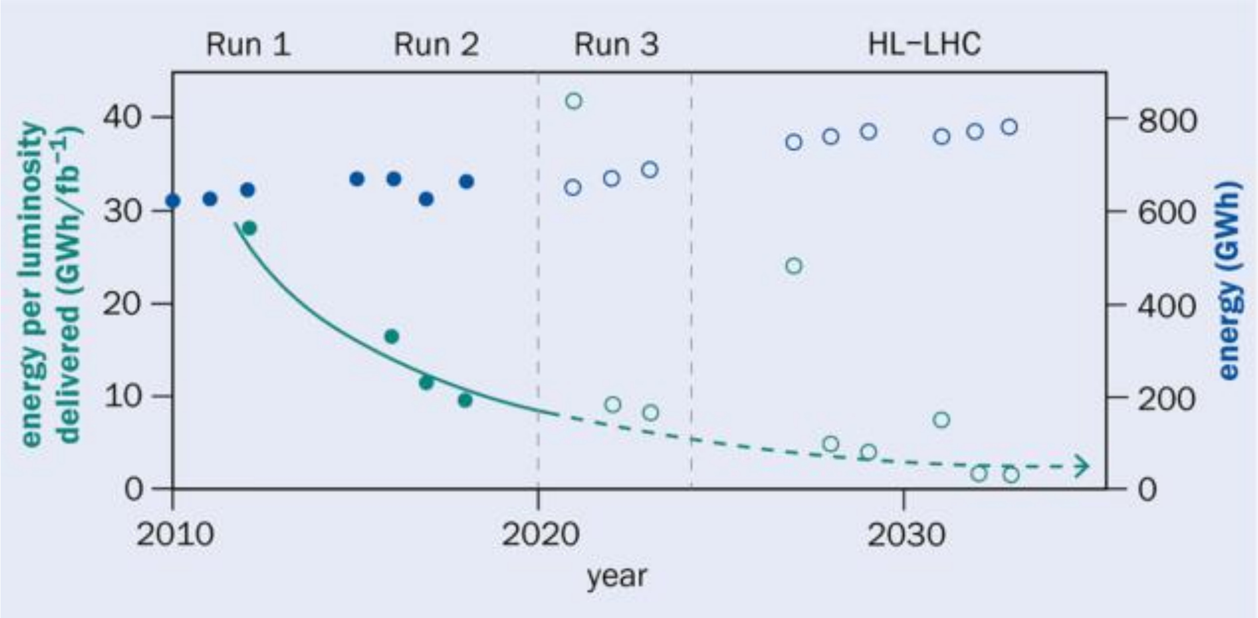


The GWh/fb$^{-1}$ metric has now been adopted by CERN as a key performance indicator (KPI) for the LHC. See this article for example

It represents the amount energy needed to **produce** a given amount of data

# $GWh/fb^{-1}$ for WLCG

In WLCG $GWh/fb^{-1}$ represent the energy needed to **analyse** the data

- look at each run rather than each year and assume that in the shutdown periods the processing of the previous run is completed

The scale on the left **(RED)** shows the energy and the scale on the right **(BLUE)** shows $GWh/fb^{-1}$ (log!)

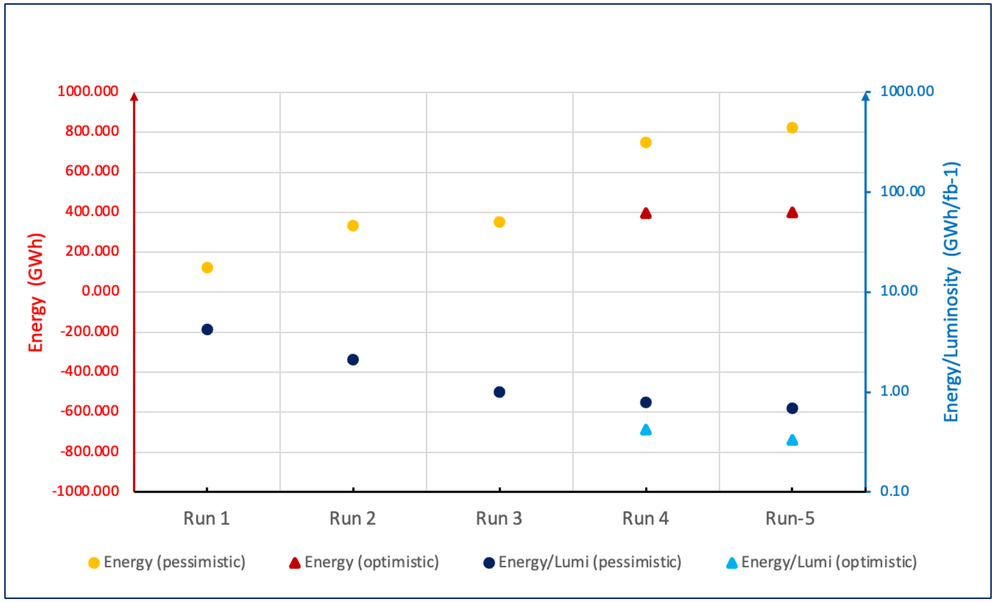


Energy needs in Run-4 and Run-5: +100% compared to Run-2 in the **pessimistic** scenario, only +10% in the **optimistic** scenario

$GWh/fb^{-1}$ decreases a factor 10 between Run-1 and Run-5.
In Run-5, $GWh/fb^{-1}$ in the **optimistic** scenario is half compared to the **pessimistic** scenario

# Power Usage Efficiency

The modernisation of the facilities generally reduces the PUE (overhead of cooling)

The effect of the PCC on the energy needs for ATLAS and CMS CPUs at CERN is on the right plot

The new CERN Prevessin Computing Center (PCC) PUE will be ~1.1

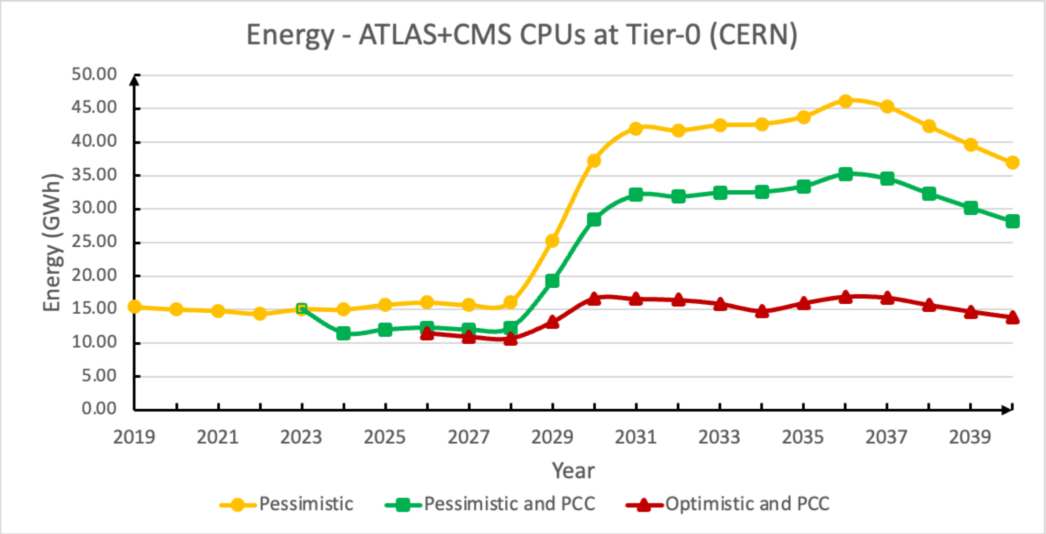


The introduction of the PCC reduces by 30% the CPU energy needs

The successful completion of the R&D program reduces the needs by another 50%

# Conclusions

- This study shows the trends and does not pretend to make predictions
- The energy needs in WLCG can be kept under control leveraging three pillars
    - ➢ The improvement in the hardware technologies. The impact depends on the hardware lifecycle strategy
    - ➢ The modernization of the facilities, going in the direction of more energy efficiency
    - ➢ The improvements in the software and computing models
- Each pillar is important, but the improvements in software and computing models are an area where everyone in the WLCG community can contribute
- In all scenarios, GWh/fb$^{-1}$ decreases over time: more "physics per kW"
- A lot more to study: non X86 CPU architectures? Hardware accelerators (e.g. GPUs?)